

Document Version

Final published version

Licence

Dutch Copyright Act (Article 25fa)

Citation (APA)

Arvis, M., Ciarella, S., Loicq, J., & Saathof, R. (2025). Single shot line-of-sight atmospheric turbulence profiling for laser satellite communications with STORM. *Applied Optics*, 64(32), 9586-9597. <https://doi.org/10.1364/AO.572341>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

In case the licence states "Dutch Copyright Act (Article 25fa)", this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership.
Unless copyright is transferred by contract or statute, it remains with the copyright holder.

Sharing and reuse

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.



Single shot line-of-sight atmospheric turbulence profiling for laser satellite communications with STORM

MARGUERITE ARVIS,^{1,*}  SIMONE CIARELLA,² JÉRÔME LOICQ,¹ AND RUDOLF SAATHOF¹ 

¹TU Delft, Space Engineering Department, Faculty of Aerospace Engineering, Delft, Netherlands

²eScience Center of the Netherlands, Amsterdam, Netherlands

*m.arvis@tudelft.nl

Received 3 July 2025; revised 10 October 2025; accepted 10 October 2025; posted 13 October 2025; published 3 November 2025

Laser satellite communication technologies are promising next-generation communications systems, offering higher data rates, more secure links, and cost-effective operations. One of the remaining challenges to tackle for ensuring sufficient link availability is atmospheric turbulence. While turbulence effects on laser links can be partly compensated for with corrective technologies and algorithms, these methods would benefit from better knowledge of turbulence profiles on the communications channel, both for system design or real-time assistance during operations. As state-of-the-art turbulence profilers are not designed to measure profiles from a laser beam on a moving satellite, this paper proposes a line-of-sight turbulence profiling tool tailored for laser satellite communications. Speckle-based turbulence observation and reconstruction via machine learning uses surrogate learning to build a model that can reconstruct optical turbulence profiles ($C_n^2(h)$) from a single shot measurement of a speckle pattern. In this paper, the first modeling results of this flexible approach demonstrate that eight-layer turbulence profiles can be reconstructed from simulated single speckle images of a star with less than 10% error on the Fried parameter. © 2025 Optica Publishing Group. All rights, including for text and data mining (TDM), Artificial Intelligence (AI) training, and similar technologies, are reserved.

<https://doi.org/10.1364/AO.572341>

1. INTRODUCTION

Since the 1990s, laser satellite communications (laser satcoms) have been seen as a promising solution for next-generation global communications systems [1,2]. Compared to classical radio-frequency communications, the shorter wave infrared spectrum produces narrower beams and lower free-space losses. This offers four main advantages: higher throughput, more secure links, optimized size, weight, and power of onboard systems, and finally, lower operational costs due to unregulated bandwidth usage [3].

However, laser satcom performance and availability suffer from the loss-inducing effects of laser beam propagation through the atmosphere [4]. In order of severity, these arise from cloud coverage, aerosol and molecular attenuation, and lastly, atmospheric turbulence. This research focuses on the latter and measuring its effects on laser links such as beam spread, beam wander, and scintillation. The resulting irradiance fluctuations on the receiver contribute to link quality degradation and potential loss of signal, amounting to overall reduced communications performance.

In order to mitigate degradations caused by turbulence, various corrective techniques, such as adaptive optics (AO) or higher-level coding strategies, are proposed [5]. Properly

designing these systems and protocols requires prior knowledge of channel quality. The latter is tightly linked to the optical turbulence profile (OTP) along the line of sight, i.e., the turbulence strength as a function of the altitude, or $C_n^2(h)$, between the ground and satellite. In particular, basic AO sizing and point-ahead capabilities are driven by the Fried parameter and the isoplanatic angle. These integrated values of C_n^2 , corresponding, respectively, to the coherence length and angle of the turbulent atmosphere, are easily computable using a differential image motion monitor (DIMM) [6].

On the other hand, a multilayer characterization of channel turbulence is more advantageous for several reasons. First and foremost, while the isoplanatic angle is an adequate starting point, phase error estimation for uplink pre-compensation is better achieved when using multilayered turbulence profiles [7]. Additionally, observing actual turbulence profiles alongside the resulting irradiance fluctuations will provide important insight to improve link performance estimations (e.g., bit error rates) as a function of turbulence altitude and strength. Moreover, collecting turbulence profiles from different satellite elevations and over urban areas could provide interesting statistical information on turbulence in typical laser satcom conditions.

Needless to say, while no method is yet designed specifically for line-of-sight turbulence profiling on laser satcom links, they are abundant for astronomy [8]. Conveniently, many of these techniques rely on starlight wavefront sensing or speckle patterns imaging, i.e., measuring irradiance interference patterns on the pupil plane. Resulting from phase aberrations generated when light propagates through turbulence cells of varying refractive indices, speckle patterns carry inherent statistical information about the turbulent state of the traversed air column. At first glance, such methods are easily applicable to laser satcoms by switching the celestial point source from stars to satellite laser beams, which also ensures a line-of-sight measurement.

One of the more widely established turbulence profiling techniques is that of scintillation detection and ranging (SCIDAR) [9]. Based on the autocorrelations of a binary star speckle pattern, SCIDAR yields high vertical resolution turbulence profiles. Moreover, generalized SCIDAR [10] allows profiling of the whole atmosphere, including the ground layer. Unfortunately, SCIDAR is inherently incompatible with laser satcoms since it requires a double light source, whereas a laser communications satellite can only act as a single source. Likewise, SLODAR [11] and one of its variants, CO-SLIDAR [12], perform similar analyses to SCIDAR using wavefront sensors but are equally incompatible due to their use of a double source.

To overcome the binary source issue, one could point to single star SCIDAR (SSS) [13,14]. Using the same principle as SCIDAR, SSS employs a time-series of single star speckle pattern images to profile whole atmospheric turbulence. Similarly, SCO-SLIDAR [15] and AC-SLODAR [16], the single-star versions of SLODAR, provide single source OTP estimations. However, while the single source is convenient for laser satcoms, using a speckle pattern time series in the low Earth orbit (LEO) satellite scenario puts more stringent requirements on the detection frame rates and workable wind speeds. Due to the high angular speeds of typical LEO laser satellites, the air column measured by the emanating laser beam will not be the same at two consecutive instances. This results in limitations of the turbulence profiling range by the detector capabilities, as detailed in Section 2.B.

Despite suffering equally from time-separated measurements, this overview would be incomplete without a mention of scintillation-based methods. These include the Shack–Hartmann image motion monitor (SHIMM) [17], a modified version of the DIMM, which provides an OTP with a four-layer vertical resolution based on scintillation measurements, i.e., measures of the total irradiance fluctuations at the receiver. Similarly, the multi-aperture scintillation sensor (MASS) [18] produces six-layer OTPs. However, even when combined with the DIMM [19,20], these techniques provide relatively poor altitude resolutions, especially near the ground layer, where turbulence is often the strongest. A different scintillation-based method, the ring-image next-generation turbulence sensor, or RINGSS [21], produces an eight-layer OTP using defocused ring images of a star, with better low-altitude resolution and greater flexibility.

Beyond their initial strengths and shortcomings, however, all the aforementioned techniques require either multiple sources or multiple measurements. While technically still feasible, the

latter techniques could be limiting for LEO laser satellites. In light of this, the purpose of this paper is to introduce a novel optical turbulence profiling technique tailored to laser satellite communications. STORM, or speckle-based turbulence observation and reconstruction via machine learning, aims to remedy the technological gap by using machine learning (ML) to determine turbulence profiles from single shot speckle patterns of a single source. This single-source-single-shot approach bypasses detector limitations as well as some physical limitations arising from wind speeds and turbulence regimes. While pending field tests, the STORM concept, development, and first simulation results are described in this paper.

For the sake of clarity, the general concept of this technique is explained in Section 2. Section 3 elaborates on the method for developing the first STORM model: the dataset generation and the machine learning model design. Finally, the results of the model training and subsequent discussions are addressed in Section 4 before concluding in Section 5.

