

Real-Time Dense Field Phase-to-Space Simulation of Imaging Through Atmospheric Turbulence

Nicholas Chimitt , Graduate Student Member, IEEE, Xingguang Zhang , Graduate Student Member, IEEE, Zhiyuan Mao , Graduate Student Member, IEEE, and Stanley H. Chan , Senior Member, IEEE

Abstract—Numerical simulation of atmospheric turbulence is one of the biggest bottlenecks in developing computational techniques for solving the inverse problem in long-range imaging. The classical split-step method is based upon numerical wave propagation which splits the propagation path into many segments and propagates every pixel in each segment individually via the Fresnel integral. This repeated evaluation becomes increasingly time-consuming for larger images. As a result, the split-step simulation is often done only on a sparse grid of points followed by an interpolation to the other pixels. Even so, the computation is expensive for real-time applications. In this article, we present a new simulation method that enables *real-time* processing over a *dense* grid of points. Building upon the recently developed multi-aperture model and the phase-to-space transform, we overcome the memory bottleneck in drawing random samples from the Zernike correlation tensor. We show that the cross-correlation of the Zernike modes has an insignificant contribution to the statistics of the random samples. By approximating these cross-correlation blocks in the Zernike tensor, we restore the homogeneity of the tensor which then enables Fourier-based random sampling. On a 512×512 image, the new simulator achieves 0.025 seconds per frame over a dense field. On a 3840×2160 image which would have taken 13 hours to simulate using the split-step method, the new simulator can run at approximately 60 seconds per frame.

Index Terms—Atmospheric turbulence, wave propagation, Zernike basis, Phase-to-Space Transform, Fourier optics.

I. INTRODUCTION

LIGHT propagating through the atmosphere suffers from distortions due to the random spatio-temporal fluctuations in the index of refraction. Over a long distance, these distortions will accumulate and degrade the image quality. The development of atmospheric turbulence mitigation algorithms has received a considerable amount of interest over the past few decades [1],

[2], [3], [4], [5], [6], [7], [8], [9]. Deep-learning-based techniques have recently been reported with some promising preliminary results [10], [11], [12], [13]. However, as in any inverse problem, the atmospheric turbulence problem requires a forward model that can accurately describe the image formation process. To this end, simulating the turbulent effect becomes a critical step toward the goal of designing algorithms, evaluating methods, and understanding the limitations of imaging systems.

For decades, simulating atmospheric turbulence is most accurately performed in the wave domain because there is no simple intensity domain model. The “gold standard” approach is the split-step propagation [5], [14], [15].¹ The idea is to split the propagation path into segments and model the phase distortion in each segment for every pixel individually. However, split-step propagation is not scalable. For a 256×256 sized image, the split-step simulator reported in [5] can evaluate a grid of size 64×64 where the remaining pixels are interpolated. The reported runtime was approximately 24.6 seconds per frame on a GPU. If the size of the image grows to 3840×2160 (4 K resolution) and the grid is *dense*, i.e. without interpolation, a rough estimate is about 13.7 hours for one image. To synthesize a training dataset containing 1000 of these sequences where each sequence has 100 frames, this will take 156 years. Recognizing the pressing need for an accurate and fast turbulence simulator, we present a method that enables turbulence simulation in real time.

A preview of our results is shown in Fig. 1. While split-step propagation is largely limited to a small grid of sparse points, the proposed method can directly generate a *dense field*. The run-time of the proposed method is approximately 0.025 seconds for a 512×512 image. For a high-definition (HD) image of size 3840×2160 , the runtime is approximately 60 seconds. As can be seen in Fig. 1, the new simulator allows us to zoom in to any region of the image while the turbulent effect is still globally correlated according to the theoretical statistics. To our knowledge, this is the first practical demonstration of an HD dense-field turbulence simulation documented in the literature.

The proposed approach, named the Dense Field Phase-to-Space (DF-P2S) simulation, is built upon the multi-aperture model by Chimitt and Chan [18], and the phase-to-space (P2S) transform by Mao et al. [19]. DF-P2S overcomes a fundamental

Manuscript received 29 June 2022; revised 12 October 2022; accepted 17 November 2022. Date of publication 2 December 2022; date of current version 6 January 2023. This work was supported in part by the Intelligence Advanced Research Projects Activity (IARPA) under Contract No. 2022-21102100004, and in part by the National Science Foundation under Grants CCSS-2030570 and IIS-2133032. 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 IARPA, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation therein. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Alin M. Achim. (Corresponding author: Nicholas Chimitt.) The authors are with the School of Electrical and Computer Engineering, Purdue University, West Lafayette, IN 47907 USA (e-mail: nchimitt@purdue.edu; zhan3275@purdue.edu; mao114@purdue.edu; stanchan@purdue.edu). Digital Object Identifier 10.1109/TCI.2022.3226293

¹There exist other modalities for the simulation of these effects that do not require computationally costly numerical wave propagation [16], [17], however, split-step is presently the most theoretically justifiable approach.




Fig. 1. [Top] This article presents a new turbulence simulation method that produces *dense-field* turbulence in real-time. For a 512×512 image, the classical split-step propagation takes 1 s to generate a 32×32 grid followed by interpolation of the field. The proposed method takes 0.03 seconds to generate the turbulence at dense-grid. [Bottom] Snapshot of a simulated turbulence image from an input image with a 4 K resolution (3840×2160 pixels).

limitation of [19] which is the size of the cross-correlation matrix (tensor). In [19], the cross-correlation tensor must be pre-computed, stored, and decomposed before running the simulator. This causes some computational overhead, however, the bigger issue is memory. The largest cross-correlation tensor that can be stored is for a spatial grid of size 32×32 using 36 basis coefficients. To simulate an image with a higher resolution, we need to interpolate the field, which limits the overall accuracy of such a simulation. Our solution to overcome this memory bottleneck is to maintain the homogeneity along the spatial dimensions of the correlation tensor and perform an approximation of the cross-correlation functions which otherwise restrict this behavior. This is based on a new observation that the exact form of the cross-correlation functions can be approximated without severely hurting the tensor statistics. As a result, we can employ Fourier-based techniques to draw *dense field* samples spatially at a low computational cost. This allows us to maintain a similar speed as P2S [19], yet gain an increase in statistical accuracy.

To summarize, this article offers two contributions:

- 1) *Real-time dense field turbulence simulation.* We report the first turbulence simulator that can simulate over a dense field and in real-time.
- 2) *New approximation to the cross-correlation function.* We show how certain off-diagonal blocks of the cross-correlation matrix used in [19] can be removed to utilize Fourier-based generation, hence enabling a significant resolution upscale.

