Predictive Learn and Apply: MAVIS application - Learn

Hao Zhang[a], Jesse Cranney[a], Nicolas Doucet[a], Yuxi Hong[c], Damien Gratadour[a,b], Hatem Ltaief[c], David Keyes[c], and François Rigaut[a]

[a]Research School of Astronomy & Astrophysics,
College of Science,
Australian National University, Australia
[b]LESIA, Observatoire de Paris,
Université PSL, Sorbonne Université, Université de Paris, Sorbonne Paris Cité,
CNRS France
[c]Division of Computer, Electrical, and Mathematical Sciences and Engineering,
Extreme Computing Research Center,
King Abdullah University of Science and Technology, KSA

ABSTRACT

The Learn and Apply reconstruction scheme uses the knowledge of atmospheric turbulence to generate a tomographic reconstructor, and its performance is enhanced by the real-time identification of the atmosphere and the wind profile. In this paper we propose a turbulence profiling method that is driven by the atmospheric model. The vertical intensity distribution of turbulence, wind speed and wind direction can be simultaneously estimated from the Laser Guide Star measurements. We introduce the implementation of such a method on a GPU accelerated non-linear least-squares solver, which significantly increases the computation efficiency. Finally, we demonstrate the convergence quality from numerically generated telemetry, we provide the end-to-end Adaptive Optics simulation results, and we highlight a time-to-solution analysis, all based on the MAVIS system design.

Keywords: adaptive optics, turbulence profiling, real-time processing, MAVIS, Learn and Apply, Stochastic Levenberg-Marquardt

1. INTRODUCTION

The MCAO Assisted Visible Imager and Spectrograph (MAVIS) is a proposed instrument for the European Southern Observatory's (ESO) Very Large Telescope (VLT) Adaptive Optics Facility (AOF).[1,2] It aims at delivering near-diffraction limited images in the visible band utilising eight Laser Guide Stars (LGS) and three Deformable Mirrors (DM). The high performance requirements lead to a tight error budget, and an optimised control scheme is one of the most critical requirements.

Since 2019, the minimum mean-square error (MMSE) method has been applied to test and predict the performance of the MAVIS system (see Agapito *et al.*[3]). Beyond that, the two-step Learn and Apply (L&A) reconstruction scheme has also been explored as one of the candidates for the MAVIS project.[4] Its first step records a set of wavefront measurements to identify the atmospheric and instrumental parameters (so-called *Learn*) and the second step computes the tomographic reconstructor (so-called *Apply*) using a minimum-phase-variance approach. Recently, it has been extended to include a predictive estimation to further optimize for the residual wavefront error caused by the servo-lag, which is also known as the predictive L&A (pL&A) method. The readers may refer to the companion paper[5] for more detailed discussion regarding to the implementation of predictive control, which is also called predictive *Apply* in the pL&A scheme. In this paper we focus on the identification of relevant atmospheric turbulence and wind profile, i.e. predictive *Learn*.

The calculation of a predictive L&A reconstructor is based on the *frozen flow* assumption,[6] and requires the knowledge of turbulence strength (C_n^2 distribution), wind speed and wind direction. Due to the rapidly evolving

Send correspondence to Hao Zhang, E-mail: hao.zhang1@anu.edu.au

nature of the atmospheric turbulence, it is highly desirable that such information is updated in real-time. During the past few decades, a variety of methods have been proposed to handle this task. Classical SLOpe Detection And Ranging (SLODAR)[7,8] uses a peak detection algorithm to identify the layer altitude and relative strengths of atmospheric turbulence directly from Shack-Hartmann WaveFront Sensor (SHWFS) measurements, and is extended to Laser Guide Star (LGS) cases by Gilles & Ellerbroek in 2010.[9] The core idea of SLODAR is to use the cross-correlation of the measurements from different LGSs to identify the contribution of turbulence as a function of altitude. Wind profiling has been explored by Wang *et al.* in 2008,[10] Sivo *et al.* in 2018,[11] and Laidlaw *et al.* in 2020,[12] respectively. Their methods are inspired by the classical SLODAR, with one or more additional steps involved to identify the wind profile using iterative fitting techniques. The computational load of these methods can be extremely heavy when the AO system becomes more complicated, so they are usually performed off-line, which does not fit the goal of predictive control. Other SLODAR or SLODAR-based methods focus on the wind-vector norm, which does not contain the wind direction information, so cannot be directly used in the calculation of predictive reconstructors.[13]