2. STORM CONCEPT

The novelty of STORM is threefold. First of all, though applicable to astronomy, this technology is primarily targeted at laser satellite communications. Second, STORM is a single shot profiler from a single source, a combination of traits not yet explored by the state of the art. Third, thanks to the use of surrogate learning, the STORM machine learning model is a flexible post-processing method that can easily be generalized to different use cases. This section serves to highlight these innovative aspects by explaining the STORM system, why it is advantageous, and how the theory aligns.

A. System Overview

The STORM technology encompasses two elements: the hardware, consisting of an imaging instrument, and the software, consisting of an image-processing technique that outputs a turbulence profile. Just like in SCIDAR, the instrument is a pixelated detector positioned in the pupil plane to collect a speckle pattern from the light source after atmospheric propagation. Unlike SCIDAR or SSS, however, STORM's post-processing produces a full turbulence profile from a single speckle image of a single source without needing a binary source or a time-series measurement. This is achieved by employing a neural network trained to extricate turbulence profiles from the information encoded in speckle patterns. An explicative diagram of the full technology is depicted in Fig. 1.

It is worth mentioning that although ML-based turbulence reconstruction solutions have been implemented in prior studies [22–24], they have sought to reconstruct local turbulence profiles or single integrated parameters from meteorological or AO inputs on a given site rather than produce a generalized tool to reconstruct a full profile from a single shot measurement on any site.

It should also be noted that for the purpose of demonstrating the STORM concept, the case study presented in this paper is that of a single star observation at a 550 nm wavelength. Hopefully, it will become clear to the reader that this same procedure (shown in Fig. 4) can easily be applied to any point

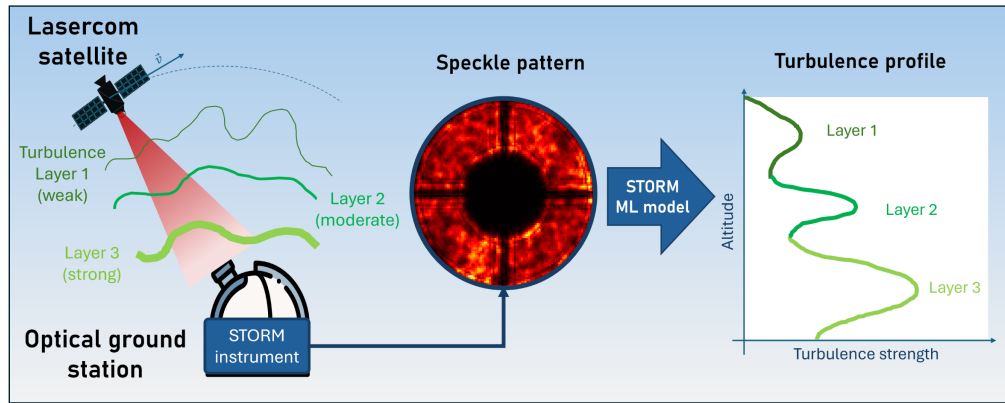


Fig. 1. STORM concept: the speckle pattern is imaged on a ground station during a satellite laser communications pass; by feeding this single image to the pretrained STORM ML model, the atmospheric turbulence profile can be reconstructed in output.

source, such as that of a satellite laser, by modifying the appropriate parameters in the development procedure and hardware. Furthermore, this work is only intended to validate the STORM technique via simulation results. Experimental results to validate the STORM instrument and the full technology will be addressed in future work.

B. Single Shot Advantage

When evaluating the applicability of current profiling techniques to laser satcoms instead, one could naturally question if a stereo- [25] or single star SCIDAR approach would be successful. In fact, imaging speckle patterns from two consecutive satellite positions could equate to a double source or a single source separated by the satellite's angular displacement. However, albeit theoretically feasible, this approach is limited by detector capabilities.

To demonstrate this, one can consider the example of a typical LEO communications satellite with a 550 km orbit performing

a high elevation pass. The average angular velocity with relation to the ground is of around $v_{\theta_{\text{sat}}} = 0.5 \text{ deg/s}$, and reaches up to $v_{\theta_{\text{sat}}} = 1 \text{ deg/s}$ at culmination [see Fig. 2(a)]. As the ground telescope tracks the satellite at the same angular rate, this means that still turbulence features at an altitude h in the measured air column move along and out of the frame by a transversal displacement $d(b, t)$ approximated by

$$d(b, t) = \frac{h * v_{\theta_{\text{sat}}}(t) + w(b)}{r_f}, \quad (1)$$

where r_f is the pupil camera frame rate. Moreover, since features are generally not still but rather displaced by wind, a conservative case is considered by supposing counter-directional wind, $w(b)$. In this study, the Bufton wind model is considered as the wind profile.

While this feature displacement is not an issue in the case of a single shot measurement (with low enough integration times), it becomes problematic for multi-frame techniques when the

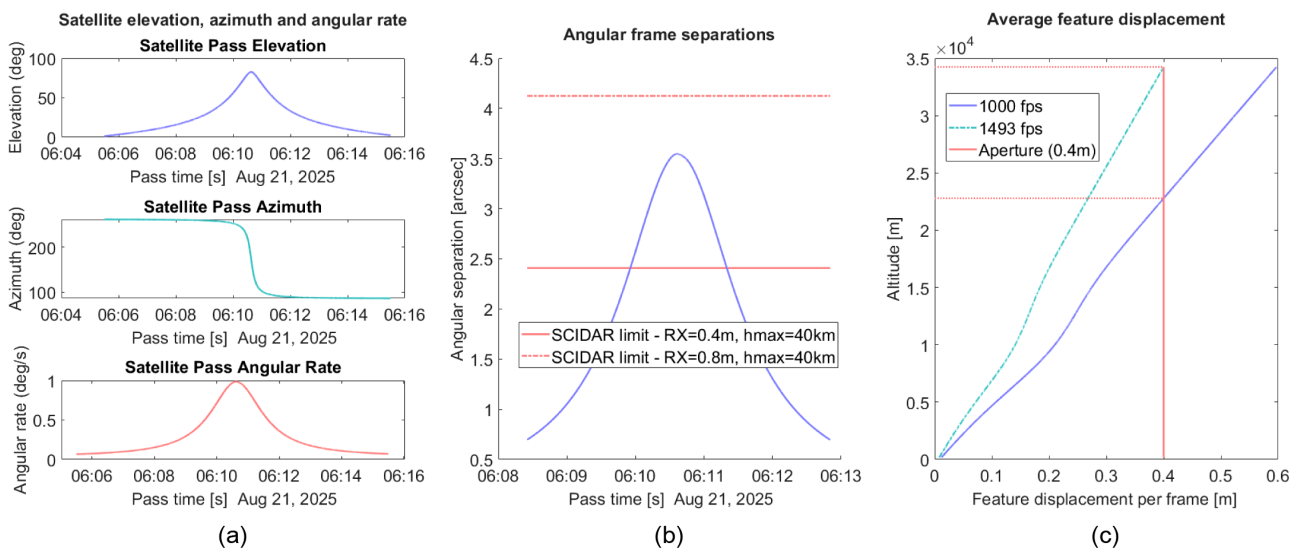


Fig. 2. (a) LEO (550 km) satellite high pass: elevation, azimuth, and angular rate. (b) Angular separation between satellite positions for consecutive frames versus physical angular limits for stereo-SCIDAR. (c) Comparison of the turbulence feature displacement across the imaging frame for different frame rates, at maximum satellite angular velocity, with counter-directional wind (Bufton model). The vertical line indicates the physical displacement limit set by the ground telescope aperture.

displacement surpasses that of the telescope aperture, D_{rx} . In other terms, when the separation between satellite positions in consecutive frames exceeds the physical angular limitations set for binary stars in SCIDAR, then the profiling range becomes limited [see Fig. 2(b)]. For reference, binary stars used for SCIDAR are typically within $3'' - 10''$ [26]. To remedy this, a minimum frame rate $r_{f \min}$ is required such that

$$r_{f \min} = \frac{h_{\max} * v_{\theta \text{sat}} + w(h_{\max})}{D_{rx}}, \quad (2)$$

where h_{\max} is the desired profiling range.

