

# Physics Guided Machine Learning for Wavefront Sensing on a Hybrid Optical Telescope

**Fabien R. Baron<sup>a</sup>, Douglas A. Hope<sup>b</sup>**  
**Stuart M. Jefferies<sup>a</sup>, Daniel Johns<sup>a</sup>**  
*Georgia State University<sup>a</sup>*  
*Georgia Tech Research Institute<sup>b</sup>*

## ABSTRACT

An essential requirement for Space Domain Awareness (SDA) is a robust capability to characterize resident space objects (RSO). Fulfilling this requirement requires imagery with high spatial and spectral resolution. Spatial resolution requires large apertures, while spectral resolution requires an analysis of reflected light from the RSO at all wavelengths. A new Fizeau imaging telescope design employing lightweight optics, called the Hybrid Optical Telescope (HOT), may offer a path forward in collecting high-resolution spatial imagery on RSOs, particularly in geostationary orbit and beyond – including Cislunar space. One of the favored architecture for a HOT consists in an array of small aperture primary mirrors organized concentrically, and physically supported by tensegrity. Imaging with a HOT will consequently require high-precision measurement and compensation of the wavefront errors induced by structure flexure and the turbulent atmosphere. Given the complexity of the wavefront control for a HOT, research has focused on building a physical model for turbulence which faithfully captures the spatial and temporal behavior of the atmosphere. The purpose is two-fold. First, we want to correct for turbulence using a deformable secondary mirror on timescales near the coherence time of the atmosphere. Second, we want to determine the feasibility of predictive capability for the wavefront sensor that relies on wind vectors for the turbulent layers and determining the flow pattern across adjacent M1-apertures in the HOT array.

## INTRODUCTION

Space plays important role in modern society in both the civilian and defense sector, thus maintaining and managing access to the space environment is a paramount, and requires a robust space domain awareness (SDA) in orbital regimes from low-earth orbit to Cislunar space. Achieving this requires telescopes with sufficient spatial and spectral resolution to match operational objectives for SDA which broadly defines three essential problems. The first is the detection and characterization of a faint closely space object (CSO) near a brighter source. Of particular interest is a CSO near geostationary satellite, which actually has a brightness of magnitude 9-13 approximately. The brightness of the CSO might be 10 orders of magnitude fainter than the primary geo-stationary satellite. The second problem involves characterizing a satellite based on its shape, attitude and its surface material composition, which requires both spatial and spectral resolution. Spatial resolution requires large aperture telescopes, whether on the ground or in space, while spectral resolution depends on measuring the amount of reflected light from the satellite at different wavelengths and comparing this know spectral signatures of known satellite materials.

The current generation of large aperture telescopes (6-10m) still lack sufficient spatial resolution to characterize geostationary targets. Obtaining sub-meter spatial resolution on geo targets will require 15-40m size apertures. Building filled apertures beyond this sizes is hampered by numerous engineering challenges that makes these systems cost prohibitive. While long-baseline optical interferometry has the required spatial resolution for imaging geostationary satellites, the approach lacks the sensitivity necessary to produce images with a clean background free from artifacts.

We propose a new telescope design which may offer a path forward to obtaining spatial and spectral resolution for the characterization of space objects in geostationary orbit.

## 1. THE HOT ARCHITECTURE

The Hybrid Optical Telescope (HOT) design is a Fizeau imager, which measures wavefront phases from the image plane of combined beams. Our Hybrid Optical Telescope concept (HOT) is based on the idea of the Exo-Life Imager




Fig. 1: Conceptual Hybrid Optical Telescope [8]

[8], which uses multiple apertures arranged in a ring pattern to achieve high-resolution imaging. Our investigation focuses on two configurations of a HOTA. The first is the pathfinder small-ELF (SELF), which has 15 lightweight 0.5m mirrors placed on a ring with a spatial resolution equivalent to a full 3.5m aperture. The SELF telescope is under construction by the Instituto de Astrofísica de Canarias (IAC) on Teide Mountain in Tenerife. Another HOTA configuration investigated in this work is a larger array with 37 1m apertures arranged on a ring, which yields the spatial resolution of a 15m full aperture and the light-gathering capacity of a 6m full aperture (see Figure 1).