In this paper, we propose a predictive *Learn* method that is capable of estimating all required atmospheric turbulence parameters for predictive tomographic reconstruction simultaneously in one step. Based on the von Kármán turbulence model[14] and the frozen flow assumption, an analytical time-delayed measurement covariance matrix can be derived from the AO system geometry. On the other hand, such a covariance matrix can be numerically calculated using the pseudo open-loop telemetry. This motivates a least squares optimisation problem, as defined in Eq. (3). A non-linear solver is introduced to iteratively fit the analytical model to the numerical covariance matrix. This process typically involves large matrix calculations that are highly parallelizable on GPU, which significantly boosts the computational efficiency.

This paper is organized as follows. In Sec. 2.1, we derive the predictive *Learn* scheme used to identify the atmospheric turbulence parameters from the von Kármán model. We also show that this scheme is capable of being further extended to estimate other useful parameters in the AO system, including the WFS translational misregistration. In Sec. 2.2, we demonstrate the implementation of the above mentioned GPU accelerated non-linear least-squares solver. In Sec. 3, we present simulation results regarding to the convergence quality of numerical data, as well as end-to-end AO simulations using the predictive *learn* results, and the time to solution analysis in the context of MAVIS. Finally, we present our conclusions in Sec. 4.

2. PREDICTIVE *LEARN* SCHEME

The L&A method is inspired by the conventional phase-based MMSE algorithm which is detailed in Le Roux *et al.*[15] L&A optimizes the DM command for a series of target directions (which are also known as "Truth Sensors") by projecting the actual WFS measurements to those directions. The projection is based on the von Kármán turbulence model and a covariance matrix is used to build the link between different directions in the metapupil, as shown in Vidal *et al.*[4] In this paper we assume that the atmosphere turbulence can be modeled by a number of vertically discrete layers, and the relative intensity of each layer is given by the C_n^2 fractions. In the L&A scheme, the covariance matrix is directly derived in the slope space and is usually block structured due to the $x-y$ arrangement slope measurements. Every element in the matrix represents a pair of sub-apertures and its value is a function of their distance and the atmospheric profile. Our companion paper[5] provides an informative discussion of the use of covariance matrices in the predictive *Apply* step, and a detailed expression of the functions can be found in Gendron *et al.*[16] For readability, we use a shorthand notation $F\{\cdot\}$ for the analytical model. Denoting the distance between i^{th} and j^{th} sub-apertures as r_{ij}, the slope covariance between these two sub-apertures is given by:

$$C_s(i,j) = F\{r_{ij}, \rho\}, \tag{1}$$

where ρ is the concatenation of atmospheric turbulence parameters. Note that some of the atmospheric turbulence related parameters (e.g. the Fried parameter, the outer scale and altitude for different layers) are not taken into consideration in the predictive *Learn* proposed in this paper as they are either possible to be identified with other methods or not required to be updated in real-time.

On the other hand, the same matrix can be numerically estimated from the AO telemetry data by:

$$\hat{C}_s = \frac{1}{N}\sum_{k=1}^{N} s_k s_k^T, \tag{2}$$

where $\hat{\cdot}$ denotes estimation, and s_k is the zero-mean slope vector acquired by concatenating all measurements from WFSs at time k. By definition, the estimation in Eq. (2) approaches the true value as $N \to \infty$, so C_s in practice always suffers from convergence noise.

Now the atmospheric turbulence profiling takes the form of a least-squares optimization:

$$\hat{\rho} = \arg\min_{\rho \in \mathbb{R}^d} \|\hat{C}_s - C_s(\rho)\|^2, \tag{3}$$

where the $\|\cdot\|^2$ operator is the Frobenius norm. The numerical $\hat{C}_s$ is used as the input dataset and an optimizer is applied to find the solution. In the restricted case where only the C_n^2 profile is to be identified, the system is linear and therefore the solution can be found using a standard linear least-squares solver.