In this study [Fig. 2(c)], a good camera frame rate of 1 kfps would limit the profiling altitude range of stereo-SCIDAR to about 22 km at peak satellite angular rates, with a minimum frame rate requirement of 1500 fps to profile up to 35 km. While such frame rates are not unrealistic, developing single shot methods has the clear advantage of releasing profiling techniques from frame rate constraints. In both cases, however, it should be noted that the satellite motion will pose an integration time requirement if blur is to be avoided on higher layers. Though it is unclear how much blur can be withstood, the integration time requirement could be as low as tens of microseconds.

C. STORM Operational Principle

While advantageous, single shot profiling must also be feasible, which this section will attempt to demonstrate. Optical turbulence is defined as a variation of the refractive index given by the structure function D_n [4] such that

$$D_n(r, t) = \langle |n(r_0 + r, t) - n(r_0, t)|^2 \rangle = C_n^2 r^{3/2}, \quad (3)$$

where n is the refractive index, r_0 is any point in a statistically homogeneous and isotropic medium, r is a normalized distance contained within the Kolmogorov inner and outer scales, and C_n^2 is the refractive index structure constant, which is more widely used as an indicator of turbulence strength. The statistical homogeneity and isotropy are assumed to extend horizontally and vertically over the length of the outer scale; thus, this formula stays valid over any relevant turbulence layer.

Naturally, the propagation of light through changing refractive indices produces beam deviations, of which the interferences create speckles on the receiver plane, as in Fig. 3. The size of said speckles is defined by the first Fresnel zone [4], $\sqrt{\lambda L}$, where λ is the wavelength, and L is the propagation distance. Moreover, as the interference itself is caused by the

refractive index delta, the speckle contrast is a function of C_n^2 . This phenomenon is visually demonstrated for a single turbulence layer in Fig. 3, where one can see that higher altitude turbulence generates larger speckles and stronger turbulence generates deeper contrasts. In practice, this means that a speckle pattern contains sufficient spatial information to reconstruct an OTP within an altitude range (h_{\max}) set by the maximum detectable speckle size, which itself is set by the telescope aperture (d) such that

$$h_{\max} = \frac{d^2}{\lambda}. \quad (4)$$

By applying the Taylor frozen turbulence hypothesis [4], multiple shot methods like single star SCIDAR use a similar theory in a temporal domain. The cross-correlation of wind-shifted speckle patterns gives the necessary altitude and C_n^2 information via the width and height of the correlation peaks, respectively. Additionally, the position of these peaks provides information about wind speeds, which the proposed spatial analysis does not. However, the spatial analysis has the advantage of not being bounded by minimum or maximum wind speeds.

While the inversion problem appears straightforward for a single turbulence layer as presented above, it becomes more complex with a multilayer atmosphere. In order to circumnavigate this complexity while demonstrating the feasibility of single shot profiling, it was decided to use machine learning. Advantageously enough, this method also has the benefit of potentially fitting a solution, which is not limited by the strength of the turbulence regime.

3. METHOD

A. Development Process

In order to develop an adequate neural network for STORM, a surrogate learning approach was used. This means that the ML model was trained to mimic the output of a numerical model, rather than that of physical results. In other words, the training data, i.e., the speckle pattern dataset, is collected through numerical wave propagation simulations as opposed to real data.

Although large databases of speckle observations exist from stereo-SCIDAR measurements [27], the simulation approach was selected for adaptability to wavelength, telescope diameter, detector resolution, and turbulence conditions. First, generating tailored training data offers the advantage of creating a model that learns to recognize turbulence profiles of any shape and strength. This ensures that STORM can accurately predict




Fig. 3. Visual demonstration of the impact of turbulence altitude and strength on a speckle pattern received at ground level: turbulence at higher altitudes generates larger speckles [as seen by the spot sizes between (a) and (b)] and stronger turbulence generates deeper contrasts [as seen by the color gradients between (b) and (c)].

any OTP, even over previously unmonitored sites such as those of urban laser satcom ground stations. Moreover, simulating artificial speckle patterns offers the flexibility to quickly generate data fitted to the right ground station parameters (e.g., telescope diameter, image resolution, wavelength), making the model easily trainable for new sites and use cases. As such, regardless of the wavelength used for this validation, the STORM model will simply need to be trained from scratch on infrared speckles in order to work for C-band laser communications.

1. First STORM Model

In this paper, a first STORM model is produced and evaluated using the following three-step methodology: to begin with, the speckle pattern dataset is generated using a wave propagation simulator. Next, an ML model is built based on successful convolutional neural networks. Finally, this model is trained with 90% of the training dataset, and its performance is evaluated using the remaining 10%.

2. General Implementation Process

For the general implementation of the STORM technique on a ground station anywhere, one should follow the process shown in Fig. 4. Minus redesigning the actual ML model, which this research has made publicly available, the steps are the same as those described in the previous paragraph. After training the model on their own use case-specific datasets, a STORM user then needs to implement adequate hardware for capturing speckle patterns from their ground station. Once speckle patterns are collected on site, they are input into the trained ML model to reconstruct local turbulence profiles.

B. Assumptions

Throughout the development of the first STORM model, the following assumptions were made both for the propagation simulation and the machine learning process:

- A1 Celestial sources (starlight or LEO satellite downlink laser beams) are approximated as plane waves at the top of the atmosphere.
- A2 The optical propagation simulation applies the von Kármán spectrum within the inertial range, using an outer scale of hundreds of meters and an inner scale of the order of the millimeter. Underlying assumptions of local homogeneity and isotropy apply within this range.
- A3 Turbulence profiles are bounded by typical ranges: $[10^{-18}; 10^{-12}] \text{ m}^{2/3}$.
- A4 The altitude range of theoretically measurable turbulence is determined solely by the ground station parameters (aperture and resolution).
- A5 The intensity of the received light is assumed sufficient within the minimum exposure time to detect clearly contrasted speckle patterns.
- A6 Variations in atmospheric attenuation are considered negligible in relation to turbulence coherence time and length.
- A7 The effects of wind and pointing jitter on imaging single shot speckle patterns are considered negligible (camera integration time is short enough to avoid blur—tens of microseconds for lasercoms, hundreds of milliseconds for stars).
- A8 The success of the machine learning concept is transposable between the visible (550 nm) and IR (1550 nm) domains (albeit with a new training dataset and a newly trained model as per the pipeline in Fig. 4). It should nonetheless be noted that the altitude range and scintillation decrease as the wavelength increases.