The HOTA architecture offers a lower-cost method to build large telescopes with a sufficient spatial resolution to achieve specific SDA objects in geostationary orbit and Cislunar space. Some of the advantages of the HOTA optical configurations include [8]:

- the design is intrinsically scalable, with each sub-aperture performing as a distinct off-axis telescope.
- the overall mechanical stiffness (and mass) can be far less than what is required to produce a single full-aperture telescope.
- wavefront sensing and adaptive secondary only need to correct wavefront errors over the small sub-apertures (0.5m and 1m) versus the entire full array.
- the wavefront measurements can be obtained from the common sub-aperture foci, which is the same optical path as the science instruments.
- the HOTA can implement PSF engineering to dynamically change the distribution of the PSF sidelobes and create dark holes to enhance the detection of faint sources near a bright primary source.

### 1.1 PSF Engineering on a HOTA

The design of the HOTA enables new science related to the detection of closely spaced objects via PSF engineering. This term generally refers to the engineering of the system PSFs via modification of the light using diffractive or reflective optical elements in pupil plane prior focusing onto the detector. The objective of the PSF engineering in the HOTA is to create an area in the image free of sidelobe flux, so that fainter objects close to a bright target can be imaged. PSF engineering is accomplished by introducing known aberrations on to the sub-apertures in the array to effectively create a dark hole in the image. In the astronomical context, the goal is to image faint companions of stars (e.g. brown dwarfs or even bright exoplanets) and beat the unfavorable flux ratios of such difficult science cases. In the SSA and SDA context, one would attempt to image closely spaced objects in geo, XGEO, or Cislunar and/or obtain both spatial and spectral resolution on resident space objects.

While in principle PSF engineering could be applied to any telescope with a segmented mirror, there are limitations on the effectiveness of the technique. For example the secondary support structure on the Keck Figure 3 produces diffractive effects in the PSF which are likely to prevent successful PSF engineering. In contrast, thanks to the off-axis design of the HOTA telescopes the light from the sub-apertures does not encounter the secondary support structure and thus produces fewer sidelobes. PSF engineering provides an additional boost to the contrast which can enable imaging of closely spaced binary objects with higher contrast.




Fig. 2: Typical performance of a HOT Imager with 37 apertures. Left: expected AO-corrected image of Anik-F3. Center: diffraction-limited PSF. Right: sub-aperture phase solution necessary to compensate for the atmospheric turbulence.

PSF Engineering on a HOT is posed as an optimization problem performed by minimizing an objective function applied to the focal plane image. The minimization a metric  $L$  is performed with respect to the wavefront phase  $\phi$  across each sub-aperture expressed as linear combination of Zernike polynomials defined on the 1m sub-aperture. This objective function is the sum of two metrics: the first metric minimizes the contrast ratio of the hole intensity to the core PSF intensity, while the second metric enforces agreement between the current PSF  $h(\phi)$ , and the diffraction-limited PSF  $h_{DL}(\phi)$  outside of the core and hole areas.

$$L(\phi) = \frac{\|M_{\text{hole}} \cdot h(\phi)\|_1}{\|M_{\text{core}} \cdot h(\phi)\|_1} + \alpha \|M_{\text{other}} \cdot (h(\phi) - h_{DL}(\phi))\|_2^2 \quad (1)$$

where  $M_{\text{hole}}$ ,  $M_{\text{core}}$  and  $M_{\text{other}}$  are respectively the binary masks for the hole, core and other parts of the PSF, and  $\cdot$  denotes the Hadamard product. The value of  $\alpha$  determines the weight between contrast and PSF fidelity to the diffraction limited values during the minimization.

In practice control of the piston and tip/tilt components of the wavefront on each aperture of a HOT can be maintained by positioning and orienting the M1-mirrors, while higher order wavefront control can be controlled using specialized deformable mirrors.

