Physics-informed PatchGAN for atmospheric turbulence Phase Screen Generation

Cristhof J R Runge, Ulisses Dias

***Abstract*— We propose a physics-informed generative adversarial network (GAN) for synthesizing Kolmogorov-based atmospheric turbulence phase screens, supporting free-space optical (FSO) link simulation. Our approach integrates a PatchGAN discriminator with spectral normalization, combined with physics-inspired loss function components, including weighted spectral loss, higher-order moment matching, and minibatch diversity to enhance realism and variability. Training leverages relativistic adversarial loss, instance noise injection, and learning rate balancing for improved stability. Statistical and visual comparisons suggest that our framework produces phase screens with Kolmogorov-like statistics and realistic diversity, offering a tool for modeling optical propagation through turbulence. By providing a differentiable and physically consistent turbulence model, we aim at enabling the integration of realistic atmospheric effects into end-to-end trainable FSO communication systems, thereby facilitating more accurate optimization and overcoming the limitations of non-differentiable classical models.**

***Keywords*— FSO, atmospheric turbulence, GAN, Kolmogorov spectrum, OAM**

I. Introduction

In emerging paradigms for end-to-end optimization of free-space optical (FSO) communication systems, it is highly desirable to model the complete transmitter-channel-receiver chain within a unified, differentiable deep learning framework. However, modeling atmospheric turbulence based on Kolmogorov theory poses significant challenges, as practical implementations often require approximations and numerical methods that can complicate the differentiation process [1]. Previous attempts to integrate turbulence effects into neural network architectures have relied on approximation techniques like gradient passthrough, which introduced modeling inaccuracies and limited the overall performance improvements. To address these limitations, we propose a generative adversarial network (GAN)-based approach to emulate turbulence phase screens. By providing a realistic, lightweight, and naturally differentiable model of turbulence effects, our method facilitates proper backpropagation through the channel and enhances the feasibility of accurate end-to-end optimization strategies.

In order to simulate such turbulence effects in practical settings, various methods for phase screen generation have been developed. Generation of atmospheric phase screens plays a central role in simulations in adaptive optics, free-space optical communication, and astronomical imaging. Traditionally, the generation of atmospheric turbulence phase screens has relied on Fourier-based spectral synthesis techniques, which employ filtering in the frequency domain using models such as the Kolmogorov or von Kármán turbulence spectra [2], [3], [4]. While these classical approaches can yield high-fidelity representations of turbulence, they are often computationally intensive, particularly when simulating large apertures or high-resolution screens [5].

In recent years, data-driven methodologies for atmospheric phase screen generation have emerged as promising alternatives to classical approaches. Harnessing advances in deep learning, generative adversarial networks (GANs) [6] have demonstrated the ability to efficiently learn and reproduce complex statistical distributions directly from data. This approach offers substantial reductions in computational requirements while maintaining, and maybe even improving, the fidelity of turbulence statistics such as those predicted by the Kolmogorov model.

Most prior works utilize standard adversarial losses as their primary training objective [7], while only a few incorporate explicit physical metrics—such as spectral penalties or statistical constraints—into their composite loss functions [8]. The use of more sophisticated statistical losses, including minibatch diversity enforcement and higher-order moment matching (skewness and kurtosis), as well as advanced GAN stabilization techniques, has been seldom explored in the context of atmospheric phase screen generation.

Here, we present a GAN architecture in which the discriminator is based on the PatchGAN paradigm and incorporates spectral normalization. The generator and discriminator are trained jointly under multiple physics- and statistics-informed regularizations. Experimental results demonstrate that this approach can generate physically plausible and highly diverse synthetic phase screens, faithfully capturing the statistical properties of atmospheric turbulence.

II. Methodology

In this section, we describe the system architecture designed to generate synthetic atmospheric phase screens, as well as its main features.

A. Network Architecture

The system is composed of a generator network and a discriminator network. The discriminator serves as a classifier tasked with distinguishing between genuine samples drawn

The authors are affiliated with the School of Technology - Unicamp, Limeira - SP - Brazil, e-mails: cristhof@unicamp.br, ulissesd@unicamp.br.

This study was financed in part by the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior - Brasil (CAPES) - Finance Code 001 and Programa de Pós Graduação em Tecnologia da Faculdade de Tecnologia da Unicamp (PPGT).

from the dataset and synthetic samples produced by the generator. Conversely, the generator is trained to produce synthetic data that is statistically indistinguishable from real samples, thereby aiming to deceive the discriminator during the adversarial learning process.