C. Dataset Generation

The surrogate data used to train STORM's machine learning model is generated via a wave propagation simulator. As input, this simulator takes in ground station parameters and a previously generated turbulence profile. The output is a speckle pattern that could theoretically have been produced from light propagating through an atmosphere with the said turbulence profile. One simulator propagation creates one datapoint: a

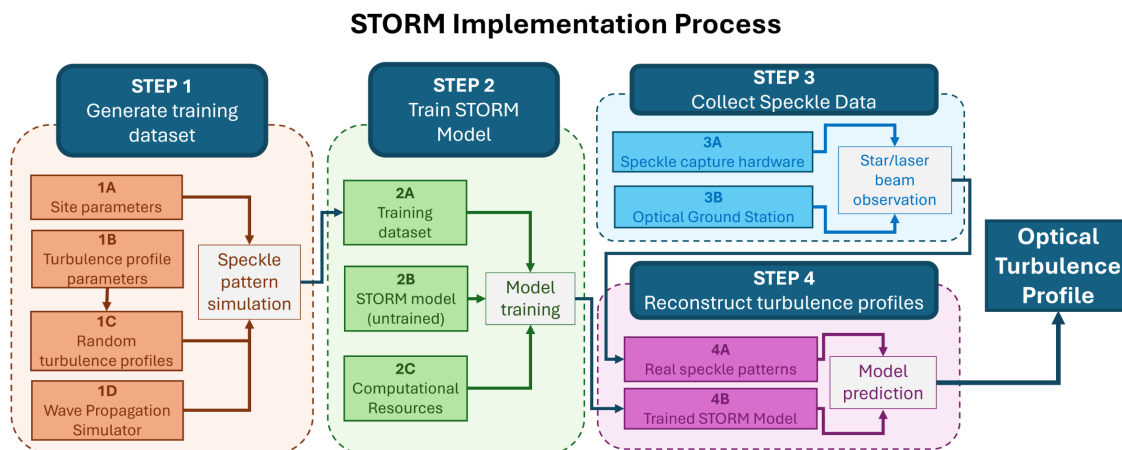


Fig. 4. Process to implement STORM on any ground station: 1. Generate a speckle pattern dataset using simulations; 2. train the STORM model on this dataset using the untrained model available from this project; 3. implement hardware on the optical ground station and collect speckle patterns; and 4 use the trained STORM model to reconstruct turbulence profiles.

Table 1. Main Simulation Parameters and Site Features

Parameter	Value
Wavelength	550 nm
Ground level	0.1 m
Maximum altitude	40 km
Negative propagation	300 m
Pointing angle	Zenith
Telescope aperture	40 cm
Secondary mirror	17.6 cm (44%)
Spider vane width	1 cm
Propagation field size	2 × 2 m (512 × 512 pix)
Output speckle size	0.4 × 0.4 m (105 × 105 pix)
Total profiles	5064
Speckles per profile	50
Total samples	253,200

speckle pattern and its tag (the turbulence profile). The generation of 250,000 datapoints constitutes the full training dataset. All the main simulation parameters are reported in Table 1.

1. Turbulence Profile Generation

The dataset generation process begins by defining a set of turbulence profiles that the model will learn to reconstruct. Making this set as complete and realistic as possible will enable the model to accurately recognize a broader range of profiles. For the purpose of these simulations, profiles are defined as turbulence integrals over eight layers of the atmosphere. The turbulence integral, J , is the integral of C_n^2 over a given portion of the atmosphere delimited by $[b_1; b_2]$:

$$J = \int_{b_1}^{b_2} C_n^2(h) dh. \quad (5)$$

The number and vertical distribution of the layers is a compromise between computational complexity and altitude resolution. In short, the more layers, the better the turbulence resolution, at the cost of higher computational demands and larger training dataset requirements. Additionally, increasing the number of layers complexifies the ML output, which, in turn, calls for even more training data. Due to these constraints, it was decided to define an eight-layer turbulence profile.

As atmospheric turbulence is generally weaker at higher altitudes and stronger in the ground layer [28], a quadratic distribution of the layers was chosen, allowing for higher vertical resolution on the lower layers. A minimum propagation distance between layers is also respected to allow phase aberrations to propagate sufficiently: the first Fresnel zone should be larger than the resolution of the next layer.

The final vertical delimitations of the eight layers are as follows: (0.1 m, 300 m, 650 m, 1 km, 2 km, 5 km, 10 km, 20 km, and 40 km).

In order to ensure a realistic and extensive set of eight-layer turbulence profiles, upper and lower bounds of turbulence strength were defined for each layer based on ranges from existing models [29,30]. Within these bounds, profiles were then randomly generated and ultimately filtered to remove those with unrealistically high Fried parameters ($r_0 > 0.5$ m) or low

Rytov indices ($\sigma_R^2 < 10^{-2}$). The Rytov index [see Eq. (9)] is an estimate of the scintillation in weak turbulence, though it defies this definition in stronger regimes.

The final distributions of the 5064 generated profiles can be seen in Fig. 5. On individual layers (see Fig. 9 in blue), the profiles show relatively logarithmically uniform distributions, whereas the total turbulence integral follows a log-normal distribution, as does consequently the Fried parameter. According to literature on measurement campaigns, true Fried parameter values do indeed follow some log-normal distribution [31–34]. It can thus be supported that the turbulence profiles generated for the dataset are realistic and cover a satisfactory spectrum of turbulence regimes.

2. Speckle Pattern Generation

Using turbulence profiles as input, speckle patterns can be generated via the wave propagation simulator described hereafter. A schematic overview of the simulation procedure is pictured in Fig. 6.

Initially defined as a uniform field representing a plane wavefront, a light source is propagated from the top of the atmosphere using the angular spectrum method (ASM) [35,36]. At the mid-altitude of each layer, a phase screen is applied. According to the layer repartition, the phase screens are inserted at (30 km, 15 km, 7.5 km, 3.5 km, 1.5 km, 825 m, 475 m, and 150 m). Each phase screen consists of filtered random noise [37] and is normalized to the turbulence strength defined by the input profile so as to replicate the effects of propagation through a turbulent layer of a given strength. Upon reaching the receiver level, a field mask is inserted to replicate the effects of the telescope spider resulting from the secondary mirror obscuration and vanes. This field is then once again propagated at a “negative” distance to mimic the generalized SCIDAR technique, which allows aberrations from surface layer turbulence to develop. Finally, the intensity of the resulting field is kept as the output speckle pattern.

One such iteration produces a single training datapoint. However, because the phase screens are generated randomly, the same turbulence profile will result in different speckle patterns for separate iterations. In this way, 50 speckle patterns were generated for each of 5064 turbulence profiles, resulting in a total of over 250,000 datapoints.

Naturally, this simulator was validated to ensure that the propagation method accurately substitutes real data. In particular, a scintillation analysis was used to compare the theoretical Rytov index to the scintillation indexes reported in [4]. Additionally, the Fried parameter was used to validate the magnitude of phase aberrations [38]. The tip-tilt aberrations have been assessed separately to ensure the subharmonics were appropriately taken into account [37]. Putting aside compromises on vertical resolution, this verification shows that exchanging real speckles for surrogate data to train the STORM ML model is a fair trade.

Generating datasets for different use cases is simply achieved by setting the simulator parameters accordingly. To favor the reproducibility of results, the parameters used for this specific dataset are listed in Table 1.

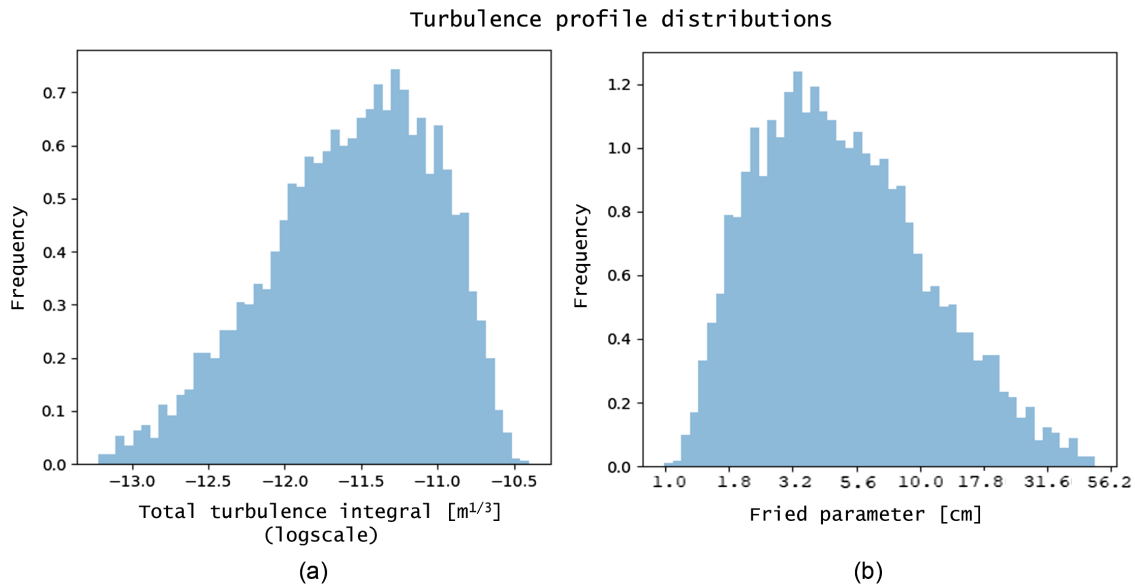


Fig. 5. Statistics on turbulence profiles of the final dataset. (a) Distribution of turbulence integral of the full profiles on the log scale. (b) Distribution of the Fried parameter for comparison.