A primary challenge in implementing a HOT architecture will be successfully phasing up the array by collecting for large piston errors across the array and closing an AO loop. While our colleagues at LIOM [5] are pursuing the use of reinforcement learning, as part our investigation we have focused on understanding the physics of turbulence and how to implement this into a ML-WFS.

## 2. ATMOSPHERE CORRELATIONS ON A HOT

### 2.1 Simple spatial and temporal correlations

Turbulence in the atmosphere cause random variations in the indices of refraction, which occur on spatial and temporal scales. The Kolmogorov-von Kármán Turbulence Model provides a statistical distribution on the how turbulence affects light propagation, based on isotropic assumptions and the energy cascade model, where the energy flows from large spatial structures to smaller structures. The power spectrum of phase fluctuations follows a power law, where in the inertial sub-range – which corresponds to the inner and outer scales of turbulence – the power spectral density of phase fluctuations follows a  $\kappa^{-11/3}$  behavior where  $\kappa$  is the spatial frequency. The inner scale  $l_0$  is typically a few millimeters to centimeters, below which energy dissipates as heat. The outer scale  $L_0$  can range from tens to hundreds of meters, representing the largest scale of turbulent motion. The temporal bandwidth of atmospheric turbulence is

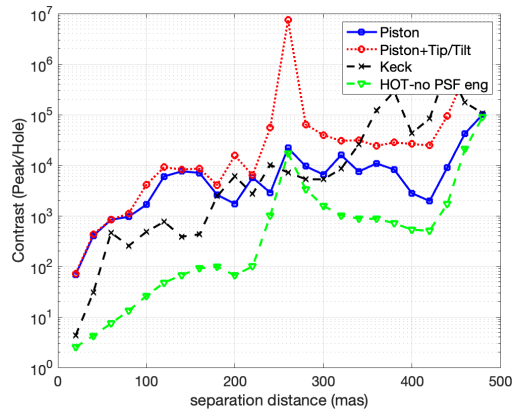


Fig. 3: Strehl and achievable contrast for a HOT with PSF engineering vs Keck full aperture (no PSF engineering possible).

described by the Greenwood frequency, which is typically around 10 – 100 Hz. Another important constraint on the temporal behavior of turbulence is described by the Taylor Frozen Flow Hypothesis, which relates the temporal power spectrum of phases fluctuations to the spatial spectrum. The frozen flow model assumes that turbulence patterns are "frozen" and simply drift across the aperture. Relative Piston phase errors between sub-apertures of the HOT are the largest source of wavefront error.

Physics-guided machine learning algorithms incorporate prior physics knowledge by encoding physical principles, equations, or constraints into the model architecture or loss functions, in order to improve the models' learning process and interpretability. Neural networks in this context can be used universal function estimators or generators.

In the case of the HOT WFS the function we need to estimate is the turbulence phase at any instant over all sub-apertures in a Hybrid optical Telescope. In the traditional application of an Shack-Hartmann WFS, the slopes of the wavefront are measured across regions of the full aperture (sub-apertures). The slopes are related to the true phases via linear matrix equation,  $A\phi = \mathbf{b}$  where  $\mathbf{b}$  are the slopes and  $\phi$  represents the atmospheric phases. The physics is limited and in fact "physics-free" wavefront correction can be achieved with reinforcement learning methods [6].

One can however intuit that for optimal correction the behavior of the atmosphere should ideally be modeled and predicted. It has been long establish that even the simple data obtained Shack-Hartmann WFS include valuable information that allows us to build operational models the atmosphere; for example the frozen flow model, where several phase screens move parallel to the ground. Even more complex effects such as boiling can be seen on a SH-WFS [10].

The Fried parameter,  $r_0$ , quantifies the area of spatial coherence of a wavefront after propagating through a turbulence regime and is computed as a uniformly weighted integral of the  $C_n^2(z)$  fluctuations. In the following we denote  $W_\phi$  the Kolmogorov-von Kármán power spectrum of the phase screens:

$$W_\phi(\kappa) = 0.033C_n^2\kappa^{-11/3} \quad \text{for } \frac{1}{L_0} < \kappa < \frac{1}{l_0} \quad (2)$$

Or first order, under frozen flow assumptions, the temporal coherence of the wavefront depends on the spatial coherence and the windspeed  $v$ ,  $\tau_0 = r_0/v$ , but more complex models of power will include boiling [2]. This paper is only concerned on demonstrating our machine learned approach for the phases is sound.