II. OUTLINE OF GENERAL SIMULATION PRINCIPLES: BUILDING BLOCKS

To keep track of the notations, we use object plane coordinates $\mathbf{x} = (x, y)$ and image plane coordinates $\mathbf{u} = (u, v)$. For

a function defined across the aperture, we use the coordinates $\xi = (\xi, \eta)$ or its polar form $\rho = (\rho, \theta)$. We also use the polar coordinates $\mathbf{s} = (s, \phi)$ to denote the displacement in the Zernike space.

The turbulence effect is modeled by convolving a *spatially varying* point spread function (PSF) to a diffraction-limited clean image. In the case of incoherent light, which is the focus of this work, the observed image $I(\mathbf{x})$ is

$$I(\mathbf{x}) = h_{\mathbf{x}}(\mathbf{u}) \sim I_g(\mathbf{u}), \quad (1)$$

where $h_{\mathbf{x}}(\mathbf{u})$ is the PSF with the subscript to emphasize that it is spatially varying, and $I_g(\mathbf{u})$ is the distortion-free image. Note that the observed image is indexed by $\mathbf{x}$ whereas the PSF and ideal image are indexed by $\mathbf{u}$. This is to emphasize that after $h_{\mathbf{x}}(\mathbf{u})$ is convolved with $I_g(\mathbf{u})$, only the *center pixel* is used to construct $I(\mathbf{x})$.

The per-pixel PSF can be generated per the Fraunhofer diffraction equation with phase error [20]. Denoting $P(\xi)$ as the aperture function, $h_{\mathbf{x}}(\mathbf{u})$ is

$$h_{\mathbf{x}}(\mathbf{u}) = \text{Fourier} \left[P(\xi) e^{-j\varphi_{\mathbf{x}}(\xi)} \right], \quad (2)$$

withholding some constants that determine the size of the PSF according to the optical parameters. Here, $\varphi_{\mathbf{x}}(\xi)$ is the phase distortion function that varies over ξ for coordinate $\mathbf{x}$ in the image. Note that $\varphi_{\mathbf{x}_0}(\xi) \neq \varphi_{\mathbf{x}_1}(\xi)$ if $\mathbf{x}_0 \neq \mathbf{x}_1$.

Given that the PSF generation (2) and the image formation (1) is relatively standard, the central focus of a simulation approach then falls upon the generation of the random phase $\varphi_{\mathbf{x}}(\xi)$ in accordance with its theoretically given statistics. There are two main categories for generating $\varphi_{\mathbf{x}}$: (i) split-step propagation [5], [14], [15], which numerically propagates a wave through a random volume, and hence *modeling the medium*; (ii) collapsed phase-over-aperture [18], which generates the phase function directly at the aperture, which we refer to as the *multi-aperture model*. We illustrate the differences in these two approaches in Fig. 2.

A. Computational Bottleneck of Split-Step

Before we discuss the two building blocks of our simulator, it would be useful to highlight the limitations of the split-step method [5], [14], [15]. The split-step method directly mirrors the physical process by which light propagates. After a point propagates through the simulated medium, it arrives at the aperture of the imaging system with a phase component $\varphi(\xi)$. In the case of turbulence, the statistics of $\varphi(\xi)$ is determined by the structure function

$$D_{\varphi}(\xi, \xi^0) = E[(\varphi(\xi) - \varphi(\xi^0))^2]. \quad (3)$$

Assuming that the random function $\varphi(\xi)$ is homogeneous and isotropic, the structure function can be simplified to

$$D_{\varphi}(|\xi - \xi^0|) = 6.88(|\xi - \xi^0|/r_0)^{5/3}, \quad (4)$$

where r_0 is the Fried parameter [21].

To numerically generate the phase φ , the split-step method uses the Kolmogorov power spectrum density (PSD) [22] (or similarly Von Karman spectrum [5], etc) to generate discrete planes of turbulent distortions, referred to as phase screens [15].




Fig. 2. Here we show the main difference between classical approaches of modeling the medium and propagation directly, [5], [14], and our approach based on [18], [19]. In split-step, a subset of pixels in the image have point sources propagated through the medium which are then interpolated between to form the image at full resolution. In our approach, every single point has its own basis vector representation, meaning a phase function is generated for each pixel *without* interpolation. We emphasize that both approaches generate phase realizations for the image formation process, but it is the generation approach that differs.

This can be done directly with knowledge of the PSD and the Fourier transform based random field generation approach. The phase screens are placed along the propagation path as shown in Fig. 2. The size of the phase screens is generated larger than the input image in order to model the spatial correlations which are determined by the overlap of the phase screen along the path.

Since the split-step models the medium directly, it is regarded as the most theoretically justifiable approach. A comprehensive discussion of this model is provided in [15]. However, the main limitation of split-step is its computational requirements. For every point source, we need to perform multiple fast Fourier transforms (FFTs). A simulation with M phase screens results in $M(W \times H)$ 2D FFTs for an $W \times H$ image, with each FFT being at the size of the image. Scaling the process of generating a large dataset is nearly impossible.

B. Building Block 1: Multi-Aperture Model

Recognizing the speed limit of the split-step simulation, Chimmitt and Chan proposed a new concept in [18] which they named the *multi-aperture model*. The idea is to skip the propagation by going to the statistical description of the resultant phase $\varphi_x(\xi)$ using the Zernike representation first proposed by Noll in 1976 [23]. The idea is to define a radius R and a vector ρ such that $R\rho$ is the polar coordinate representation of ξ . Then, the phase φ_x can be represented as

$$\varphi_x(R\rho) = \sum_{j=1}^{\infty} a_{x,j} Z_j(\rho) \approx \sum_{j=1}^N a_{x,j} Z_j(\rho), \quad (5)$$

with $Z_j(\rho)$ as the j th Zernike function and $a_{x,j}$ as its respective coefficient. We emphasize that the phase φ_x is location specific, i.e., the phase function at $\mathbf{x}$ is different from the phase at $\mathbf{x}^0$ if $\mathbf{x} \neq \mathbf{x}^0$. Therefore, the Zernike coefficients $a_{x,j}$ are different from $a_{x^0,j}$. However, the basis function

$Z_j(\rho)$ is shared for all locations. In this article, we set $N = 36$, though this can easily be increased at the cost of some speed.

The coefficients $a_{x,j}$ are zero-mean Gaussian, and have a correlation matrix (technically, a tensor stored in the matrix form) which we denote the $(\mathbf{x}, \mathbf{x}^0, i, j)$ th element of the matrix A as

$$[A]_{\mathbf{x}, \mathbf{x}^0, i, j} = E[a_{x,i} a_{x^0,j}]. \quad (6)$$

This notation stresses that the matrix A is location dependent on the pair $\mathbf{x}$ and $\mathbf{x}^0$. In the special case when the correlation matrix is spatially invariance, we write $[A]_{\mathbf{x}, \mathbf{x}^0, i, j}$ as $[A]_{\mathbf{x}-\mathbf{x}^0, i, j}$. If we further assume that $\mathbf{x}^0 = \mathbf{x}$, the correlation matrix becomes $[A]_{0, i, j} = E[a_{x,i} a_{x,j}]$, which is equivalent to the covariance matrix given by Noll [23].