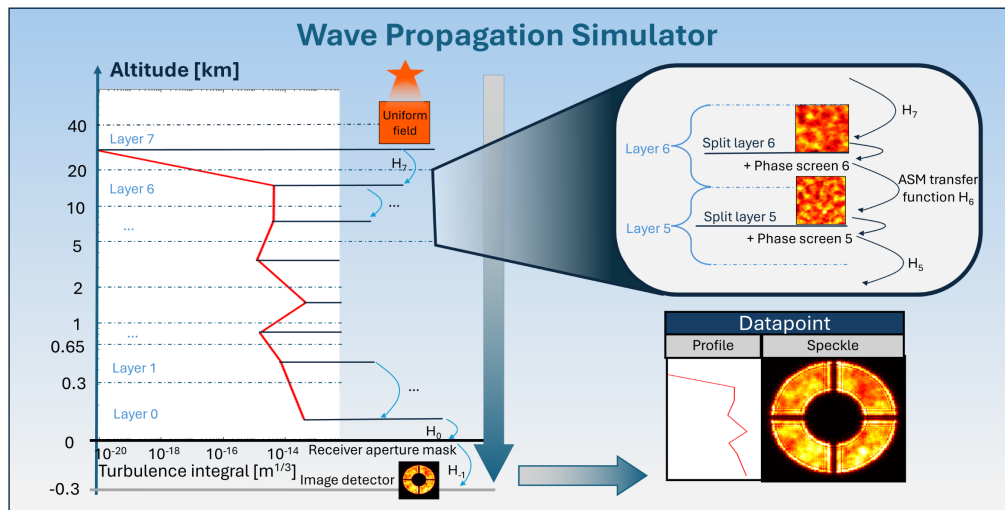


Fig. 6. Diagram of the wave propagation simulation. The randomly generated profile is input into the model along with the simulation parameters. A uniform field is then propagated through different phase screen layers from the top of the atmosphere using ASM. The final speckle pattern and input profile constitute one datapoint.

D. Machine Learning Model

To reconstruct turbulence profiles from speckle patterns, the STORM ML model is built and trained using the *speckcn2* package [39] developed for this purpose. These models take a speckle pattern as input and predict the values of the turbulence integrals per predefined layer, also referred to here as the screen tags. Training is performed by minimizing the following loss function:

$$\mathcal{L} = \sum_i (J_{i,\text{true}} - J_{i,\text{predicted}})^2, \quad (6)$$

which represents the mean-squared error (MSE) between the true and predicted screen tags. Using these predictions, additional turbulence parameters, such as the Fried parameter

[Eq. (7)], the isoplanatic angle [Eq. (8)], and the Rytov index [Eq. (9)], are computed according to definitions provided in the next subsection. While these parameters could be included directly in the loss function, their strong dependence on the dominant screen skews the model's predictive focus toward narrow edge cases, resulting in worse overall performance, as demonstrated in Section 4.B.2.

On another note, the chosen models leverage the principle of equivariance [40], which exploits the rotational symmetry inherent in the speckle patterns. This can be achieved through two approaches: (1) weak equivariance: augmenting the data with random rotations, and (2) full equivariance: using specialized models such as spherical equivariant CNNs [41,42]. Experimental results demonstrate superior performance when

fine-tuning a prebuilt model, ResNet-50 [43], on weakly equivariant data compared to training a fully equivariant model from scratch, which is also more computationally expensive. For this reason, the main results presented in this paper are based on a weakly equivariant model. Nonetheless, a comparison to a fully equivariant one is presented in Section 4.B.1.

Aside from this, all the standard choices in defining the machine learning pipeline are followed, including the use of the Adam optimizer, dropout layers, and ReLU activation functions. The final architecture and important features of the STORM model used in this study are summarized below:

1. Training data processing methods
 - Detector noise: white Gaussian noise was added to the speckles to emulate detector noise and random shot noise proportional to the square-root of the intensity.
 - Equivariance: (weak) data augmentation was achieved through random $\pi/2$ rotations.
2. Model design elements
 - Base architecture: we utilize a pre-trained ResNet50.
 - Final block: converts ResNet outputs into screen predictions using a sequence of: (i) a dropout layer, (ii) a linear layer, and (iii) a hard-sigmoid activation.
 - Optimizer: Adam optimizer.
 - Early stopping: not used.
 - Learning rate scheduler: not used.
 - Learning rate: fixed at $1e-4$.
3. Training parameters
 - Training duration: 100 epochs.
 - Batch size: 32.
 - Loss function: mean-squared error (MSE) loss.

The actual model is made available in *speckcn2* [39], with details about optional parameters and a comprehensive explanation of the model in the documentation and YAML configuration file. This configuration can be directly used as a template for the *speckcn2* model to reproduce the results of this paper but also to fine-tune the model for different tasks.

E. Validation Metrics

In the interest of quantifying the performance of the trained ML model, the output profiles are compared to the input ones via meaningful turbulence indicators: the Fried parameter (r_0), the isoplanatic angle (θ_0), and the Rytov index (σ_R^2). The exact interpretations and definitions of these parameters can be found in [4]. In this study, these indicators are bilaterally computed from the true and predicted turbulence profiles using the respective expressions below [4]

$$r_0 = \left(0.423k^2 \sec(\zeta) \sum J_i \right)^{-3/5}, \quad (7)$$

$$\theta_0 = \frac{\cos^{8/5}(\zeta)}{(2.91k^2 \int C_n^2(b)(b - b_0)^{5/3} db)^{3/5}}, \quad (8)$$

and

$$\sigma_R^2 = 2.25k^{7/6} \sec^{11/6}(\zeta) \int C_n^2(b)(b - b_0)^{5/6} db, \quad (9)$$

where k is the wavenumber, ζ is the pointing angle, and b_0 is the ground altitude. Albeit of limited accuracy for low-resolution profiles, C_n^2 values are derived from the turbulence integral using the averaged value over the layer thickness.

The predominant advantage of using these parameters as metrics becomes apparent for field tests. In fact, the Fried parameter and isoplanatic angle are measured with a DIMM, and the Rytov index can be approximated from the scintillation index measured by a scintillometer. Both DIMMs and scintillometers are commonly found on monitored sites, whereas full turbulence profiles are less common. Therefore, using these parameters as validation metrics will make field tests more straightforward to validate.

As a validation criterion, the model is deemed reasonably accurate if the above metrics computed for the predicted profiles are within a target 10% error relative to those computed for the true profiles. In fact, this value is taken such that in the case of a laser satcom link, the incurred link budget error is minimal (estimated around 1 dB) [38,44], though this estimate would benefit from a complete study.

4. RESULTS AND DISCUSSION

A. STORM Model Results

Following training, the performance of the STORM ML model was evaluated in a random subset of validation samples unused for training. The results of one such sample are shown in Fig. 7. In this example, the OTP predicted from the speckle pattern on the left overlaps almost perfectly with the true profile used to generate the speckle. Furthermore, the 49 other speckles produced from this same profile—and also unused during training—reveal that the predictions statistically follow the trend of the true profile.

Looking at the larger statistics and using the aforementioned validation metrics, overall very satisfactory results can be seen in Fig. 8. In fact, an error inferior to 10% is clearly obtained on the Fried parameter, isoplanatic angle, and Rytov index.