### 3. REGULARIZATION FOR PHASE RECONSTRUCTION FROM FOCAL PLANE IMAGES

We investigated how the phase may be fully reconstructed using focal plane data augmented by simple WFS information (estimates of wind speeds, number of atmospheric layers,  $l_0$ ,  $L_0$   $r_0$ ), but not directly WFS data. Reconstructing the phase from focal plane data is a well-known ill-posed and non-convex problem. It is unlike the reconstruction of the object where positivity can be imposed and an adequate starting point found from time-averaged data. In a regularized maximum likelihood context, priors on phases can be imposed either as "hard priors" or "soft priors" with are

analogous to the analysis versus synthesis regularization approaches in Compressed sensing. We describe both in the following. For convenience we will assume the reconstruction of square phase screens with known spatial covariance and a single atmospheric phase screen layer, but the discussion easily generalizes to spatio-temporal regularization of phase cubes and atmospheres with multiple layers. In the frozen flow von Kàrmàn regime, this requires knowing  $l_0$ ,  $L_0$ , and  $r_0$  and the wind velocity vector, all of which can be readily provided by an analysis of WFS data either through traditional means (temporal autocorrelation of the WFS phase estimates). In the near future these could also be probably learned from focal plane data.

### 3.1 Analysis approach

In the conventional "soft prior" or analysis approach, the likelihood of the data is minimized directly over the phase pixel values, but alongside a regularization term that attempts to enforce our expectations on the phase solution (the spatio-temporal correlations). The most likely phase  $\hat{\phi}$  is then given by:

$$\hat{\phi} = \underset{\phi \in \mathbb{R}^{N \times N}}{\operatorname{argmin}} \chi^2(\phi) + \phi^\top C_\phi \phi \quad (3)$$

where  $C_\phi$  is the covariance matrix of the phase. Because of Parseval's theorem the second term is more readily computed and minimized in Fourier space (see e.g. [3]).

$$\phi^\top W_\phi \phi = \sum_{\kappa \in \mathbb{C}^{N \times N}} \frac{|\tilde{\phi}_\kappa|^2}{W_{\phi_\kappa}} \quad (4)$$

where  $W_\phi$  is power spectral density representation of the covariance matrix in Fourier space, and  $\tilde{\phi} = \mathcal{F}\phi$ . This is exactly the same as imposing a Normal distribution of unknowns via the minimization of their square Cartesian norm  $\|\cdot\|_2^2$ , except here the Fourier transform of the phase unknowns is constrained to follow a Normal distribution, weighted by a physics-based spectro-temporal covariance. This is essentially a Tikohnov-like regularization. Provided the minimization of Eq. (3) does not fall into local minima, when using this soft regularization approach the solution at convergence is indeed likely to have the right properties, with some caveats.

One can note that the regularization terms in Eq. (4) are sums of residuals and therefore follow a  $\chi^2$  distribution with  $N^2$  degrees of freedom. Consequently they obey a normal distribution as  $\mathcal{N}(N^2, 2N^2)$  that a more constraining regularization could impose. Second, in practice solving these minimization problems with gradient-based algorithms leads to solutions stuck in local minima of  $\chi^2(\phi)$  that are far from having the right spatio-temporal structures. A strategy can probably be established to over-emphasize the regularization weight first, then slowly decrease it back.