Based on the decomposition in (5), we can simulate the turbulence in three steps [18]: (i) Collapse the screens; (ii) Draw spatially correlated Zernike vectors; (iii) Draw inter-modally correlated Zernike vector. Therefore, we have effectively converted the split-step propagation into a sampling problem of the Zernike coefficients.

This multi-aperture model provides roughly a $6\times$ increase in speed over split-step. However, the major drawback is that the spatial correlation for the higher order ($a_{x,i}$ where $i \geq 3$) Zernike terms are not computationally feasible. The generation of the spatial and intermodal correlated vectors requires a large covariance matrix, which exponentially increases in size with the size of the image. Therefore, only tilts are reasonably implementable for this method. Furthermore, the image formation process is the same as split-step, requiring us to generate one PSF for one pixel repeatedly over the entire field of view plus convolutions to produce the final image.

C. Building Block 2: Phase-to-Space

While the multi-aperture approach reduced the computation by collapsing the propagation into a sampling problem, the PSFs still need to be generated according to (2) and the spatially varying convolution still needs to be performed via (1). To address these computational limitations, Mao et al. [19] introduced the phase-to-space (P2S) transform that reformulates the PSF via a *spatial* decomposition:

$$h_x(u) = \sum_{m=1}^M \beta_{x,m} \phi_m(u), \quad (7)$$

where $\phi_m(u)$ is a *spatial* basis function for the PSF. In [19], these basis functions are learned by running the principle component analysis on a database on PSFs. The parameter $\beta_{x,m}$ denotes the coefficient associated with each basis $\phi_m(u)$. What (7) says is that the spatially varying convolution in (1) can now be replaced by a set of spatially invariant convolutions using $\phi_m(u)$. To form the image, we simply compute the per-pixel coefficient $\beta_{x,m}$ and form the weighted average:

$$I(x) = h_x(u) \sim I_g(u) = \sum_{m=1}^M \beta_{x,m} \underbrace{\phi_m(u) \sim I_g(u)}_{\text{spatially invariant convolution}}, \quad (8)$$

where we should emphasize the speed improvement due to the spatially invariant convolution. However, the biggest challenge here is that there is no analytic expression that allows us to translate the Zernike basis $Z_j(\rho)$ to the spatial basis $\phi_m(u)$. The solution proposed in Mao et al. [19] is to train a small neural network to convert the coefficients from the Zernike coefficients $\mathbf{a}_x = [a_{x,1}, a_{x,2}, \dots, a_{x,N}]$ (which are sampled from $\mathbf{A}$) to the PSF coefficients $\beta_x = [\beta_{x,1}, \beta_{x,2}, \dots, \beta_{x,M}]$.

The phase-to-space transform eliminates the necessity of taking any FFT for the formation of a PSF, dramatically speeding up the generation process by 1000 $\times$ compared to split-step. However, the limitation of the phase-to-space transform is the memory requirement. For an image of size $W \times H$, the generation using 36 Zernike coefficients requires the construction of a matrix that is $36HW \times 36HW$. What is worse is that we can not leverage the homogeneity property (analogous to the wide-sense stationary in random processes) and so we cannot use Fourier transforms to decompose the correlation matrix and draw random samples. The workaround solution in [19] is to focus on a small grid of points $H/d \times W/d$ where d is the sampling ratio, followed by decomposing the $36HW/d^2 \times 36HW/d^2$ matrix and interpolating between these points. The grid sizes, in most of the demonstrations, are in the range of 16×16 to 32×32 due to memory limitations. In this article, we resolve this memory issue by completely eliminating the interpolation by approximation, thus enabling a truly dense field turbulence generation.

D. Other Simulation Approaches

While split-step is regarded as the most accurate simulation method, other approaches do exist. The most closely related one is the brightness function simulation [16], [24]. These methods

are based on the solution to ray optics simulation through a series of phase screens. Their comparisons with the split-step can be found in [25]. Additionally, empirically-driven methods such as Hunt et al. [17] model the atmospheric distortion basis representations. There are also low-order approximation methods such as [6], [26], [27], which can generate visually plausible images but fail to satisfy the turbulence statistics.

III. DENSE FIELD PHASE-TO-SPACE SIMULATION (DF-P2S SIMULATION)

In this section, we outline the proposed DF-P2S simulation. DF-P2S uses the same backbone as the original P2S in [19] but offers a dense field simulation not possible with P2S.

The overall simulation can be broken down into two steps as shown in Fig. 3:

- 1) *Generate correlated Zernike coefficients.* The goal of this step is to use Building Block 1: Multi-Aperture Model to draw random but correlated Zernike coefficients according to the correlation matrix.
- 2) *Apply the Zernike coefficients using the P2S transform.* This step uses Building Block 2: Phase-to-Space Transform to convert the correlated Zernike coefficients to the spatial basis coefficients. Consequently, the resulting image can be formed by (8).

The biggest innovation of DF-P2S is a new way to accomplish the first step. We ask the question: how can we efficiently draw the Zernike coefficients from $\mathbf{A}$ without suffering from the memory bottleneck? To answer this question, it is important to investigate the structure of the correlation matrix $\mathbf{A}$.

A. Intuition of the New Method

The idea of drawing correlated Zernike coefficients according to a correlation matrix is not difficult if the number of grid points is small. Consider a multivariate Gaussian random variable with mean $\mathbf{E}[\mathbf{y}] = \mathbf{0}$ and the correlation matrix $\Sigma = \mathbf{E}[\mathbf{y}\mathbf{y}^T]$. To draw a random vector $\mathbf{y}$ from this distribution, we can decompose the correlation matrix $\Sigma = \mathbf{U}\mathbf{S}\mathbf{U}^T$ via the eigen-decomposition, and define $\Sigma^{\frac{1}{2}} = \mathbf{U}\mathbf{S}^{\frac{1}{2}}\mathbf{U}^T$. Then, starting with a white noise vector $\mathbf{e} \sim \text{Gaussian}(\mathbf{0}, \mathbf{I})$, the transformed vector $\mathbf{y} = \Sigma^{\frac{1}{2}}\mathbf{e}$ will satisfy the desired property that $\mathbf{E}[\mathbf{y}] = \mathbf{0}$ and $\mathbf{E}[\mathbf{y}\mathbf{y}^T] = \Sigma$.

