



HAL
open science

MAVIS: predictive Learn and Apply as a supervisory solution

Hao Zhang, Jesse Cranney, Damien Gratadour, Nicolas Doucet, Francois Rigaut

► **To cite this version:**

Hao Zhang, Jesse Cranney, Damien Gratadour, Nicolas Doucet, Francois Rigaut. MAVIS: predictive Learn and Apply as a supervisory solution. Adaptive Optics for Extremely Large Telescopes 7th Edition, ONERA, Jun 2023, Avignon, France. 10.13009/AO4ELT7-2023-097 . hal-04402869

HAL Id: hal-04402869

<https://hal.science/hal-04402869>

Submitted on 18 Jan 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



MAVIS: predictive Learn and Apply as a supervisory solution

Hao Zhang^a, Jesse Cranney^a, Damien Gratadour^{a,b}, Nicolas Doucet^a, and
Francois Rigaut^a

^aAustralian National University, College of Science, Research School of
Astronomy & Astrophysics, Canberra, Australia, 2611

^bLESIA, Observatoire de Paris, Université PSL, Sorbonne Université, Université
de Paris, Sorbonne Paris Cité, CNRS France

ABSTRACT

Learn and Apply is a 2-step reconstruction scheme for Adaptive Optics (AO) instruments, and has recently been extended to include a predictive step, i.e., predictive Learn and Apply (pL&A). It estimates the pseudo-real-time atmospheric turbulence profile directly from the AO telemetry buffer (the so-called Learn step), and then performs tomographic reconstruction based on the outputs (the so-called Apply step). We implement the entire pL&A pipeline in end-to-end simulations, where both the Learn and the Apply steps are performed in turns. We present the results of these simulations, as well as the challenges that were encountered and overcome – in particular, the challenges which have had an impact on the design of the MAVIS AO module.

Keywords: Adaptive Optics, turbulence profiling, real-time control, predictive control, Learn and Apply

1. INTRODUCTION

The MCAO Assisted Visible Imager and Spectrograph (MAVIS) is part of the next generation of Very Large Telescope instrumentation [31], and is set to deliver an image quality close to the diffraction limit in the visible wavelengths over a wide (30" squared) field of view, allowing deeper and sharper observations than ever before. These ambitious goals lead to a tight error budget, and are met by the current system design (the readers may refer to [26] for a comprehensive review). To fully measure the three-dimensional volume of atmospheric turbulence, MAVIS employs 8 laser guide stars (LGSs) located on a 17.5" radius circle for high-order (HO) aberration components, the corresponding Shack-Hartmann wavefront sensors (SH WFSs) have 40 sub-apertures along the telescope pupil. The tip-tilt (TT) and truth sensing is carried out by using 3 natural guide stars (NGSs) located within the 120" telescope technical FoV, and the corresponding WFSs have only 1×1 or 2×2 sub-apertures. As reported in [2], such NGS asterism achieves 50% sky coverage at the galactic pole with approximately 25 mas residual jitter. The real-time wavefront compensation is performed by three HO deformable mirrors (DMs) conjugated at [0, 6, 13.5] km. MAVIS will be installed on the highly performant Adaptive Optics Facility (AOF)

Further author information: (Send correspondence to F. R.)
F. R.: E-mail: francois.rigaut@anu.edu.au

of the VLT UT4, taking advantage of its extensive operational experience and the state-of-the-art instruments, e.g., the 4 Laser Guide Star Facility (4LGSF, see [16]) and the Deformable Secondary Mirror (DSM, see [11]).

The temporal error, caused by the constantly evolving atmospheric turbulence and the delay introduced by wavefront sensing, algorithms execution, and DM response, is one of the most critical components of the overall error budget, and has been extensively studied in the context of various instruments, see, e.g., [27], [4] and [9]. Regarding the MAVIS MCAO system, Agapito [2, 1] has presented that the expected temporal error yields approximately 38 nm (high-order) and 11 nm (low-order) under the assumption of bright NGSs with good asterism, whereas the overall error budget is 107 nm for high-order, and 30 nm for low-order (LO). Unlike other major error sources such as fitting error (error caused by the finite spatial sampling) and generalised fitting error (difference between the reconstructed turbulent layers and the projected phase on a finite number of DMs), the temporal error can be effectively modelled and mitigated with a well-designed control scheme, thus motivating the research of predictive control.

The Learn and Apply (L&A) approach was initially proposed in [30] and then demonstrated on-sky in [13] for the tomographic reconstruction in multi-object AO (MOAO) systems. It inherits the key concepts from the well-accepted MMSE approach (see, e.g., [6, 10, 20]), but accelerates the computation by operating directly from the WFS slope space. The two major steps in L&A are (i) the *Learn* step, where the recent (open-loop or pseudo-open-loop) WFS slopes are recorded to identify the atmospheric turbulence profile, and (ii) the *Apply* step, where the tomographic reconstructor is calculated from the results of the previous step with an MMSE process. Recently, Zhang [34] and Cranney [3] have introduced a predictive step to the L&A pipeline, i.e., the predictive L&A (pL&A). As demonstrated by the end-to-end numerical simulations, pL&A substantially mitigates the impact of temporal error, without any extra hardware or heavy computational load in hard real-time.