### 3.2 Synthesis approach

The second approach to regularization is the "hard prior" or the "synthesis" approach. The phase is expressed itself as a function of more fundamental latent variables that are sought instead of the phase pixels. A basic example is given by using the latent variables  $z$  provided by the classic phase screen generator:

$$\phi = \operatorname{Re} \{ \mathcal{F}^{-1} W_\phi z \} \quad (5)$$

so that the best phase is found first by minimizing the  $\chi^2$  over  $z$ :

$$\hat{\phi} = \operatorname{Re} \{ \mathcal{F}^{-1} W_\phi \hat{z} \} \quad \text{s.t.} \quad \hat{z} = \underset{z \in \mathbb{C}^{N \times N}}{\operatorname{argmin}} \chi^2(\operatorname{Re} \{ \mathcal{F}^{-1} W_\phi z \}) \quad (6)$$

Note that Eq. (6) is not strictly speaking compressed sensing. With the default FFT approach we would be using more latent variables ( $N \times N$  complex numbers) than there are phase unknowns ( $N \times N$  reals). Since phase screens are real-valued, in practice we can employ the inverse real-valued FFT, so that the number of complex latent variables to seek is  $(N \div 2 + 1) \times N$ . Better, one can also use an inverse real-valued non-equispaced Fourier transform – or a phase screen generation method such as the sparse spectrum method [4] – to minimize the number of latent coefficients, and in the processing conform better to the von Kàrmàn or Kolmogorov structure functions. When using a so-called type 3 non-uniform FFT to map values from the Fourier plane to the phase screen support, one can choose to estimate either only the latent coefficients that will affect the phase stadium or a larger area in the hope of making predictions. In practice with sparse sampling, 500 coefficients are sufficient to recover  $1024 \times 1024$  phase screens. We underline that this small number of parameters opens the possibility of solving the problem with e.g. Monte-Carlo Markov Chains, though we did not attempt it and only implemented optimization via the gradient-based VMLMB algorithm [11]. We present our results in section 4.1.

### 3.3 Synthesis approach with normalizing flows

The main drawback of the previous approach is the lack of flexibility imposed by the usage of a fixed spatial phase covariance  $W_\phi$ . One way to generalize the method could be using Gaussian processes. Instead we explored an even more expressive framework, that of generative normalizing flow networks. Normalizing flows, or invertible networks, are neural networks that learn how to construct a bijection between any wanted distribution and the normal distribution via a discrete numbers of transforms. In our work, a GLOW [7] network  $G$  is trained on a large number of phase screens generated from diverse sets of von Kármán parameters  $l_0$ ,  $L_0$  and  $r_0$ , and learns the network weights necessary to transform from phase screen space to a latent variable space, where the latent variable  $z$  are normally distributed. We set up our reconstruction problem to use frozen flow to limit the training to the learning of 2D phase screens, but in the future one could imagine training the network on translating and boiling screens. One can immediately see that for fixed von Kármán parameters,  $G$  is actually learning how to replicate Eq. (5). But the strength of the method is that  $G$  is learning screens across all possible values of von Kármán parameters. Since  $G$  is invertible,  $G^{-1}$  acts as a generative model: when applied to a normally-distributed 2D array of the correct size, it will generate phase screens with a von Kármán structure.

$$\hat{\phi} = G^{-1}(z) \quad s.t. \quad \hat{z} = \underset{z \in \mathbb{R}^{N \times N}}{\operatorname{argmin}} \chi^2(G^{-1}(z)) \quad (7)$$

Normalizing flows scale fairly well with the size of the phase screens to learn, though not as well as the sparse sampling method of phase generation does in the previous section. The number of unknowns to learn during the training phase is very significant (millions), since the network is learning across a wide parameter space. Once trained, the number of unknowns is the same as in the analysis method. We note that our implementation is still relatively experimental at this stage.

## 4. RESULTS & DISCUSSION

### 4.1 Simulated phase reconstruction: analysis vs synthesis

We simulated observations of focal plane speckle data above a individual HOT subaperture. Here we were concerned with optimally recovering the phases above each subaperture. These reconstructed phase will differ from the truth by a global piston, but this ambiguity will be lifted when using the full HOT aperture in Fizeau combination since we can recover the differential pistons. And under the assumption of frozen flow and phase boiling, the reconstructed screens could be reused when they flow from one telescope to the next.