In this work, we refer to simulated phase screens generated via classical spectral synthesis using the Kolmogorov power law as 'real' samples, in accordance with GAN terminology. While these are not experimentally measured screens, they serve as the statistical reference for training the generator.

Throughout training, the generator and discriminator networks are optimized in an adversarial fashion. The discriminator seeks to maximize its ability to correctly distinguish real samples from generated ones, while the generator aims to maximize the discriminator's likelihood of misclassifying generated samples as real. This adversarial interplay drives both networks toward improved performance, ultimately resulting in the generation of realistic synthetic phase screens.

1) Generator: The generator network is designed to transform a latent vector z of dimension 100, randomly sampled from a standard normal distribution, into a two-dimensional phase screen of size 64×64, corresponding to the synthetic turbulence phase screen presented as input to the discriminator. In the first stage, the input vector passes through a fully connected linear layer that projects it into a feature vector of length $128 \times 8 \times 8 = 8192$, which is then reshaped into a tensor of dimensions $[128, 8, 8]$.

This intermediate tensor is processed through a sequence of three transposed convolutional layers to incrementally increase the spatial resolution while reducing the number of channels. The first transposed convolution layer upsamples the feature map to 16×16 while reducing the depth to 64 channels, followed by batch normalization and a ReLU activation. The second layer further increases the resolution to 32×32 with 32 channels, again followed by batch normalization and ReLU activation. The final transposed convolution produces a single-channel output at the target resolution of 64×64, yielding the synthetic phase screen. Batch normalization is applied after each upsampling step except the last, and all intermediate activations employ ReLU nonlinearity. The architecture does not apply an explicit output activation function, so the range of generated phase values is determined by the preceding layers. This generator is trained end-to-end to capture both the spectral and statistical properties of atmospheric turbulence phase fluctuations, as modeled by the Kolmogorov theory.

2) PatchGAN Discriminator: The discriminator employs a PatchGAN architecture to distinguish between real and synthetic phase screens by focusing on local image statistics. The network receives a single-channel input image of size 64×64 and processes it through a sequence of four convolutional layers. Each convolution is equipped with spectral normalization to enforce Lipschitz continuity, which has been shown to significantly improve the stability of GAN training [10].

The initial convolutional layer expands the input to 32 channels at half the spatial resolution, followed by a Leaky ReLU activation. This is followed by a second convolution that increases the channel count to 64 while maintaining the reduced spatial size, accompanied by batch normalization and another Leaky ReLU activation. The third convolutional layer further increases the channel count to 128, also followed by batch normalization and Leaky ReLU activation.

The final convolution maps the features to a single output channel, producing a 7×7 feature map instead of a single scalar. Each unit of this output map provides a local real/fake prediction for overlapping 46×46 patches across the input phase screen. This patch-based approach enforces the learning of high-frequency local patterns and textures, which is critical for accurately modeling the fine-grained structure of atmospheric turbulence. By leveraging spectral normalization, batch normalization, and Leaky ReLU activations, the discriminator achieves robustness against training instabilities and overfitting, effectively guiding the generator to produce phase screens that are both globally and locally consistent with the statistics of the Kolmogorov model.

B. Loss Functions

The incorporation of physics-inspired loss terms is central to our generative modeling strategy. This subsection provides an overview of the loss functions and quantitative metrics employed to guide the training process. Our objective is to synthesize phase screens that faithfully reproduce the physical and statistical properties of atmospheric turbulence.

To this end, the generator is trained using a composite loss function comprising several complementary terms. These include an adversarial realism constraint, spectral fidelity enforcement, higher-order moment matching, and diversity promotion, each designed to encourage the generation of samples consistent with both theoretical expectations and empirical characteristics. Formally, the generator's total loss is expressed as a weighted sum of the following components:

$$\mathcal{L}_G = \mathcal{L}_{\text{GAN}} + \lambda_{\text{spec}}\,\mathcal{L}_{\text{spec}} + \lambda_{\text{mom}}\,\mathcal{L}_{\text{mom}} + \lambda_{\text{div}}\,\mathcal{L}_{\text{div}} \quad (1)$$

where λ_* are empirically chosen hyperparameters.