2.1 Prediction in *Learn*

In every frame of an AO loop, the WFS measurement and the subsequent corresponding DM command cannot happen simultaneously. There is always a lag due to the time required for WFS readout, DM command calculation and application. The evolution of atmospheric turbulence during this period results in a residual phase error, which can be mitigated using a predictive control scheme. The frozen flow assumption proposes that the temporal evolution of atmospheric turbulence can be regarded as a pure translational shifting, which allows us to generate the time shifted covariance matrix by including a wind shift term in Eq. (1):

$$C_s(i, j, \delta t) = F\{\|\overrightarrow{r}_{ij} + \overrightarrow{r}_w\|, \rho\}, \tag{4}$$

where $\vec{r}_w$ has components $(v_x \delta t, v_y \delta t)$.

The corresponding numerical covariance is then defined as:

$$\hat{C}_s(\delta t) = \frac{1}{N}\sum_{k=1}^{N} s_k s_{k+\delta t}^T. \tag{5}$$

Now ρ becomes the concatenation of C_n^2, v_x, v_y for each layer, the least-squares solver defined in Eq. (3) can be applied so that all turbulence related parameters can be estimated simultaneously. Additionally, the model described by Eq. (4) can be further extended to fit other parameters that are used in AO control. As an example, the translational misregistration of WFSs, including displacement in (x, y) direction, and rotation against the pupil center, appear as additional terms in the overall distance between sub-apertures.

The least squares fitting process used here can be time consuming. Though the total number of parameters to be identified is relatively small, the input dataset is quite large. For the MAVIS system where 8 LGSs are used with (40×40) sub-apertures per LGS, the size of the covariance matrix yields $(20k \times 20k)$ and every element of the matrix is considered as one term of the cost function, Eq. (3). A variety of techniques can be applied to speed up the process, and in Sec. 2.2 we focus on a recently proposed non-linear solver.

2.2 Stochastic Levenberg-Marquardt Algorithm in *Learn*

Previous works employed the conventional Levenberg-Marquardt (LM) algorithm[4] to solve the system defined in Eq. (3). This algorithm is a second-order optimizer that is widely used in the machine learning community for solving non-linear systems. As a variant of the Gauss-Newton method, LM uses a regularization factor μ along with the Hessian H of the parameters to find the minimizer. The typical update h_{lm} of the parameters ρ in one iteration is the solution of:

$$(J_k^T J_k + \mu I) h_{lm} = -g_k, \tag{6}$$

where J_k is the Jacobian matrix of $C_s(\rho)$ at k^{th} iteration, $J_k^T J_k$ is the approximated H, I is the identity matrix, and g is the gradient of $F(\rho)$. More detailed discussion on the conventional LM algorithm can be found in [17,18]. It is worth noting that this algorithm operates on the full dataset (i.e., the numerical covariance matrix as defined in Eq. (5)), therefore the computational power required is quite large and usually takes up to hundreds of seconds to find the solution for large scale problems such as the MAVIS MCAO system.

Stochastic first-order algorithms such as Stochastic Gradient Descent have been widely used in the last decade since the rise of machine learning and deep learning. These algorithms use only a subset of the dataset to approximate the gradient, thus accelerating the computation per iteration. This approach inspires the idea of the Stochastic LM (SLM) method, which is further justified by the apparent redundancy within the covariance matrices. Sub-aperture pairs that have the same (x, y) displacement in space have exactly the same expected covariance value, which leads to a large number of duplicated elements in those matrices. Thus an approximated gradient can be derived from a random subset of the dataset. Eq. (6) can also be used to describe the SLM algorithm update except that g and J are calculated using a random subset of elements in the covariance matrix.

Recently, Hong *et al.*[19] demonstrated the SLM method showing a 50-fold acceleration compared to the conventional LM algorithm with noiseless analytical datasets.

3. SIMULATIONS

In this section, we present the predictive *Learn* performance using simulation results.

3.1 Implementation Details

Our simulation consists of three steps.

1. Firstly, a large set of WFS measured slopes (so-called slope buffer) are collected from the end-to-end numerical AO simulator, COMPASS.[20,21] The AO system configuration chosen is based on the designed parameters for the MAVIS system, as detailed in Sec. 3.2. The total amount of slope buffer is equivalent to 1,000 seconds of on-sky observation time. It is assumed that the atmospheric turbulence configuration (i.e., C_n^2 distribution, wind speed, wind direction, *etc.*) remains the same during this period.