We generated a sequence of one hundred 5ms-frames of the satellite model object presented in our previous AMOS paper [1] for a unit telescope of  $1m$ , atmosphere with  $l_0 = 3cm$ ,  $L_0 = 20m$  and  $r_0 = 9cm$ , and wind speed of  $8m/s$  (these values were chosen partially so that the figure is representative). The atmospheric turbulence is really benign compared to the typical speckle imaging conditions on a larger telescope, which highlights one of the strengths of the HOT architecture, the favorable  $D/r_0$  on the subapertures. The image size to be reconstructed in these simulations were  $128 \times 128$  pixels, and the phase stadium support was  $640 \times 128$  pixels.

Figure 4 presents our results. In the analysis approach we directly reconstruct the  $\sim 36000$  pixels of the phase stadium. The gradient descent algorithm is initialized with starting phase stadium inferred from the centroid method. In this method for each frame we find the centroid, which provides initial guesses for the tip and tilt at the corresponding location in the phase stadium. The optimal starting phase screen stadium is then given as the least-square solution to the problem of observing the centroids for all the frames. The algorithm quickly converges toward a local minimum solution which has relatively correct medium and long frequencies. But it also displays numerous artefacts, and – when zoomed-in – very few high spatial frequency Kolmogorov structures. The artefacts are mostly due to local minima arising from wrapping of the phase.

In the synthesis approach we reconstructed 500 coefficients which in turn were used to generate the phase stadium solution. The simulation was run in supervised mode, i.e. assuming the truth values known for the parameters  $l_0$ ,  $L_0$ ,  $r_0$  and wind speed and direction. For the initial guess we inverted the initial phase guess from centroids, using the inverse non-equispaced Fourier transform coefficients. Convergence to the solution in the synthesis approach is much slower than in the analysis approach in terms of number of iterations. But, as expected, the solutions for each iteration comport with the von Kármán power spectral behavior. We note we could initialize the coefficients with random initial values and yet still converge to a solution close to the truth. In the synthesis approach we found the

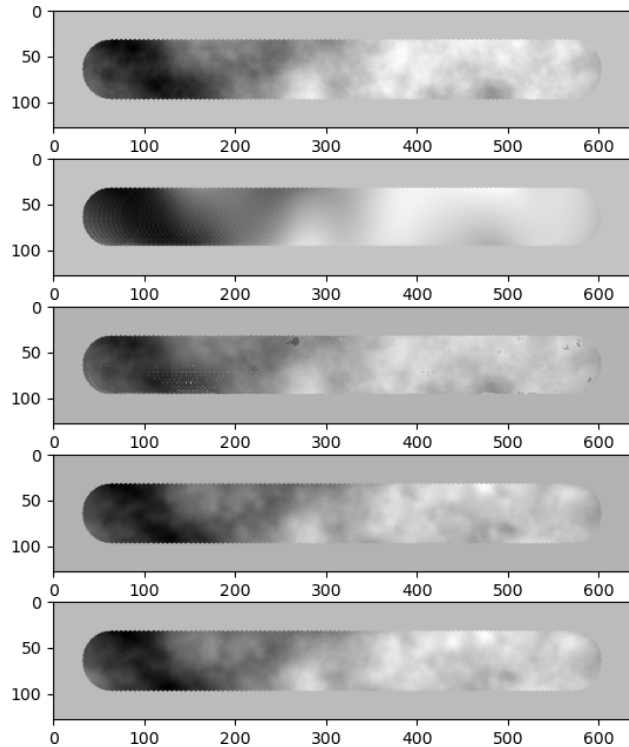


Fig. 4: Phase stadia in our reconstructions. From top to bottom: true phase, starting phase estimated from centroids, reconstruction obtained with the analysis method, reconstruction obtained with the synthesis approach, and reconstruction obtained with regularization using a GLOW network.

phase to be naturally unwrapped, since the correlations are enforcing neighboring pixels to co-vary strongly on short spatial scales.

In the invertible network approach, we trained our network on a set of phase stadia with von Kármán structures, with a range of  $l_0$  between 1cm and 5cm and  $L_0$  between 10m and 100m. The solution is very close to what the synthesis approach provides, but with the added flexibility that the parameters are unsupervised. Convergence is also much faster, making it the best solution overall.