When the size of the grid is small so that the dimension of the correlation matrix Σ is small, the matrix $\Sigma^{\frac{1}{2}}$ can be generated using standard numerical techniques such as the Cholesky factorization. However, for a large grid of points, storing the matrix Σ and running the factorization would become infeasible. One exception is that if $\mathbf{y}$ is homogeneous (wide-sense stationary), then the correlation matrix Σ is *circulant* and so the eigen-decomposition is equivalent to the Fourier transform. In this case, generating the random vector $\mathbf{y}$ can be implemented via **

$$\mathbf{y} = \mathbf{U}\mathbf{S}^{\frac{1}{2}}\mathbf{U}^T\mathbf{e} = \mathbf{F}^{-1}(\mathbf{S}^{\frac{1}{2}}\mathbf{F}(\mathbf{e})), \quad (9)$$

where $\mathbf{F}$ denotes the discrete-time Fourier transform, and the diagonal matrix $\mathbf{S}$ is the Fourier spectrum of one row (or column) of the correlation matrix Σ .




Fig. 3. Here we present the overall DF-P2S simulation pipeline. We first generate a Zernike vector for *every* point in the image, a feat only possible through our simulation technique. Next, we feed these Zernike vectors parallelly into the P2S transform. The P2S transform maps the phase coefficients to PSF coefficients, which are then used to implement a fast version of a spatially varying convolution.




Fig. 4. The Zernike Space is a representation describing the phase decomposition by its basis functions. For every point on an object, or analogously every pixel in an image, there is a Zernike vector that describes the distortion across the aperture. This motivates us to define the Zernike space as a tensor.

The significance of the homogeneity is that it allows us to speed up the sampling process by performing all computations in the Fourier space. In addition, the memory bottleneck is resolved because we do not need to construct the full correlation matrix Σ and run the Cholesky factorization. The question is: For the Zernike correlation matrix we are considering in this article, does it have any kind of wide sense stationarity? If not, can we approximate it using something that has such a property?

B. Structure of the Correlation Matrix

In this subsection, we take a closer look at the structure of the Zernike correlation matrix A . As defined in (6), the $(\mathbf{x}, \mathbf{x}^0, i, j)$ th entry of the correlation matrix at the location-pair $(\mathbf{x}, \mathbf{x}^0)$ and the Zernike mode-pair (i, j) is $[A]_{\mathbf{x}, \mathbf{x}^0, i, j} = E[a_{\mathbf{x}, i} a_{\mathbf{x}^0, j}]$. Thus, A is a four-dimensional tensor, with two dimensions allocated to the pair of spatial coordinates (organized through a column-wise stack), and two dimensions allocated to the pair of Zernike modes. One way to visualize the Zernike space is to consider the illustration shown in Fig. 4. For a fixed pixel location $\mathbf{x}$, there

is a vector $\mathbf{a}_{\mathbf{x}} = [a_{\mathbf{x},1}, \dots, a_{\mathbf{x},N}]$ where $n = 1, \dots, N$ denotes the Zernike mode index. As we move to another pixel location $\mathbf{x}^0$, the vector becomes $\mathbf{a}_{\mathbf{x}^0} = [a_{\mathbf{x}^0,1}, \dots, a_{\mathbf{x}^0,N}]$.

For a fixed Zernike mode pair (i, j) , the correlation is limited to the spatial axis. This will give us the (i, j) th slice of the four-dimensional tensor

$$A_{i,j} = \begin{bmatrix} [A]_{\mathbf{x}_1, \mathbf{x}_1, i, j} & \dots & [A]_{\mathbf{x}_p, \mathbf{x}_1, i, j} \\ \vdots & \ddots & \vdots \\ [A]_{\mathbf{x}_p, \mathbf{x}_1, i, j} & \dots & [A]_{\mathbf{x}_p, \mathbf{x}_p, i, j} \end{bmatrix}, \quad (2)$$

where $\mathbf{x}_1, \dots, \mathbf{x}_p$ are the p coordinates in the grid. Assuming that the Zernike coefficients across the field of view is homogeneous [18], the $(\mathbf{x}, \mathbf{x}^0)$ th element $[A]_{i,j}(\mathbf{x}, \mathbf{x}^0)$ will be the function of $\mathbf{x} - \mathbf{x}^0$ instead of the absolute positions $(\mathbf{x}, \mathbf{x}^0)$. In this case, we can write $[A]_{i,j}(\mathbf{x}, \mathbf{x}^0)$ as $[A]_{i,j}(\mathbf{s})$ where $\mathbf{s} = (\mathbf{x} - \mathbf{x}^0)/D$ with D being the aperture diameter². For a grid of points $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_p$, the matrix $A_{i,j}$ takes the form

$$A_{i,j} = \begin{bmatrix} [A]_{\mathbf{s}_0, i, j} & [A]_{\mathbf{s}_1, i, j} & \dots & [A]_{\mathbf{s}_p, i, j} \\ [A]_{\mathbf{s}_1, i, j} & [A]_{\mathbf{s}_0, i, j} & \dots & [A]_{\mathbf{s}_{p-1}, i, j} \\ \vdots & \vdots & \ddots & \vdots \\ [A]_{\mathbf{s}_p, i, j} & [A]_{\mathbf{s}_{p-1}, i, j} & \dots & [A]_{\mathbf{s}_0, i, j} \end{bmatrix}. \quad (3)$$

Because of homogeneity, $A_{i,j}$ can be decomposed via the Fourier transform. This “slice” of the tensor is representative of a single Zernike coefficient field.

²The necessity to normalize the coordinate using D comes from the fact that the geometry of the optical system has a certain impact to the Zernike coefficients. In particular, two aperture diameters D_1 and D_2 will lead to two different Zernike space because the imaging systems are viewing through different effective slices as a function of aperture size. Therefore, we define a standardized unitless vector corresponding to correlation length in accordance with [18] to be $\mathbf{s} = \frac{\mathbf{L}(\boldsymbol{\theta} - \boldsymbol{\theta}^0)}{D} = \frac{\mathbf{x} - \mathbf{x}^0}{D}$, where $\boldsymbol{\theta}, \boldsymbol{\theta}^0$ are two vectors pointing from the center of the imaging system to points $\mathbf{x}, \mathbf{x}^0$, respectively.

For a fixed location pair (x, x^0) , we can obtain another slice of the correlation matrix

$$A_{x, x^0} = A_s = \begin{bmatrix} [A]_{s,1,1} & \dots & [A]_{s,1,N} \\ \vdots & \ddots & \vdots \\ [A]_{s,N,1} & \dots & [A]_{s,N,N} \end{bmatrix}.$$

The special case where $s = 0$, which gives matrix A_0 is exactly the Noll matrix [23] that specifies the correlation between the Zernike modes at a single pixel location x . In general, the matrix A_s does not have the circulant structure and cannot be decomposed via the Fourier transform.