2. Subsets of the slope buffer are used to compute different numerical covariance matrices as input. The proposed predictive *Learn* and the SLM algorithm are implemented in the Supervisor module with HIgh Performance Software (SHIPS), which is a software platform for the soft real-time component of tomographic AO systems.

   SLM algorithms are implemented by using CPU while the most time-consuming operations (calculating $F(\rho)$ and calculating Hessian and gradient) are offloaded to GPU accelerators. The CUDA programming model is used to implement these kernels introduced in Sec. 2.2.

   We apply predictive *Learn* on these numerical covariance matrices.

3. At last, the results from step 2 are used to generate predictive *Apply* reconstructors. We run the end-to-end AO simulations again in COMPASS with these reconstructors as a demonstration of predictive *Learn* performance.

3.2 AO System Configuration

The MCAO system configuration used in this paper follows the MAVIS system design.[2] Here we list only the key parameters that are directly related to the predictive *Learn* process.

The telescope diameter is 8 meter, the AO system operates at 1 kHz, and the overall seeing at 500 nm is 0.8 arcsec (at zenith). 8 LGS SHWFS located at 17.5 arcsec radius are used, with 40×40 sub-apertures for each SHWFS. The readout noise is set to 0.5 e^-. Therefore the total number of slope measurements per frame is 19,584. A single Gaussian sodium distribution is assumed with LGS altitude at 90 km and flux 88 photon/sub-aperture/ms. The atmospheric profile used in all simulations is listed in Table 1. The predictive *Apply* reconstructor used in step 3 is optimised for 30 arcsec Science Field of View (FoV).

Table 1. Cn2 Profile and wind velocities.

Altitude [m]	Cn2 Fraction	Wind Velocity [m/s]	Outer scale [m]
30	0.59	6.6∠0°	25
140	0.02	5.9∠10°	25
281	0.04	5.1∠20°	25
562	0.06	4.5∠25°	25
1125	0.01	5.1∠0°	25
2250	0.05	8.3∠10°	25
4500	0.09	16.3∠20°	25
7750	0.04	30.2∠25°	25
11000	0.05	34.3∠0°	25
14000	0.05	17.5∠10°	25

3.3 Results

Given the AO system and atmosphere configuration mentioned above, the performance of predictive *Learn* is explored with regard to a variety of parameter settings as illustrated in Table 2.

Usually the convergence quality of the numerical covariance matrix is directly related to T but also depends on the turbulence condition during this period, which is marked by Ind in our simulation. For each T selection, 4 numerical covariance matrices with different Ind are generated while ensuring no overlapping in the slope buffer. Due to the similarity of turbulence between adjacent samples in time (e.g., s_k and s_{k+1}), it is not necessary to collect 100% of the measurements to generate the numerical covariance matrix (lowering the communication overheads). This leads to the use of the decimation rate r (e.g. $r = 10\%$ means taking only 1 measurement out of 10 continuously collected measurements). In the following simulations, we keep $r = 10\%$ as it reduces the computational load while keeping good Signal-to-Noise ratio. dT is usually chosen based on the frame rate of AO system and the prior experience of turbulence on site. In this paper we set $dT = 0.02$ s, $dp = 1\%$, and $\mu = 3 \times 10^{-5}$ in all simulations.

Table 2. Parameters explored in simulations.

Parameter List		Notation/unit	Definition
Buffer Parameters	Integration Time	T / s	Length of numerical buffer in seconds
	Decimation Rate	r / %	Percentage of slope buffer to be used in covariance matrix calculation
	Buffer Index	Ind / N/A	ID of numerical buffer marked by starting indices in total buffer
	Time Delay	dT / s	Time delay for the calculation of covariance matrices
SLM Parameters	Iteration Time	T_{max} / s	Max time allowed for SLM algorithm
	Data Percentage	dp / %	Percentage of data used in each SLM iteration
	Regularization Factor	μ / N/A	Fixed value used in Eq. 6

Figure 1 shows the long exposure Strehl Ratio averaged over Science FoV (F-LESR) in end-to-end simulations with tomographic reconstructors computed from predictive *Learn* results. The LESR standard deviation from various Ind is given by the error bar. For comparison, we simulated the same system using exact knowledge of the atmospheric parameters with a reconstructor generated by both classical L&A, and pL&A. These achieved 25.7% and 32.4% F-LESR at 550 nm respectively.