Besides these general results, some other interesting features can be observed:

- First, the relative error on the Fried parameter and the isoplanatic angle does not evolve proportionally to the strength of the turbulence. This shows that the model is, at first glance, equally accurate regardless of the turbulence strength and speckle sizes.
- For the Rytov index, on the other hand, the relative error increases with the overall turbulence strength. More specifically, the Rytov index is underestimated for lower values and overestimated for larger values. Hence, the model is slightly more affected by the turbulence strength when predicting scintillation.
- Finally, if the distribution of the true turbulence integral is compared to the predicted one, as in Fig. 9, the model predictions follow a log-normal distribution, whereas the true values are more uniformly distributed. In other words, the highest and lowest turbulence strengths tend to be underestimated. This phenomenon results from regularization to avoid overfitting and is a typical feature in ML. Since the model attempts to minimize loss, it will naturally avoid extremities, as they have a

Single speckle prediction vs true profile

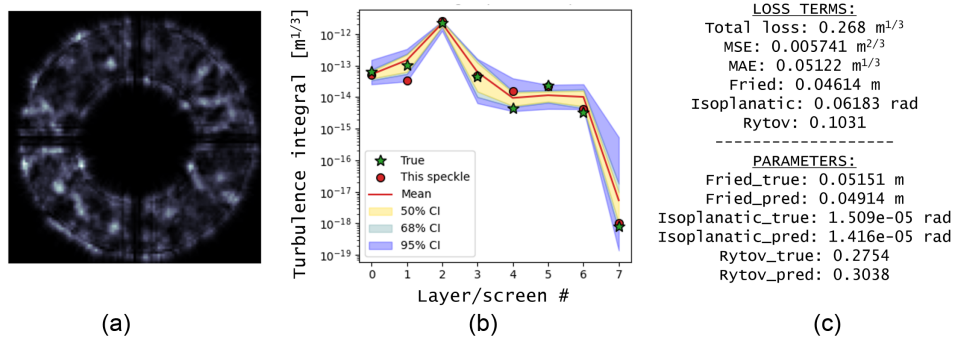


Fig. 7. Example of a good turbulence profile prediction (b) made from the speckle in (a). The statistical distribution of the predictions for the 49 other speckles generated from the same profile is also shown [shaded zones in (b)]. The true and predicted values of different parameters for this individual speckle prediction, including the total loss, are found in (c).




Fig. 8. Training results in terms of relative errors on the validation metrics: Fried parameter (a), isoplanatic angle (b), and Rytov index (c). Everything in the shaded green zone is below 10% error.

higher loss risk for strong turbulence and a lower gain for weaker turbulence.

Overall, profile trends are generally predicted correctly, with more difficulties when strong outlying layers are present or when the Rytov variance is exceptionally strong. It should be noted that the error on the predicted Fried parameter is naturally driven by errors in predictions of the strongest turbulence layers. This also results in occasionally large discrepancies between predicted and true values for weak turbulence layers. While mitigation strategies remain to be elaborated regarding such order-of-magnitude errors, this misestimation of weak turbulence is not expected to have such a significant impact on link performance prediction. This postulate nonetheless remains to be studied in further detail.

Although these results are a promising indication that single shot OTP reconstruction works, experimental campaigns will shortly be conducted to evaluate the STORM method in the field. In fact, as the training is performed on simulated data, the impact of non-modeled atmospheric phenomena is yet to be determined. In that same line of analysis, the current model was limited to an eight-layer profile due to computational constraints. The sensitivity of this discretization to the precision of the turbulence profile estimate is not investigated in this paper.

B. ML Model Trade-Offs

As expressed in Section 3, the first STORM model is a result of multiple iterations in which several ML models were traded off. The results of the main candidate models are presented here.

1. Fully Equivariant Models

To begin with, the weak equivariant model used as the main STORM model is put against a fully equivariant model based on [41,42]. It should be noted that for the fully equivariant model, the same hyperparameters optimized for the weakly equivariant model are applied to produce the best results. Table 2 shows how both models compare computationally (time- and memory-wise) when using an Nvidia v100 GPU. This shows that even if the weakly equivariant model contains ~50 times the number of parameters, it is still faster to run, due to its better optimization.

Performance is then evaluated by examining the fully equivariant model's Fried parameter predictions, as shown in Fig. 10(a). Comparing this with Fig. 8(a), it is clear that the fully equivariant model is not as effective as the weakly equivariant model. Given its lower accuracy and higher computational demands, this supports the decision to use the weakly equivariant model for STORM.

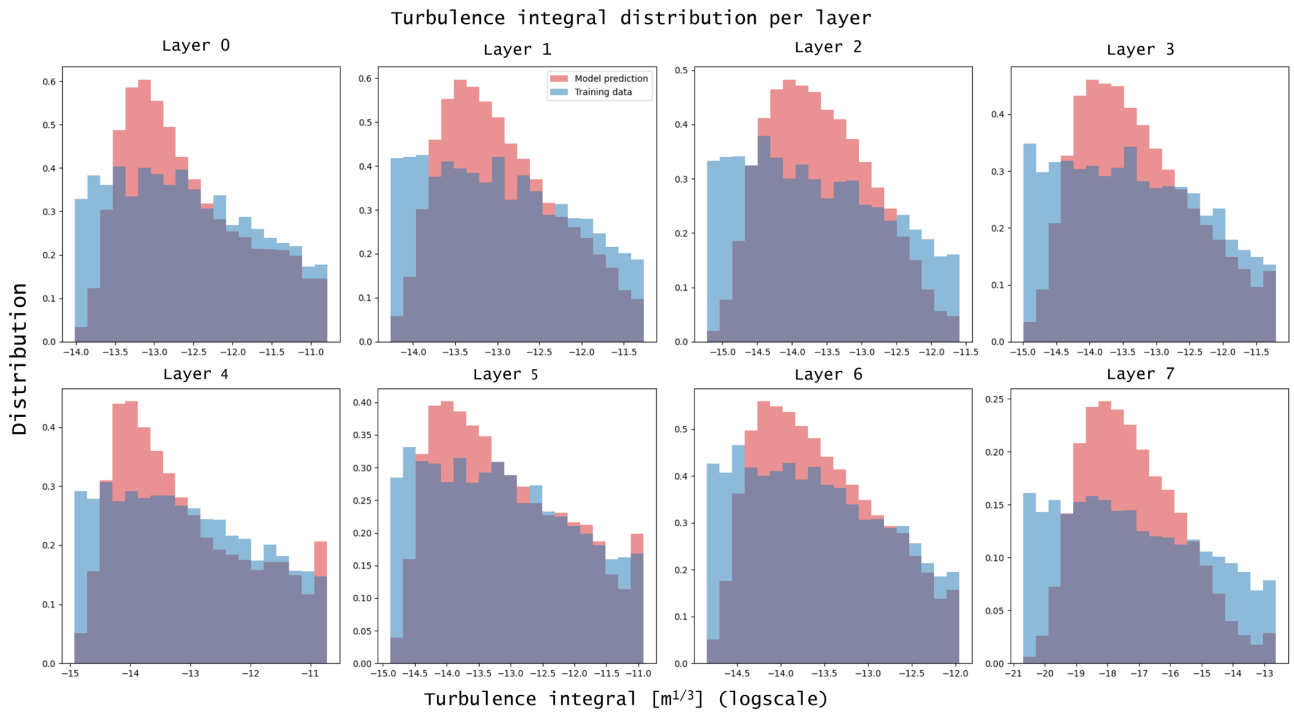


Fig. 9. Distribution of the turbulence integral (J) for each layer of the dataset profiles (0 is the lowest layer): the actual value/the tag (in blue) and the values as predicted by the trained STORM ML model (in red).