Finally we must note that our experiments with simulations with much larger numbers of pixels exacerbate all the problems and behavior described above: local minima for analysis, very slow convergence for synthesis, but comparatively fast and parameter-free machine learned reconstruction.

## 5. FUTURE WORK

To achieve imaging with a HOT that can be of interest to SSA, we have developed a software framework that can tackle WFS cophasing and deconvolution problems. This software base will be used to train additional machine learning codes such as KalmanNet [9]. HOT pupil configurations present a large number of adjacent telescopes, so phase screen reconstruction can exploit some of the spatially and temporally correlated structures as the frozen flow carries out the phase from one telescope to the next. The work presented here was concerned with recovering the spatial structure of the phase, but this will be extended to full spatio-temporal regularization – including both the wind-driven motion as well as phase screen boiling terms.

## 6. ACKNOWLEDGEMENTS

This material is based upon work supported by the Air Force Office of Scientific Research under award number FA9550-23-1-0536. Any opinions, findings, conclusions, or recommendations expressed in this material are those of

the authors and do not necessarily reflect the views of the United States Air Force.

- [1] F. Baron, S. Jefferies, D. Shcherbik, R. Hall, D. Johns, and D. Hope. Hyper-Spectral Speckle Imaging of Resolved Targets. In S. Ryan, editor, *Proceedings of the Advanced Maui Optical and Space Surveillance (AMOS) Technologies Conference*, page 46, September 2023.
- [2] A. Berdja and J. Borgnino. Modelling the optical turbulence boiling and its effect on finite-exposure differential image motion. *Monthly Notices of the Royal Astronomical Society*, 378(3):1177–1186, July 2007.
- [3] L. Blanco and L. M. Mugnier. Marginal blind deconvolution of adaptive optics retinal images. *Optics Express*, 19(23):23227, November 2011.
- [4] Mikhail Charnotskii. Comparison of four techniques for turbulent phase screens simulation. *J. Opt. Soc. Am. A*, 37(5):738–747, May 2020.
- [5] Justin Fletcher and Peter Sadowski. Towards jointly learned control policies and image recovery for distributed aperture telescopes. In Genshe Chen and Khanh D. Pham, editors, *Sensors and Systems for Space Applications XV*, volume 12121, page 1212108. International Society for Optics and Photonics, SPIE, 2022.
- [6] Yann Gutierrez, Johan Mazoyer, Laurent M. Mugnier, Olivier Herscovici-Schiller, and Baptiste Abeloos. Image-based wavefront correction using model-free reinforcement learning. *Optics Express*, 32(18):31247, August 2024.
- [7] Durk P Kingma and Prafulla Dhariwal. Glow: Generative flow with invertible 1x1 convolutions. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018.
- [8] J. R. Kuhn, S. V. Berdyugina, J.-F. Capsal, M. Gedig, M. Langlois, G. Moretto, and K. Thetraphi. The Exo-Life Finder Telescope (ELF): design and beam synthesis concepts. In Heather K. Marshall and Jason Spyromilio, editors, *Ground-based and Airborne Telescopes VII*, volume 10700, page 1070015. International Society for Optics and Photonics, SPIE, 2018.
- [9] Guy Revach, Nir Shlezinger, Xiaoyong Ni, Adria Lopez Escoriza, Ruud J. G. van Sloun, and Yonina C. Eldar. KalmanNet: Neural Network Aided Kalman Filtering for Partially Known Dynamics. *IEEE Transactions on Signal Processing*, 70:1532–1547, January 2022.
- [10] David Saint-Jacques and John E. Baldwin. Taylor’s hypothesis: good for nuts. In Pierre Léna and Andreas Quirrenbach, editors, *Interferometry in Optical Astronomy*, volume 4006 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, pages 951–962, July 2000.
- [11] Éric Thiébaud. Optimization issues in blind deconvolution algorithms. In Jean-Luc Starck and Fionn D. Murtagh, editors, *Astronomical Data Analysis II*, volume 4847 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, pages 174–183, December 2002.