It is shown that the predictive *Learn* results are expected to achieve 32.1% F-LESR when using at least $T = 100$ s slope buffer and $T_{max} \geqslant 10$ s*, the standard deviation (STD) is 0.4% F-LESR, which suggests that the predictive *Learn* is quite robust in this case. Compared to the reference pL&A F-LESR, the wavefront RMS error introduced in the predictive *Learn* process is approximately 7.5 nm.

*In this paper, the T_{max} is based on the use of a single NVIDIA Quadro RTX 5000 GPU.

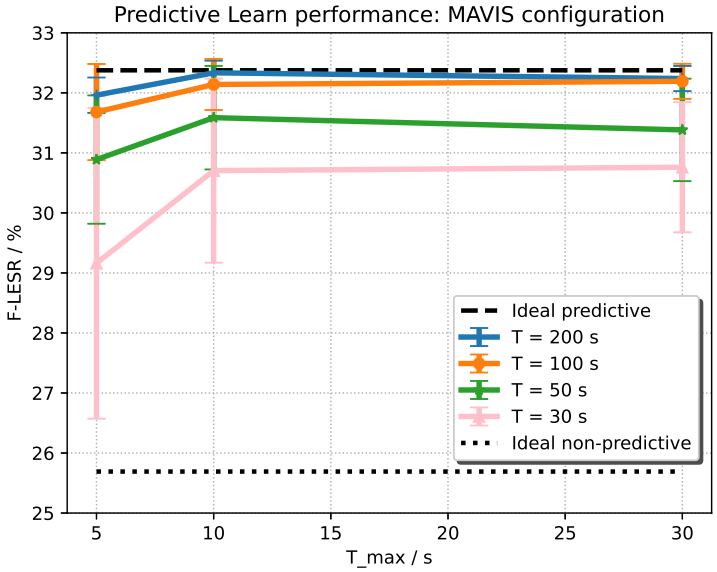

Figure 1. F-LESR from End-to-end simulation using reconstructor calculated from predictive *Learn* results as a function of T_{max}. Black dashed lines are the reference F-LESR acquired with ideal predictive and non-predictive L&A reconstructors.

For shorter slope buffers, the resulting F-LESRs were $31.6 \pm 0.9\%$ (T = 50 s) and $30.7 \pm 1.5\%$ (T = 30 s), which outperform the non-predictive L&A case by 39.8 nm and 37.0 nm RMS error respectively.

3.4 Discussion

Figure 1 implies that buffers with smaller T are expected to give lower F-LESR and higher STD, which suggests that the predictive *Learn* performance relies heavily on the convergence quality of slope buffer.

One way to study the convergence quality of a given numerical buffer is using the diagonal entries of the corresponding numerical covariance matrix. As introduced in Sec. 2, the covariance matrix is block-structured and the values are determined by the WFS sub-aperture coordinates. Specifically, the diagonal entries are the temporal auto-covariances of the measurements at each sub-aperture, and thus only depend on their axis of measurement (relative to the wind-profile) and not on their spatial coordinates. To see this, note that Eq. (4) does not depend on sub-aperture coordinates when $i = j$ (a.k.a., the diagonal).

Due to the stochastic nature of the turbulence, the convergence noise level varies significantly from buffers with different T, as well as buffers with the same T but collected at different times during a simulated observation. As a result, the diagonal values of numerical covariance matrices tend to follow a Gaussian distribution, and the STD of which can be used to describe the convergence quality of the slope buffer. Figure 2 is the histogram of diagonal values of 2 different numerical covariance matrices, along with Gaussian fits. Their corresponding F-LESR are 27.9% (blue histogram, labeled "bad") and 32.6% (orange histogram, labeled "good") respectively when $T_{max} = 10$ s. It is shown that their convergence quality in the $y - y$ direction are similar, yet the distributions in the $x - x$ direction are notably different, with the "bad" case deviating from the Gaussian distribution significantly.