The prior notation allows us to easily identify the homogeneous and the non-homogeneous components of the tensor. The entire Zernike correlation tensor A can be written as

$$A = \begin{bmatrix} A_{1,1} & \dots & A_{1,N} & A_{s_0} & \dots & A_{s_p} \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ A_{N,1} & \dots & A_{N,N} & A_{s_p} & \dots & A_{s_0} \end{bmatrix},$$

The entire tensor structure is non-homogeneous. Thus, it is not possible to draw samples via the FFT method, which highlights a fundamental limitation of [18] and [19].

C. Generation of the Random Zernike Fields

Due to the non-homogeneity of A , we propose an approximate tensor $\mathbf{A}$ which we claim captures a majority of the statistical behavior while having the property of homogeneity. Reserving the accuracy of this approximation for Section III-D, we propose the following method for the generation of the Zernike fields:

- 1) Generate $i = \{1, 2, \dots, N\}$ unit-variance, spatially correlated random fields according to $A_{i,i}$. Note this generation uses only *autocovariance* functions. At this stage, our N fields are independent, thus utilizing the FFT-based method based on their homogeneity.
- 2) Perform a point-wise mixing of the random fields according to the Noll matrix A_0 . This mixing is done pixel-wise (across the coefficient index dimension) per pixel.

This generation process will give us random fields which are in accordance with the autocovariance functions, however, for the cross-covariance terms there will exist some deviation. This is most simply presented by the form of the covariance structure of the tensor,

$$\mathbf{A} = L \begin{bmatrix} A_{1,1} & 0 & \dots & 0 \\ 0 & A_{2,2} & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & A_{N,N} \end{bmatrix} L^T, \quad (10)$$

where $LL^T = A_0$. We note the resultant matrix is no longer diagonal. Therefore, the off-diagonal entries of $\mathbf{A}$ differ from A , which will be the focus of our numerical analysis.

We provide a visualization of the resulting covariance matrix in Fig. 5 using a simplistic example in the case of an 8×8 image with 3 coefficient fields. Initially, white noise is generated for an $8 \times 8 \times 3$ random volume, which is then spatially correlated for




Fig. 5. We give a visual representation of the covariance structure as it changes through the generation process. First, white noise has a covariance structure applied to it, resulting in spatially correlated, yet independent random fields. Next, along the index axis, a mixing matrix is applied via Cholesky decomposition. Finally, the resulting fields are correlated both spatially and along the index dimension and have the covariance structure shown.

each slice according to its autocovariance function. The resulting covariance matrix has a considerable amount of zero-entries corresponding to the other 8×8 random fields. After this, the fields are mixed according to the 3×3 covariance matrix along the final index axis, after which the covariance matrix becomes much denser. This same principle is extended to the case of an $W \times H$ image with N coefficient fields.

D. The Impact of Cross-Correlation Functions

The core innovation in this work is the ability to utilize property of homogeneity for individual fields $a_{s,i}$ (or equivalently, $\mathbf{a}_{s,i}$). This allows us to quickly draw samples using FFT-based generation followed by a mixing matrix. To justify this approximation, we begin with the following considerations:

- 1) The correlation functions $A_{i,j}$ for $i \neq j$ sharply decay and approximately vanish for $s > 4$;
- 2) The off-diagonal terms for the matrix $A_{x,x}$ (corresponding to indices $i \neq j$) contain a small proportion of the energy.

Together, these two properties suggest that the overall energy lost by the removal of the cross-correlation functions will be negligible. Intuitively, the reason for this lies within the sharp decay of the correlation functions (1), which are already significantly smaller than their counterparts (2).

With this observation, we turn to the numerical analysis performed to more concretely justify this approximation. We wish to measure the energy that will be contained within our approximation and compare it with the energy contained if the entire cross-correlation functions were to be retained. To do so, we integrate outward in the Zernike space and quantify the deviation of $\mathbf{A}$ from A as a function of s . While the correlation functions $A_{i,j}$ have angular dependencies, we integrate over the angular components for ease of presenting our comparison. We note that along any particular direction individually, the results do not vary considerably. Mathematically, we write the energy contained with all functions considered via the Frobenius norm as

$$E(s) = \int_0^{2\pi} \int_0^s ds^0 \text{Tr} A_{s^0}^T A_{s^0}, \quad (11)$$




Fig. 6. The results from our numerical evaluation of the three proposed energy functions. The axis for $E(s)$ and $\tilde{E}(s)$ are on the right, while the axis for residual difference energy function $E^{(-)}(s)$ is on the left. We note the apparent match between the two individual energy functions. More importantly, there is a comparably small magnitude for the difference function, suggesting a reasonable match.




Fig. 7. A heat chamber is built to collect real turbulence data. The heat chamber consists of three blocks where the total optical path length is 20 meters. The temperature along the path can be controlled by the number of heat lamps that are switched on. At one end of the heat chamber, we put a monitor displaying the target pattern. An 800 mm lens camera is placed at the other end of the setup to capture the image.

where we've integrated spatially outwards up to s , as we intend show the accuracy of the approximation as the size of our field in the Zernike space grows. As s becomes large with respect to the spread of the function, $E(s)$ will converge. The energy within our approximation can be written similarly as

$$\tilde{E}(s) = \int_0^s \int_0^{2\pi} ds^0 \text{Tr} \left[\mathbf{A}_{s^0}^T \mathbf{A}_{s^0} \right] . \quad (12)$$

In addition to comparing the energies separately, we may also consider measuring the difference between them. To this end, we propose to measure

$$E^{(-)}(s) = \int_0^s \int_0^{2\pi} ds^0 \text{Tr} \left[\mathbf{A}_{s^0}^T \mathbf{A}_{s^0} - \mathbf{A}_{s^0}^T \mathbf{A}_{s^0} \right] , \quad (13)$$

where we are only interested in the magnitude of the residual, therefore measuring its absolute value.

These three functions will help to provide insight into the accuracy of our approximation. Ideally, we want $E(s)$ and $\tilde{E}(s)$ to match. However, this alone is not enough to claim accuracy, as this is only a way of measuring the total joint behavior. We additionally then propose to use $E^{(-)}(s)$ to measure the total index-wise residuals. An optimal result for $E^{(-)}(s)$ would be for it to vanish at every point. We present the results of this numerical analysis in Fig. 6 which we note is cumulative, therefore reaching

a steady state means no additional errors will be incurred. We also choose to not include the first three Zernike coefficients (piston and x, y -tilts) for this plot, as these terms dominate the plot significantly in magnitude, though the conclusion of the analysis is unchanged. With this, we present the following observations:

- 1) At $s = 0$ there is no loss in energy/statistical accuracy (e.g. for a single point we are perfectly accurate);
- 2) A majority of the errors are for small s , which is expected as the correlation functions have not yet vanished at this point;
- 3) For separations $s > 4$, there is no additional loss in energy and remains at steady state. The energy of the difference $E^{(-)}(s)$ is two orders of magnitude smaller than the total energy, suggesting its overall accuracy.