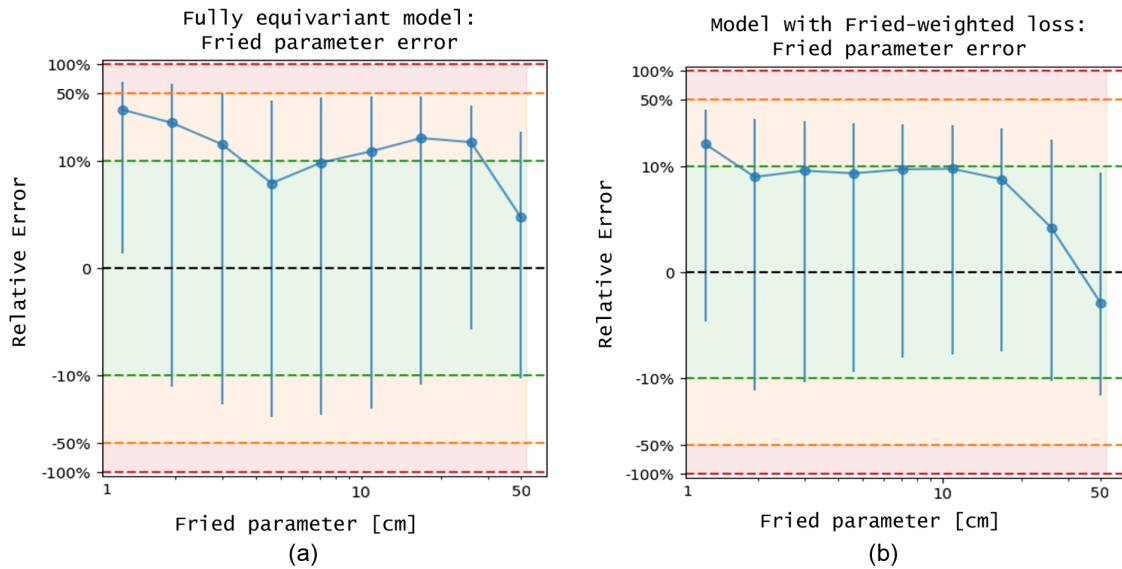


Fig. 10. (a) Relative error of the Fried parameter as a function of its value using a fully equivariant model. The data reports worse performance than the weakly equivariant model [Fig. 8(a)]. (b) Relative error of the Fried parameter as a function of its value using $\alpha = 0.5$ that weighs the value of the Fried parameter into the loss function. The results are slightly worse than those for $\alpha = 0$ [Fig. 8(a)], and the model is \sim three times slower.

2. Different Loss Functions

The second noteworthy model trade-off concerns the use of tailored loss functions that include custom metrics. In particular, instead of Eq. (6), the following loss function is considered:

$$\mathcal{L} = \alpha [\mathcal{F}_{\text{true}}(J) - \mathcal{F}_{\text{pred}}(J)]^2 + (1 - \alpha) \sum_i (J_{i,\text{true}} - J_{i,\text{predicted}})^2, \tag{10}$$

where $\mathcal{F}(J)$ is the Fried parameter as a function of the screen tags J , and $\alpha \in [0, 1]$ controls the impact of the Fried loss term.

For any $\alpha > 0$, the inference time increases by a factor of ~ 3 because the backpropagation has to go through the integral equation that defines the Fried parameter from the model prediction. More importantly, any value of $\alpha \neq 0$ degrades the performance of the model, even in the prediction of \mathcal{F} itself, as reported in Fig. 10(b). The proposed interpretation is that since \mathcal{F} is dominated by $\max_i(J_i)$, i.e., the most turbulent screen,

Table 2. Equivariant Model Comparison: Computational Efficiency

	Time Per Epoch	Number of Model Parameters
Weak equivariance (ResNet50)	~230 s/epoch	23,518,152
Full equivariance (SCNN)	~290 s/epoch	510,024

the introduction of \mathcal{F} in the loss term skews the training toward optimizing for single screens, making the whole procedure less stable and general. For this reason, the loss function defined by Eq. (6) is preferred.

5. CONCLUSION

In conclusion, STORM, or speckle-based turbulence observation and reconstruction via machine learning, is a newly proposed line-of-sight turbulence profiling technique that offers several noteworthy features.

First, this method is specially designed for the needs of laser satellite communications, an application not targeted by traditional astronomy techniques. In fact, while astronomical imaging has its own accuracy needs, the STORM concept is fueled by the need to understand the specific effects of turbulence on communications performance in order to design better mitigation systems. Hence, line-of-sight turbulence profiling using a single laser source from LEO satellites is a basic requirement covered by STORM.

Second, and perhaps most remarkably, STORM can accurately reconstruct OTPs from single speckle images of a single source. Whereas current methods require two sources or multiple images, this paper has shown that the STORM model can successfully reconstruct eight-layer profiles from a simulated single star speckle pattern with an accuracy of over 90% on the Fried parameter, isoplanatic angle, and Rytov index. In addition to ensuring the feasibility of line-of-sight profiling while tracking LEO satellites, this trait relaxes some hardware requirements concerning detector frame rates. However, reducing the number of input sources may inherently come at the cost of overall reduced accuracy.

Finally, STORM's surrogate learning-based development approach grants the model a great amount of flexibility. In fact, the use of synthetic data for training guarantees that the model can learn to reconstruct an almost infinite range of profiles. In other words, the model is not limited to any turbulence regime since it fits a solution within the turbulence ranges covered by the training dataset. This unique ability to characterize strong turbulence can be considered an advantage over current methods, which are usually limited to weak turbulence. Furthermore, this methodology can easily be generalized to any hardware setup by adapting the training data to the ground station parameters. Lastly, although the vertical resolution for these specific simulations was set to eight layers, models with higher vertical resolutions could certainly be obtained, granted higher computational resources. Therefore, the STORM methodology offers a flexible approach that can produce a turbulence profile predicting model tailored to the use case.

In this paper, STORM's theoretical OTP reconstruction capabilities have been established via simulation results. Thus, while pending field tests, it is only fair to say that STORM is a promising technology for better understanding atmospheric turbulence effects on laser satcom links, optimizing corrective technologies, and, consequently, improving the overall availability of laser satellite communications.

Funding. Netherlands eScience Center (NLESC.SSIML.2022c.021); Toegepaste en Technische Wetenschappen, NWO (COLI TKI); Nederlandse Organisatie voor Toegepast Natuurwetenschappelijk Onderzoek (COLI TKI).

Acknowledgment. The authors would like to acknowledge the financial and in-kind contributions of HTSM, TNO, and the partners of the COLI TKI. Personal thanks also go to M. Pierzyna for his good counsel and positive support in machine learning and the science of atmospheric turbulence.

Disclosures. The authors declare no conflicts of interest.

Data availability. The training dataset presented in this article is publicly available on 4TU.ResearchData in Ref. [45]. The code used to generate the STORM ML model is also available on Zenodo in Ref. [39].