All simulations presented in this paper are based on the use of a single NVIDIA Quadro RTX 5000 GPU. It is reasonable to assume that the supervisor hardware for dedicated systems such as MAVIS would be more powerful than this. The reference SLM paper[19] gives an informative scalability analysis for SLM algorithm which demonstrates a decent linear acceleration when the SLM process is parallelized on modern multiple NVIDIA GPU-based platforms.

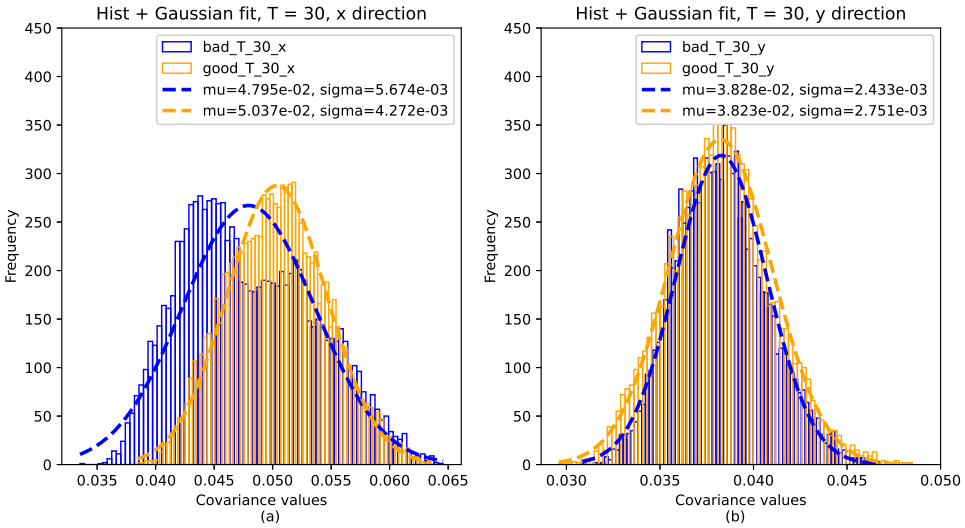

Figure 2. Histogram of numerical covariance matrices diagonal values, along with Gaussian fit, $T = 30$s. (a) $x - x$ slope covariance; (b) $y - y$ slope covariance

4. CONCLUSION

The pL&A reconstruction technique shows great potential in the scope of the MAVIS project. The real-time identification of atmospheric turbulence parameters is required for the computation of the predictive *Apply* reconstructor. Using the predictive *Learn* method, it is possible to acquire such information from a 100 s slope buffer requiring only 10 s for the SLM algorithm to converge on a reliable set of system parameters based on the MAVIS system configuration and modest GPU hardware. The expected long-exposure Strehl ratio over the 30 arcsec MAVIS science FoV is 32.1% at 550 nm using predictive *Learn* result obtained from a 100 s slope buffer, which is only 0.2% lower than the ideal predictive case when the true atmospheric profile is known. When the slope buffer is reduced to 30 s long, the performance drops only to 30.7%. These results indicate that the predictive *Learn* step of pL&A is feasible and effective on a system such as MAVIS, and further studies will be undertaken for a variety of different atmospheric profiles, as well as further convergence optimizations within the SLM algorithm.

REFERENCES

[1] F. Rigaut, R. McDermid, G. Cresci, *et al.*, "MAVIS Conceptual Design," in *(submitted to) Ground-based and Airborne Instrumentation for Astronomy VIII*, International Society for Optics and Photonics, SPIE (2020).

[2] F. Rigaut, D. Brodrick, G. Agapito, *et al.*, "Toward a Conceptual Design for MAVIS," in *6th International Conference on Adaptive Optics for Extremely Large Telescopes, AO4ELT 2019*, (2019).

[3] G. Agapito, D. Vassallo, C. Plantet, *et al.*, "MAVIS: System Modelling and Performance Prediction," in *Adaptive Optics Systems VII*, International Society for Optics and Photonics, SPIE (2020).

[4] F. Vidal, E. Gendron, and G. Rousset, "Tomography Approach for Multi-Object Adaptive Optics," *JOSA A* **27**(11), A253–A264 (2010).

[5] J. Cranney, H. Zhang, N. Doucet, *et al.*, "Predictive Learn and Apply: MAVIS application - Apply," in *Adaptive Optics Systems VII*, International Society for Optics and Photonics, SPIE (2020).