These three observations suggest that our approximation retains a majority of the field's statistical behavior, while allowing for considerable speed-up.

E. Spatio-Temporally Correlated Fields

The extension of the proposed ideas to spatio-temporally correlated fields is possible through the adoption of Taylor's frozen flow hypothesis [28]. Under this hypothesis, a displacement in time t may be written as a spatial displacement y through the relationship $y = vt$, with v as the mean transverse velocity of the turbulent medium. Taylor's frozen flow hypothesis allows us to extend the results by considering $z = x + vt$. This will give us a spatio-temporal tensor A_{z, z^0} of analogous structure to the previously described A_{x, x^0} .

A downside of the frozen flow hypothesis is that we need to pre-compute and store multiple three-dimensional random fields. The workaround solution is to enforce the temporal correlation via an auto-regressive process. Therefore, the random seed e used to generate the Zernike fields as in (9) are updated according to $e_t = \gamma e_{t-1} + \sqrt{1 - \gamma^2} n$. Here we've introduced subscript t to indicate a temporal index along with n as a zero-mean, unit-variance Gaussian vector and $\gamma \in [0, 1]$. While the theoretical justification of such a model is limited, the speed of simulation is marginally impacted. We observe the visual effects under the auto-regressive approximation to be suitable in practice.

IV. EXPERIMENTS

In this section, we discuss comparisons between the proposed DF-P2S scheme and split-step. We view the P2S method [19] as a fast way to generate our PSFs, which may be optionally replaced by (2) at the cost of speed. Therefore, we do not compare directly to this method. We begin with visual comparisons with split-step, which we emphasize are helpful as a reality check, though are not as useful as a statistical comparison. As a result, we take care in comparing our simulation results to the theoretically predicted statistics, which we find to be more informative. Finally, we compare the runtime and resolution of the DF-P2S approach compared against other methods.




Fig. 8. Comparison of our simulation approach against our observed heat chamber data and split-step. We observe a reasonable match for both cases, with our approach having the capability to easily operate on high-resolution images. Taking the first row as our example, DF-P2S can operate under a second while split-step would take an estimated time of 20 minutes. Our estimate is based on the grid size used by the split-step.

A. Visual Comparisons

Visual comparisons require a ground truth image to simulate, as well as turbulent pairs. For this purpose, we use a heat chamber, where we have a series of heating lamps along a 20-meter path, as shown in Fig. 7. This allows us to easily use a digital image displayed on a monitor at the end of the path.

We compare our method with our heat chamber data, as well as split-step [5], [15]. We present a few examples of our visual comparisons in Fig. 8. Visually, we can observe a reasonable match between both generation methods, though we have up-scaled the distortions of split-step according to [5] to match the image resolution.

The direct visual comparison is a useful reality check; it serves to show that a method can produce visually realistic images. However, there is no easily quantifiable approach to validating a simulator via visual comparison. This motivates us to emphasize other modes of validation.

B. Statistical Validations

In analyzing the simulation approach, adherence to the desired statistical behavior is a key metric by which we judge the quality of our simulation approach. The validation here is divided into two categories: (1) aperture statistics and (2) spatial statistics. For statistics on the aperture, the key comparison is our generated statistics plotted against the theoretical structure function. With respect to spatial statistics, our comparison is with the known tilt statistics; there is a limitation on the known behavior of the spatial statistics with respect to the blur without the approximation in [18], so no direct comparison is possible.

Aperture Statistics: For the evaluation of aperture statistics, the most important function to match is that of the structure function given by (4). By matching this function, as we show




Fig. 9. A comparison between the theoretical structure function (solid curves) at different distortion levels vs. our generated statistics (dotted lines). We observe a reasonable match across multiple levels of distortions.

in Fig. 9, we are able to match any statistical value that can be written in terms of the structure function, such as the Fried parameter. We show a good match to the structure function at varying levels of turbulent distortions.

We also perform another experiment on the temporal averages of the distortions, the short exposure (SE) and long exposure (LE) optical transfer functions (OTFs). The OTF in general is defined as

$$H(\xi) = P(\xi)e^{-j\varphi(\xi)} \sim P(-\xi)e^{j\varphi(-\xi)}, \quad (14)$$

which is the autocorrelation operation of the overall pupil function and phase distortion. The LE OTF is then given by a temporal average over realizations of individual OTFs. Mathematically, the LE OTF is given as

$$H_{LE}(\xi) = E \left[\frac{1}{h} \sum_{i=1}^h P(\xi)e^{-j\varphi(\xi)} \right] \sim P(-\xi)e^{j\varphi(-\xi)}. \quad (15)$$




Fig. 10. Comparing the long and short exposure OTFs at different distortion levels with the theoretical curves (solid lines) and our statistics (dotted lines).

The SE OTF is similarly defined using the temporal average, however, this measure uses a “tilt-corrected” phase function ϕ given by

$$\phi(\xi) = \phi(\xi) - \alpha^T \xi, \quad (16)$$

where $\alpha^T \mathbf{f}$ is the plane of best fit, effectively removing the tilt. The remaining phase distortion ϕ will then only describe the high-order distortions. This differs from the LE OTF, which includes both blur and tilt/shifting. The SE is then described as

$$H_{SE}(\xi) = E \left[P(\xi) e^{-j\phi(\xi)} \right] \sim P(-\xi) e^{j\phi(-\xi)}. \quad (17)$$

The LE and SE OTFs have analytic expressions [28], which we compare with the results of our simulation in Fig. 10 which also includes the diffraction OTF. We observe a wide variety of D/r_0 levels that we perform with sufficient accuracy. Some slight deviation may be improved by subharmonic methods such as those described by Schmidt [15].

Spatial Statistics: For further validation, we also make a comparison with spatial statistics, of which we choose the tilt correlation and differential tilt variance. The tilt correlation measures the similarity between tilts as a function of position, therefore monotonically decreasing, while the differential tilt variance measures the difference between the tilt, and as a result, monotonically increases. The expression for tilt correlation has been carried out by Basu (now Bose-Pillai) et al. [29], which is what we use for our comparison. However, as the expression is analytically cumbersome, we state the definition as given by




Fig. 11. Comparing theoretical (solid and dashed) vs. empirical spatial tilt statistics (dotted). We note that this is for one configuration of camera and geometry, though among our tests we can match to this level of accuracy in a wide variety of scenarios. The limitation of this match is inherent to the approximation of [18].