a) Adversarial (PatchRSGAN) Loss:: We employ the relativistic average PatchGAN loss [11], using the softplus function as a numerically stable alternative to the standard negative log-likelihood. The adversarial term for the generator is:

$$\mathcal{L}_{\text{GAN}} = \mathbb{E}_{x_f}\left[\text{softplus}\left(-(D(x_f) - \mathbb{E}_{x_r} D(x_r))\right)\right] \quad (2)$$

where $D(x)$ denotes the PatchGAN discriminator output, x_r are real samples, and x_f are generated ("fake") samples. As previously mentioned, 'real' refers to Kolmogorov-simulated phase screens used as a reference in GAN training, whereas 'fake' designates the outputs generated by the network under optimization during adversarial learning.

b) Spectral Loss:: To ensure that the power spectra of generated phase screens match the target Kolmogorov spectrum, we include a weighted ℓ_2 loss between their normalized

radially averaged spectra:

$$\mathcal{L}_{\text{spec}} = \alpha \sum_{k=k_{\text{min}}}^{k_{\text{max}}} w_k \left[S_k(G(z)) - S_k^*\right]^2 \tag{3}$$

where $S_k(G(z))$ and S_k^* denote the spectrum of a generated and a reference screen at frequency bin k, w_k are cosine weights, and α is a scaling factor.

c) Higher-Order Moment Loss:: Recognizing that simple first- and second-order statistics are insufficient for capturing phase screen realism, we penalize discrepancies in skewness and kurtosis (which are functions of third and fourth central moments) between real and fake samples:

$$\begin{aligned}\mathcal{L}_{\text{mom}} = &\left|\text{skew}(G(z)) - \text{skew}(x_{\text{ref}})\right| \\ &+ 0.25\ \left|\text{kurt}(G(z)) - \text{kurt}(x_{\text{ref}})\right|\end{aligned} \tag{4}$$

where $\text{skew}(x)$ and $\text{kurt}(x)$ are the sample skewness and excess kurtosis, respectively, over all pixels.

d) Diversity Loss:: To discourage mode collapse and encourage diversity among the generated samples within a batch, we include a negative (normalized) covariance trace penalty:

$$\mathcal{L}_{\text{div}} = -\alpha \frac{1}{BHW} \text{Tr}\left[\text{Cov}(G(z_1), \ldots, G(z_B))\right] \tag{5}$$

where Tr denotes the trace, Cov is the sample covariance, B is the batch size, H, W are the spatial dimensions and α is a scaling factor.

e) Note:: A structure function loss, penalizing mismatch in two-point statistics, was implemented for analysis but omitted from generator updates in the current work ($\lambda_{\text{struct}} = 0$).

C. Training

Training a generative adversarial network (GAN) can be formally cast as a two-player minimax game, wherein the generator G and the discriminator D pursue opposing objectives. This dynamic is encapsulated by the value function $V(G, D)$:

$$\min_G \max_D \ V(G, D) \tag{6}$$

At equilibrium—that is, a Nash equilibrium in the language of game theory—the following conditions hold:

- The discriminator cannot consistently distinguish real samples from synthetically generated ones, thereby assigning equal probability to both classes: $D(a) = 0.5$ for any input a.
- The generator has succeeded in accurately reproducing the target data distribution, so that the distribution of generated samples $P_{G(z)}$ matches the empirical data distribution $P_a(a)$.

In our setting, a classical Kolmogorov phase screen generator, based on spectral synthesis, is used to create a training dataset comprising 1,000 phase screens, each with a spatial size of 64×64 pixels. This dataset serves as the reference for supervising the adversarial training process.

During initial training experiments, we observed a tendency for the discriminator to become excessively confident—quickly converging to degenerate solutions in which all inputs, including real samples, were misclassified as fake. This "saturation" of the discriminator reduced the magnitude of gradients provided to the generator, leading to vanishing feedback and poor generative performance. To counteract this instability, we adopted an asymmetric training schedule in which, for each discriminator update, the generator is updated four times. This adjustment helps to equilibrate the adversarial dynamics and ensures that neither network outpaces the other, thus mitigating the risks of instability and mode collapse while promoting robust convergence.

Furthermore, we introduce *instance noise* as an additional stabilization technique. Specifically, controlled Gaussian noise is injected into both real and generated images prior to presentation to the discriminator. Adding such noise serves to smooth the data distributions and the discriminator's decision boundary, particularly in the early stages of training. This regularization prevents the discriminator from fixating on spurious high-frequency artifacts or trivial statistical differences, thus compelling the generator to align its output with the target data distribution at both global and structural levels [9].