In this paper, we implement the entire pL&A pipeline in a closed supervisory loop, where both the Learn and the Apply steps are performed in turns. Sec. 2 derives the pL&A supervisory scheme by addressing a set of fundamental problems, including the issue of TT indetermination, the pseudo-open-loop (POL) implementation, the mitigation of tilt anisoplanatism (TA), and the MCAO supervisor pipeline. Sec. 3 presents the end-to-end numerical simulations to demonstrate the performance of pL&A as a closed-loop MCAO supervisor based on the MAVIS system configuration. In Sec. 4, we discuss the simulation results and propose future work to be done in the upcoming phases of the MAVIS project.

2. THE PREDICTIVE LEARN AND APPLY SCHEME

L&A’s fundamental assumption is that optimising the DM commands is straightforward enough when an on-axis guide star (GS) is available, where the control matrix is the pseudo-inverse of the interaction matrix \mathcal{D} . Though the on-axis GS is generally not available in MCAO, L&A introduces one or multiple virtual GSs (VGS, in the literature it is usually referred to as a truth sensor, e.g., in [21, 12, 5], but we prefer to call it VGS to avoid ambiguity, as ”truth sensors” is generally used in AO for Non-Common Path Aberration control applications). The VGS is usually considered as a bright HO NGS, therefore free from noise, cone effect, and LO uncertainty. In the case of MCAO, one VGS is not sufficient. A constellation of multiple VGSs is usually chosen to optimise the performance in the FoV of interest. Fig. 1 shows a 3×3 VGSs constellation for the MAVIS MCAO system. Note that the number, constellation, and individual weight of VGSs may vary according to the performance requirements and the computational resources.

Now, the goal of the tomographic reconstruction process is to link the VGSs to actual WFSs. Specifically, we consider a matrix \mathcal{R} that reconstructs the slopes for all VGSs, denoted s_v , from the slopes of actual WFSs, denoted s_m :

$$s_v = \mathcal{R}s_m, \quad (1)$$

where both slopes are assumed to be open-loop (i.e., independent of any DM influence). The linear least mean square error (LLMSE) [18] estimation of \mathcal{R} is written as (neglecting the noise covariance term):

$$\mathcal{R} = \langle s_v, s_m^T \rangle \langle s_m, s_m^T \rangle^{-1}, \quad (2)$$

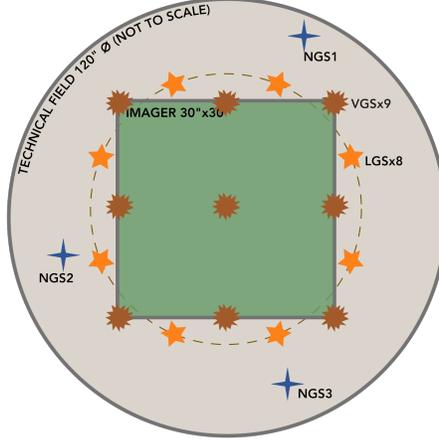


Figure 1. A typical VGS constellation for the MAVIS MCAO system.

where $\langle \cdot \rangle$ denotes the ensemble average, and the superscript \cdot^T denotes the transpose. By definition, $\langle s_v, s_m^T \rangle$ is the covariance matrix between VGSs and actual WFSs, $\langle s_m, s_m^T \rangle$ is the covariance matrix between actual WFSs and themselves,

Under observational conditions, $\langle s_m, s_m^T \rangle$ can be populated from either numerical method (i.e., with the recent telemetry data) or pseudo-analytical method (i.e., from a well-defined atmospheric turbulence model), while $\langle s_v, s_m^T \rangle$ is available only in the pseudo-analytical method, as the VGSs are not present in the actual instrumental setting. As an illustration, the predictive Learn (pLearn) step takes as input the numerical $\langle s_m, s_m^T \rangle$, and estimates a set of atmospheric turbulence parameters, including the fractional C_n^2 , the wind speed and direction for each turbulent layer, via an iterative fitting process. The *Learnt* parameters are then used in the predictive Apply (pApply) step to calculate the pL&A reconstructor matrix for the closed-loop AO control. The readers are pointed to [34, 3] for detailed derivations of the predictive tomography model based on the von Kármán spectrum [32] and Taylor's *frozen flow* hypothesis [29], whereas in this paper we focus on the fundamental limitations encountered in the closed-loop implementation.