Fried [30]

$$E[\alpha(0)\alpha(\vartheta)] \approx \int \int d\mathbf{r} d\mathbf{r}^0 W(\mathbf{r}) W(\mathbf{r}^0) \mathbf{r} \cdot \mathbf{r}^0 D_\varphi(\mathbf{r} - \mathbf{r}^0, \vartheta) \quad (18)$$

with a somewhat modified structure function from [29], [30]. We present a comparison with these theoretical statistics with our generated statistics in Fig. 11. In addition to tilt correlations, we may also calculate differential tilt variance (DTV) from the tilt correlation via

$$E[(\alpha(0) - \alpha(\vartheta))^2] = 2 E[\alpha^2(0)] - E[\alpha(0)\alpha(\vartheta)], \quad (19)$$

which quantifies the overall variance in the distortions as a function of position of the Zernike tilt terms. We again observe a match across a wide variety of D/r_0 values. The deviation of these results is a result of a Taylor series performed in the analysis of [18] upon the structure function. The tilt correlation and DTV changes with different combinations of imaging geometries and camera parameters, though in our tests we can typically match to this degree of accuracy assuming a constant C_ρ . For more details regarding this approximation, we would refer the reader to [18].

C. Speed and Resolution Comparisons

For our runtime comparison, we consider the generation contained in Fig. 2. That is, we want to compare the time to generate only the phase distortions, not the application of the point spread functions. We choose not to include this as this can be replaced in either approach with either the analytic formula (2) or the P2S network. We feel this is the most fair comparison of an atmospheric turbulent simulation tool, as the core goal is to produce the phase distortions. Therefore, Fig. 12(a) reflects the time to generate the turbulent phase distortions at varying resolutions between our approach and split step. We then show additional resolution for our simulation in Fig. 12(b) to which split-step is not scalable. Notably, our simulator can generate an image at an image resolution of 512×512 in the same time as approximately a 10×10 image with split-step, which would then typically be upsampled $4\times$, resulting in a 40×40 image.




Fig. 12. A comparison of the runtimes of split-step [5] and ours. (a) Low resolution. (b) High resolution. Note that split-step cannot handle any high resolution.




Fig. 13. Graphics User Interface (GUI). Our GUI can be connected to a standard camera and generate dense-field turbulence effects at 7 fps for a 512×512 image.

D. GUI for Real-Time Simulation

As a result of our simulation being computationally efficient, a real-time generation of turbulent images can be developed. We show a real-time demonstration using a camera and GPU in which we can display the distorted video stream, pixel displacements, and an 8×8 grid of sub-sampled PSFs in Fig. 13. Given

some interfacing with the camera and calculation of statistics in real-time, this is not performed at the same speed as reported in Fig. 12. However, we can achieve approximately 8 frames per second (FPS) for a 512×512 image on a nVidia GeForce GTX 1080 Ti GPU with the GUI and additional statistical information displayed. To our knowledge, this is the first simulator in the literature that has the capability of performing a simulation with this level of accuracy at this speed, making this the first real-time demo of turbulent imaging simulation. Additionally, we can modify the imaging geometry and turbulence strength in real-time, for which we have tunable knobs in our GUI.

V. CONCLUSION

In this article, we've proposed the DF-P2S simulator for imaging through atmospheric turbulence. The DF-P2S approach can generate real-time turbulent realizations for standard 512×512 images and is capable of generating large-scale datasets for training purposes. In addition to its speed, it maintains competitive accuracy against with the traditional split-step simulation. This is leveraged by the approximation on the Zernike space covariance tensor, which we justify via numerical experiments. Statistically, the DF-P2S simulation can generate distortions that match with their theoretically predicted curves. We further demonstrate with the speed of our simulator, we can perform these actions in nearly real-time.


REFERENCES

- [1] N. Anantrasirichai, A. Achim, N. G. Kingsbury, and D. R. Bull, "Atmospheric turbulence mitigation using complex wavelet-based fusion," *IEEE Trans. Image Process.*, vol. 22, no. 6, pp. 2398–2408, Jun. 2013.
- [2] M. C. Roggemann, C. A. Stoudt, and B. M. Welsh, "Image-spectrum signal-to-noise-ratio improvements by statistical frame selection for adaptive-optics imaging through atmospheric turbulence," *Opt. Eng.*, vol. 33, no. 10, pp. 3254–3264, 1994.
- [3] N. Anantrasirichai, A. Achim, and D. Bull, "Atmospheric turbulence mitigation for sequences with moving objects using recursive image fusion," in *Proc. IEEE Int. Conf. Image Process.*, 2018, pp. 2895–2899.
- [4] G. Gilles and S. Osher, "Wavelet burst accumulation for turbulence mitigation," *J. Electron. Imag.*, vol. 25, no. 3, 2016, Art. no. 033003.
- [5] R. C. Hardie, J. D. Power, D. A. LeMaster, D. R. Droege, S. Gladysz, and S. Bose-Pillai, "Simulation of anisoplanatic imaging through optical turbulence using numerical wave propagation with new validation analysis," *Opt. Eng.*, vol. 56, no. 7, 2017, Art. no. 071502.
- [6] X. Zhu and P. Milanfar, "Removing atmospheric turbulence via space-invariant deconvolution," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 35, no. 1, pp. 157–170, Jan. 2013.
- [7] Z. Mao, N. Chimitt, and S. H. Chan, "Image reconstruction of static and dynamic scenes through anisoplanatic turbulence," *IEEE Trans. Comput. Imag.*, vol. 6, pp. 1415–1428, 2020.
- [8] Z. Mao, A. Jaiswal, Z. Wang, and S. H. Chan, "Single frame atmospheric turbulence mitigation: A benchmark study and a new physics-inspired transformer model," in *Proc. Eur. Conf. Comput. Vis.*, 2022, pp. 430–446.
- [9] S. H. Chan, "Tilt-then-blur or blur-then-tilt? clarifying the atmospheric turbulence model," *IEEE Signal Process. Lett.*, vol. 29, pp. 1833–1837, 2022.
- [10] R. Yasarla and V. M. Patel, "Learning to restore images degraded by atmospheric turbulence using uncertainty," in *Proc. IEEE Int. Conf. Image Process.*, 2021, pp. 1694–1698.
- [11] C. P. Lau, H. Souri, and R. Chellappa, "ATFaceGAN: Single face semantic aware image restoration and recognition from atmospheric turbulence," *IEEE Trans. Biom., Behav., Ident. Sci.*, vol. 3, no. 2, pp. 240–251, Apr. 2021.
- [12] N. G. Nair and V. Patel, "Confidence guided network for atmospheric turbulence mitigation," in *Proc. IEEE Int. Conf. Image Process.*, 2021, pp. 1359–1363.