Jointly, these modifications contributed to a more stable and productive adversarial optimization process, yielding improved sample quality and more faithful reproduction of turbulence statistics.

III. Simulation Results

After implementing the previously described stabilization techniques, the adversarial dynamics reached a desirable equilibrium. As illustrated in Figure 1 , the discriminator's success rate converged to approximately 0.5 after 960 discriminator update steps. Given that each epoch comprises 32 discriminator updates, this indicates that around 30 epochs are required for convergence. Notably, Figure 1 also demonstrates that this equilibrium between the discriminator and the generator is maintained up to at least 50 epochs. This sustained behavior aligns with the theoretical Nash equilibrium, in which the discriminator is unable to reliably distinguish between real and generated samples—assigning them equal probability.

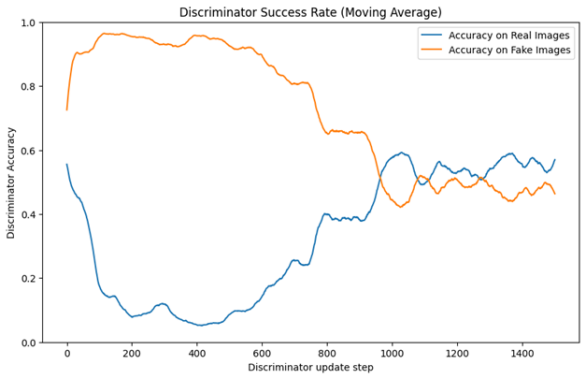

Fig. 1. Evolution of Discriminator Sucess Rate during training

Further evidence of successful training is observed in the close match between the empirical distributions of phase values for both real and generated phase screens. As illustrated in Figure 2, the histogram of the generated data closely

parallels that of the true Kolmogorov phase screens, indicating strong statistical similarity and suggesting that the generator successfully avoided mode collapse.

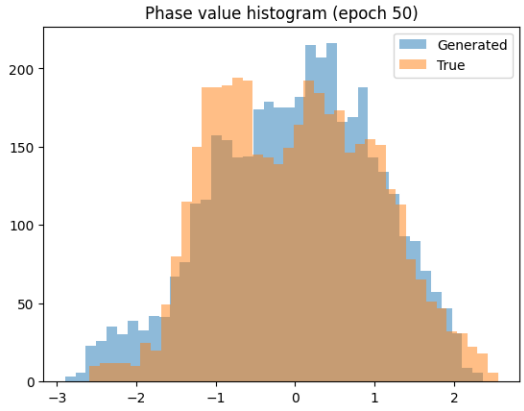

Fig. 2. Comparison of Generated vs. True Phase Distribution Histograms (Values: Point Count per Bin, Scale: Radians)

Finally, Figure 3 presents qualitative visual comparisons between representative synthetic and real phase screens. The generated images faithfully reproduce the spatial structures and complexity characteristic of turbulence-induced phase variations, underscoring the effectiveness of the proposed approach in capturing both local and global statistical properties.

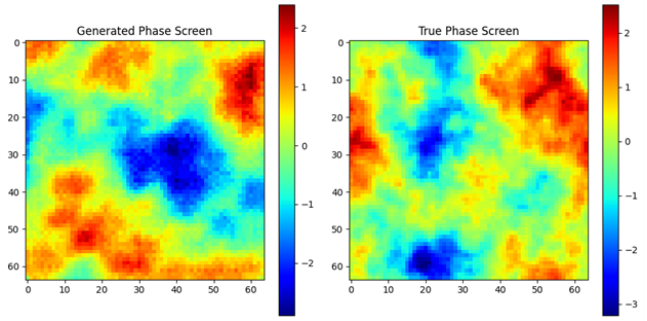

Fig. 3. Synthetic (Generated) and Real (True) Phase Screen after training; scale in radians

The quantitative and qualitative analyses presented in this section validate the effectiveness of the proposed framework in generating physically realistic phase screens. In the next section, we discuss the broader implications of these findings and the potential for further improvements.

IV. Discussion

The use of physics-based and higher-order statistical losses remains seldom explored in the neural phase screen generation literature. Our findings suggest that such regularizations are effective in ensuring the physical and statistical plausibility of synthetic screens. Additionally, the incorporation of relativistic adversarial loss and advanced GAN stabilization techniques—commonly employed in computer vision, but rare in this domain—proved to be beneficial for training stability and sample quality.