REFERENCES

1. K. E. Wilson, J. R. Lesh, and T.-Y. Yan, "GOPEX: a laser uplink to the Galileo spacecraft on its way to Jupiter," *Proc. SPIE* **1866**, 138–146 (1993).
2. K. E. Wilson, J. R. Lesh, K. Araki, *et al.*, "Overview of the ground-to-orbit lasercom demonstration (GOLD)," *Proc. SPIE* **2990**, 23–30 (1997).
3. H. Hemmati, "Near-Earth laser communications," in *Near-Earth Laser Communications*, 2nd ed. (CRC Press, 2020), pp. 1–40.
4. L. C. Andrews and R. L. Phillips, *Laser Beam Propagation through Random Media* (SPIE, 2005).
5. R. Saathof, R. den Breeje, W. Klop, *et al.*, "Optical technologies for terabit/s-throughput feeder link," in *IEEE International Conference on Space Optical Systems and Applications (ICSOS)* (2017), pp. 123–129.
6. A. Tokovinin, "From differential image motion to seeing," *Publ. Astron. Soc. Pac.* **114**, 1156 (2002).
7. P. Lognoné, O. Farley, R. Griffiths, *et al.*, "Atmospheric C_n^2 profile requirements for GEO-feeder uplink AO optimisation," in *Optica Imaging Congress (3D, AOMS, COSI, ISA, pcAOP)* (Optica Publishing Group, 2024), paper PTh4E.6.
8. G. Lombardi, J. Navarrete, and M. Sarazin, "Review on atmospheric turbulence monitoring," *Proc. SPIE* **9148**, 678–689 (2014).
9. J. Vernin and F. Roddier, "Experimental determination of two-dimensional spatiotemporal power spectra of stellar light scintillation evidence for a multilayer structure of the air turbulence in the upper troposphere," *J. Op. Soc. Am.* **63**, 270–273 (1973).
10. R. Avila, J. Vernin, and E. Masciadri, "Whole atmospheric-turbulence profiling with generalized scidar," *Appl. Opt.* **36**, 7898–7905 (1997).
11. R. W. Wilson, "SLODAR: measuring optical turbulence altitude with a Shack-Hartmann wavefront sensor," *Mon. Not. R. Astron. Soc.* **337**, 103–108 (2002).
12. C. Robert, J.-M. Conan, L. M. Mugnier, *et al.*, "Near ground results of the CO-SLIDAR C_n^2 profiler," *J. Phys. Conf. Ser.* **595**, 012030 (2015).
13. J. L. Caccia, M. Azouit, and J. Vernin, "Wind and C_n^2 profiling by single-star scintillation analysis," *Appl. Opt.* **26**, 1288–1294 (1987).
14. D. Coburn, D. Garnier, and J. C. Dainty, "A single star SCIDAR system for profiling atmospheric turbulence," *Proc. SPIE* **5981**, 105–114 (2005).
15. N. Védrenne, V. Michau, C. Robert, *et al.*, " C_n^2 profile measurement from Shack-Hartmann data," *Opt. Lett.* **32**, 2659–2661 (2007).
16. X. Ran, L. Zhang, and C. Rao, "AC-SLODAR: measuring daytime normalized optical turbulence intensity distribution based on slope autocorrelation," *Mon. Not. R. Astron. Soc.* **528**, 3981–3991 (2024).
17. S. Perera, R. W. Wilson, T. Butterley, *et al.*, "SHIMM: a versatile seeing monitor for astronomy," *Mon. Not. R. Astron. Soc.* **520**, 5475–5486 (2023).

18. V. Kornilov, A. A. Tokovinin, O. Vozyakova, *et al.*, "MASS: a monitor of the vertical turbulence distribution," *Proc. SPIE* **4839**, 837–845 (2003).
19. V. Kornilov, A. Tokovinin, N. Shatsky, *et al.*, "Combined MASS–DIMM instruments for atmospheric turbulence studies," *Mon. Not. R. Astron. Soc.* **382**, 1268–1278 (2007).
20. A. Tokovinin and V. Kornilov, "Accurate seeing measurements with MASS and DIMM," *Mon. Not. R. Astron. Soc.* **381**, 1179–1189 (2007).
21. A. Tokovinin, "Measurement of turbulence profile from defocused ring images," *Mon. Not. R. Astron. Soc.* **502**, 794–808 (2021).
22. A. Y. Shikhovtsev, A. V. Kiselev, P. G. Kovadlo, *et al.*, "Estimation of astronomical seeing with neural networks at the Maidanak Observatory," *Atmosphere* **15**, 38 (2024).
23. J. Osborn, D. Guzman, F. J. d. C. Juez, *et al.*, "First on-sky results of a neural network based tomographic reconstructor: Carmen on Canary," *Proc. SPIE* **9148**, 1541–1546 (2014).
24. J. Smith, T. Fujii, J. Cranney, *et al.*, "Fried parameter estimation from single wavefront sensor image with artificial neural networks," *arXiv* (2025).
25. H. W. Shepherd, J. Osborn, R. W. Wilson, *et al.*, "Stereo-SCIDAR: optical turbulence profiling with high sensitivity using a modified SCIDAR instrument," *Mon. Not. R. Astron. Soc.* **437**, 3568–3577 (2014).
26. A. Tokovinin, "Study of the SCIDAR concept for adaptive optics applications," Tech. Rep. VLT-TRE-UNI-17416-0008 (European Southern Observatory, 1998).
27. J. Osborn, T. Butterley, D. Föhning, *et al.*, "Characterising atmospheric optical turbulence using stereo-SCIDAR," *J. Phys. Conf. Ser.* **595**, 012022 (2015).
28. L. B. Stotts and L. C. Andrews, "Improving the Hufnagel-Andrews-Phillips refractive index structure parameter model using turbulent intensity," *Opt. Express* **31**, 14265–14277 (2023).
29. F. Quatresooz and C. Oestges, " C_n^2 modeling for free-space optical communications: a review," *IEEE Access* **13**, 21279–21305 (2025).
30. R. E. Good, R. R. Beland, E. A. Murphy, *et al.*, "Atmospheric models of optical turbulence," *Proc. SPIE* **0928**, 165–186 (1988).
31. R. Racine and B. L. Ellerbroek, "Profiles of nighttime turbulence above Mauna Kea and isoplanatism extension in adaptive optics," *Proc. SPIE* **2534**, 248–257 (1995).
32. E. Masciadri, J. Stoesz, S. Hagelin, *et al.*, "Optical turbulence vertical distribution with standard and high resolution at Mt Graham," *Mon. Not. R. Astron. Soc.* **404**, 144–158 (2010).
33. V. Kornilov, B. Safonov, M. Kornilov, *et al.*, "Study on atmospheric optical turbulence above Mount Shatdzhatmaz in 2007–2013," *Publ. Astron. Soc. Pac.* **126**, 482 (2014).
34. J.-L. Prieur, G. Daigne, and R. Avila, "SCIDAR measurements at Pic du Midi," *Astron. Astrophys.* **371**, 366–377 (2001).
35. K. Khare, M. Butola, and S. Rajora, "Wave equation and diffraction of light," in *Fourier Optics and Computational Imaging*, K. Khare, M. Butola, and S. Rajora, eds. (Springer International Publishing, 2023), pp. 121–142.
36. J. D. Schmidt, *Numerical Simulation of Optical Wave Propagation with Examples in MATLAB* (SPIE, 2010).
37. M. Charnotskii, "Comparison of four techniques for turbulent phase screens simulation," *J. Op. Soc. Am. A* **37**, 738–747 (2020).
38. R. J. Noll, "Zernike polynomials and atmospheric turbulence," *J. Op. Soc. Am.* **66**, 207–211 (1976).
39. S. Ciarella, V. Azizi, L. Orozco, *et al.*, "speckcn2," Zenodo (2024), <https://doi.org/10.5281/zenodo.11447920>.
40. T. Cohen and M. Welling, "Group equivariant convolutional networks," in *Proceedings of the 33rd International Conference on Machine Learning* (PMLR, 2016), pp. 2990–2999.
41. G. Cesa, L. Lang, and M. Weiler, "A program to build E(n)-equivariant steerable CNNs," in *International Conference on Learning Representations (ICLR)* (2022).
42. M. Weiler and G. Cesa, "General E(2)-equivariant steerable CNNs," in *Advances in Neural Information Processing Systems* (Curran Associates, 2019), Vol. **32**.
43. K. He, X. Zhang, S. Ren, *et al.*, "Deep residual learning for image recognition," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (IEEE, 2016), pp. 770–778.
44. L. C. Andrews, R. L. Phillips, R. J. Sasiela, *et al.*, "Strehl ratio and scintillation theory for uplink Gaussian-beam waves: beam wander effects," *Opt. Eng.* **45**, 076001 (2006).
45. M. Arvis, R. Saathof, and S. Ciarella, "STORM training dataset," 4TU.ResearchData (2024), <https://10.4121/f41caeea-0be8-4247-9ef2-06956a3edf6b>.