- [13] X. Zhang, Z. Mao, N. Chimitt, and S. H. Chan, "Imaging through the atmosphere using turbulence mitigation transformer," Accessed: Oct. 11, 2022. [Online.] Available: <https://arxiv.org/abs/2207.06465>
- [14] J. P. Bos and M. C. Roggemann, "Technique for simulating anisoplanatic image formation over long horizontal paths," *Opt. Eng.*, vol. 51, no. 10, 2012, Art. no. 101704.
- [15] J. D. Schmidt, *Numerical Simulation of Optical Wave Propagation: With Examples in MATLAB*. Bellingham, WA, USA: SPIE Press, 2010.
- [16] S. L. Lachinova, M. A. Vorontsov, V. V. Dudorov, V. V. Kolosov, and M. T. Valley, "Anisoplanatic imaging through atmospheric turbulence: Brightness function approach," *Proc. SPIE*, vol. 6708, 2007, pp. 111–119.
- [17] B. R. Hunt, A. L. Iler, C. A. Bailey, and M. A. Rucci, "Synthesis of atmospheric turbulence point spread functions by sparse and redundant representations," *Opt. Eng.*, vol. 57, no. 2, 2018, Art. no. 024101.
- [18] N. Chimitt and S. H. Chan, "Simulating anisoplanatic turbulence by sampling intermodal and spatially correlated Zernike coefficients," *Opt. Eng.*, vol. 59, no. 8, 2020, Art. no. 083101.
- [19] Z. Mao, N. Chimitt, and S. H. Chan, "Accelerating atmospheric turbulence simulation via learned phase-to-space transform," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2021, pp. 14759–14768.
- [20] J. W. Goodman, *Introduction to Fourier Optics*, 3rd ed. Placerville, CA, USA: Roberts and Company, 2005.
- [21] D. L. Fried, "Optical resolution through a randomly inhomogeneous medium for very long and very short exposures," *J. Opt. Soc. Amer.*, vol. 56, no. 10, pp. 1372–1379, 1966.
- [22] V. I. Tatarski, *Wave Propagation in a Turbulent Medium*. New York, NY, USA: Dover Publications, 1961.
- [23] R. J. Noll, "Zernike polynomials and atmospheric turbulence," *J. Opt. Soc. Amer.*, vol. 66, no. 3, pp. 207–211, Mar. 1976.
- [24] M. A. Vorontsov and V. Kolosov, "Target-in-the-loop beam control: Basic considerations for analysis and wave-front sensing," *J. Opt. Soc. Amer. A*, vol. 22, no. 1, pp. 126–141, Jan. 2005.
- [25] S. L. Lachinova, M. A. Vorontsov, G. A. Filimonov, D. A. LeMaster, and M. E. Trippel, "Comparative analysis of numerical simulation techniques for incoherent imaging of extended objects through atmospheric turbulence," *Opt. Eng.*, vol. 56, no. 7, 2017, Art. no. 071509.
- [26] C. P. Lau, Y. H. Lai, and L. M. Lui, "Restoration of atmospheric turbulence-distorted images via RPCA and quasiconformal maps," *Inverse Problems*, vol. 37, no. 7, Mar. 2019, Art. no. 074002.
- [27] K. R. Leonard, J. Howe, and D. E. Oxford, "Simulation of atmospheric turbulence effects and mitigation algorithms on stand-off automatic facial recognition," *Proc. SPIE*, Oct. 2012, pp. 182–198.
- [28] M. C. Roggemann and B. M. Welsh, *Imaging Through Atmospheric Turbulence*. Milton Park, Abingdon-on-Thames, U.K.: Taylor & Francis, 1996.
- [29] S. Basu, J. E. McCrae, and S. T. Fiorino, "Estimation of the path averaged atmospheric refractive index structure constant from time lapse imagery," *Proc. SPIE*, vol. 9465, May 2015, Art. no. 94650T.
- [30] D. L. Fried, "Varieties of isoplanatism," *Proc. SPIE*, vol. 0075, 1976, pp. 20–29.



Nicholas Chimitt (Graduate Student Member, IEEE) received the B.S. degree in electrical engineering from Purdue University Northwest, Hammond, IN, USA, in 2017. He is currently working toward the Ph.D. degree in electrical and computer engineering with the School of Electrical and Computer Engineering with Purdue University, West Lafayette, IN, USA. His research interests include computational imaging, wave propagation through random media, and statistical image modeling.



Xingguang Zhang (Graduate Student Member, IEEE) received the B.E. degree in optical science and engineering from Zhejiang University, Hangzhou, China, in 2018, and the M.S. degree in electrical and computer engineering from Purdue University, West Lafayette, IN, USA, in 2021. He joined the i2Lab in 2021. His research interests include turbulence modeling and mitigation and machine learning for video restoration.



Zhiyuan Mao (Graduate Student Member, IEEE) received the B.Eng. degree in electronics and communications engineering (with first-class Hons.) from the Hong Kong University of Science and Technology, Hong Kong, in 2017. He is currently working toward the Ph.D. degree in electronics and communications engineering with Purdue University, West Lafayette, IN, USA. His research interests include computational imaging, image processing, deep learning, and computer vision.



Stanley H. Chan (Senior Member, IEEE) received the B.Eng. degree in electrical engineering from the University of Hong Kong, Hong Kong, in 2007, the M.A. degree in mathematics, and the Ph.D. degree in electrical engineering from the University of California San Diego, La Jolla, CA, USA, in 2009 and 2011, respectively. From 2012 to 2014, he was a Postdoctoral Research Fellow with Harvard University, Cambridge, MA, USA. He joined Purdue University, West Lafayette, IN, USA, in 2014, where he is currently an Elmore Associate Professor of electronics and communications engineering. His research interests include single-photon imaging, imaging through atmospheric turbulence, and computational photography. He is the author of a popular undergraduate textbook *Introduction to Probability for Data Science*, Michigan Publishing 2021. The book is part of the University of Michigan's, Ann Arbor, MI, USA, free textbook initiative to disseminate high-quality educational materials to students and families around the world. Since the launch of the book in Fall of 2021, the book has saved more than \$3 million USD for families globally. He is a Senior Area Editor of the IEEE TRANSACTIONS ON COMPUTATIONAL IMAGING. He was an Associate Editor for the OSA Optics Express in 2016–2018, and an Elected Member of the IEEE Signal Processing Society Technical Committee in Computational Imaging in 2015–2020.