[6] L. Poyneer, M. van Dam, and J.-P. Véran, "Experimental verification of the frozen flow atmospheric turbulence assumption with use of astronomical adaptive optics telemetry," *JOSA A* **26**(4), 833–846 (2009).

[7] R. W. Wilson, "SLODAR: Measuring Optical Turbulence Altitude with a Shack–Hartmann Wavefront Sensor," *Monthly Notices of the Royal Astronomical Society* **337**(1), 103–108 (2002).

[8] T. Butterley, R. W. Wilson, and M. Sarazin, "Determination of the Profile of Atmospheric Optical Turbulence Strength From SLODAR Data," *Monthly Notices of the Royal Astronomical Society* **369**(2), 835–845 (2006).

[9] L. Gilles and B. L. Ellerbroek, "Real-Time Turbulence Profiling With a Pair of Laser Guide Star Shack–Hartmann Wavefront Sensors for Wide-Field Adaptive Optics Systems on Large to Extremely Large Telescopes," *JOSA A* **27**(11), A76–A83 (2010).

[10] L. Wang, M. Schöck, and G. Chanan, "Atmospheric turbulence profiling with slodar using multiple adaptive optics wavefront sensors," *Applied optics* **47**(11), 1880–1892 (2008).

[11] G. Sivo, A. Turchi, E. Masciadri, *et al.*, "Towards an Automatic Wind Speed and Direction Profiler for Wide Field Adaptive Optics Systems," *Monthly Notices of the Royal Astronomical Society* **476**(1), 999–1009 (2018).

[12] D. J. Laidlaw, *Turbulence and Wind Velocity Profiles From Adaptive Optics Telemetry: a General and Scalable Solution for Extremely Large Telescopes*. PhD thesis, Centre for Advanced Instrumentation, Durham University (2020).

[13] O. Martin, C. Correia, E. Gendron, *et al.*, "William Herschel Telescope Site Characterization Using the Moao Pathfinder Canary On-Sky Data," in *Adaptive Optics Systems V*, **9909**, 99093P, International Society for Optics and Photonics (2016).

[14] A. Tokovinin, "From Differential Image Motion to Seeing," *Publications of the Astronomical Society of the Pacific* **114**(800), 1156 (2002).

[15] B. Le Roux, J.-M. Conan, C. Kulcsár, *et al.*, "Optimal Control Law for Classical and Multiconjugate Adaptive Optics," *JOSA A* **21**(7), 1261–1276 (2004).

[16] É. Gendron, A. Charara, A. Abdelfattah, *et al.*, "A Novel Fast and Accurate Pseudo-Analytical Simulation Approach for MOAO," in *Adaptive Optics Systems IV*, **9148**, 91486L, International Society for Optics and Photonics (2014).

[17] K. Madsen, H. B. Nielsen, and O. Tingleff, "Methods for Non-Linear Least Squares Problems," (2004).

[18] N. Doucet, D. Gratadour, H. Ltaief, *et al.*, "Efficient Supervision Strategy for Tomographic AO Systems on E-ELT," in *Adaptive Optics for Extremely Large Telescopes V (AO4ELT5)], This conf.(Oct. 2017)*,

[19] Y. Hong, E. Bergou, H. Zhang, *et al.*, "Stochastic Levenberg-Marquardt for Solving Optimization Problems on Hardware Accelerators," in *(submitted to) 35th IEEE International Parallel & Distributed Processing Symposium*, IEEE (2021).

[20] "COMPASS: an Efficient, Scalable and Versatile Numerical Platform for the Development of ELT AO Systems, author=Gratadour, D and Puech, M and Vérinaud, C and Kestener, P and Gray, M and Petit, C and Brulé, J and Clénet, Y and Ferreira, F and Gendron, E and others, booktitle=Adaptive Optics Systems IV, volume=9148, pages=91486O, year=2014, organization=International Society for Optics and Photonics,"

[21] F. Ferreira, D. Gratadour, A. Sevin, *et al.*, "COMPASS: an Efficient GPU-based Simulation Software for Adaptive Optics Systems," in *2018 International Conference on High Performance Computing Simulation (HPCS)*, 180–187 (2018).