Although the structure function loss was not incorporated into the present results, the proposed combination of advanced statistical and adversarial modeling strategies already represents a significant advancement in the state of the art for atmospheric turbulence simulation using data-driven methods. Future investigations will explore the inclusion of explicit structure function regularization to further enhance physical fidelity.

The insights gained from this work highlight both the strengths and the current limitations of the proposed approach, setting the stage for future advancements. We summarize the main contributions and potential directions for further investigation in the next section.

V. Conclusions and Future Work

We presented a novel framework for the artificial generation of Kolmogorov-based atmospheric turbulence phase screens, leveraging generative adversarial networks with multiple physics- and statistics-informed loss functions. The synergy between spectral, higher-order moment, diversity, and relativistic adversarial losses, combined with advanced training techniques, constitutes a meaningful advancement beyond previously published approaches.

Our results demonstrate that the proposed method effectively captures both the global and local statistical properties of atmospheric turbulence, producing synthetic phase screens that closely match Kolmogorov-based references in both structure and distribution.

Although the explicit use of structure function loss was deferred in this study, its incorporation remains a promising direction for future work aimed at further enhancing the physical fidelity of synthetic turbulence modeling. Future efforts will also explore scaling to larger screen sizes and investigating the transferability of the learned turbulence representations to downstream tasks, such as end-to-end optimization of free-space optical communication systems.

References

[1] C. J. Roosen Runge and U. Dias, “End-to-end optimization of neural networks for orbital angular momentum-based transmission system design,” *Optical Engineering*, vol. 64, no. 4, pp. 048106–048106, 2025.

[2] R. G. Lane, A. Glindemann, and J. C. Dainty, “Simulation of a Kolmogorov phase screen,” *Waves in Random Media*, vol. 2, no. 3, pp. 209–224, 1992.

[3] Jingsong Xiang, “Fast and accurate simulation of the turbulent phase screen using fast Fourier transform,” *Optical Engineering*, vol. 53, no. 1, p. 016110, 2014. Available: https://doi.org/10.1117/1.OE.53.1.016110

[4] Peng Jia, Dongmei Cai, Dong Wang, and Alastair Basden, “Real-time generation of atmospheric turbulence phase screen with non-uniform fast Fourier transform,” *Monthly Notices of the Royal Astronomical Society*, vol. 450, no. 1, pp. 38–44, Apr. 2015. Available: https://doi.org/10.1093/mnras/stv602

[5] P. Jia, D. Cai, D. Wang, and A. Basden, “Simulation of atmospheric turbulence phase screen for large telescope and optical interferometer,” *Monthly Notices of the Royal Astronomical Society*, vol. 447, no. 4, pp. 3467–3474, 2015. doi:10.1093/mnras/stu2655.

[6] Chenda Lu, Qinghua Tian, Xiangjun Xin, Lei Zhu, Qi Zhang, Haipeng Yao, Huan Chang, and Ran Gao, “Orbital Angular Momentum (OAM) Recognition with Generative Adversarial Network (GAN) based Atmospheric Modeling,” in *Proc. 2021 Optical Fiber Communications Conference and Exhibition (OFC)*, pp. 1–3, 2021.

[7] Shyam Nandan Rai and C. V. Jawahar, "Removing Atmospheric Turbulence via Deep Adversarial Learning," *IEEE Transactions on Image Processing*, vol. 31, pp. 2633–2646, 2022.

[8] Dongxiao Zhang, Junjie Zhang, Yinjun Gao, and Taijiao Du, "Optical Field-to-Field Translation Under Atmospheric Turbulence: A Conditional GAN Framework with Embedded Turbulence Parameters," *Photonics*, vol. 12, no. 4, Article 339, 2025. Available: https://www.mdpi.com/2304-6732/12/4/339

[9] Martin Arjovsky and Léon Bottou, "Towards Principled Methods for Training Generative Adversarial Networks," arXiv preprint arXiv:1701.04862, 2017. Available: https://arxiv.org/abs/1701.04862

[10] Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida, "Spectral normalization for generative adversarial networks," in *Proceedings of the International Conference on Learning Representations (ICLR)*, 2018. Available: https://arxiv.org/abs/1802.05957

[11] A. Jolicoeur-Martineau, "The relativistic discriminator: a key element missing from standard GAN," in *Proceedings of the International Conference on Learning Representations (ICLR)*, 2019.