2.1 Tip-tilt indetermination

During the uplink propagation of a LGS, the laser beam is distorted by atmospheric turbulence, which means its angular position on-sky is not determined, and the TT measurements from an LGS becomes useless, i.e., the TT indetermination ([25]). In practice, we remove the average TT from each WFS by removing the averaged x component from all x slopes, and the same for y slopes. Note that this is a linear function of the slope vector, and can be embedded as a matrix-vector multiplication:

$$s_m^{TTF} = \mathcal{Q}s_m, \quad (3)$$

where \mathcal{Q} is the so-called TT filtering (TTF) matrix.

Similarly, for the covariance matrix of TT-filtered slopes:

$$\begin{aligned} \langle s_m, s_m^T \rangle^{TTF} &= \left\langle (s_m^{TTF}) (s_m^{TTF})^T \right\rangle \\ &= \left\langle \mathcal{Q} (s_m) (s_m)^T \mathcal{Q}^T \right\rangle \\ &= \mathcal{Q} \langle s_m, s_m^T \rangle \mathcal{Q}^T, \end{aligned} \quad (4)$$

where $\langle s_m, s_m^T \rangle$ is the TT-included covariance matrix. Eq. (4) is valid for either the whole covariance matrix or a block of it, but not for individual covariance values.

The TTF process can be highly expensive in terms of computations, especially in the pLearn algorithm. As discussed in [17], the current pLearn algorithm is accelerated by the stochastic Levenberg-Marquardt (SLM)

algorithm where every iteration uses a subset of the numerical covariance matrix, substantially reducing the total number of calculations. However, when considering the TT indetermination issue, the TTF process must be applied at each iteration of the SLM algorithm, requiring to generate the whole covariance matrix instead of only a subset. If not handled properly, SLM loses its advantage in computation speed.

An efficient way to circumvent the TTF issue is reconstructing the *effective* TT information from the NGS measurements. This can be achieved by constructing a cross-covariance matrix between the HO and LO measurements. NGS WFSs generally have fewer sub-apertures due to the limited flux: in the case of MAVIS, the NGS WFSs have only 1×1 or 2×2 sub-apertures. In this case, all NGSs sub-apertures are not perfect squares nor fully illuminated, limiting the use of existing models (see, e.g., [21]).

We hereby propose a generalised wavefront sensing model that takes the illumination condition of different sub-apertures into consideration. A detailed derivation and analysis of the generalised model will be submitted as a separate paper, here we give a brief summary of the key concepts. First off, a fine grid is defined in each sub-aperture area, as shown in Fig. 2 (a). Due to the telescope spider, central obstruction and pupil shape, a sub-aperture is sometimes partially illuminated, especially for large NGS sub-apertures. In this case, introducing the fine grid allows us to consider only the valid area of a sub-aperture in the covariance model. We will denote d the cell width of the grid, L the width of a square sub-aperture, and we have $d \in (0, L]$. For simplicity, we assume L is a multiple of d , and the number of cells along the telescope pupil is given by $n_{\text{sub-ap}} = L/d$, where $n_{\text{sub-ap}}$ is the number of sub-apertures along the pupil. The x -slope measured by a sub-aperture is defined by the average wavefront phase difference over its valid area

$$s_x = \frac{\lambda}{2\pi N_c} \sum_{i=1}^{N_c} \frac{\varphi_{li} - \varphi_{ri}}{d}, \quad (5)$$

where N_c is the number of valid cells in this sub-aperture, φ_{li} and φ_{ri} are the phase on the left and right sides of i th cell, respectively.

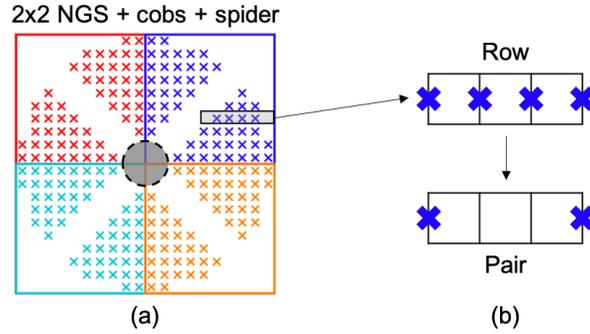


Figure 2. (a) Geometry of a 2×2 WFS with telescope central obstruction and spider, grids are defined on the valid area only; (b) illustration of the left or right side of a cell/cell-pair.

As shown in Fig. 2 (b), the right side of one cell is equivalent to the left side of the next L cell in the same row (this can be extended to the y -slope case by replacing left/right to bottom/top and row to column). That is to say, populating a pseudo-analytical slope-space covariance matrix from all cells is equivalent to using only the effective *cell-pairs*, so the overall computational load of the generalised model is still affordable in soft real-time.

Now we consider a set of VGSs that has the same (i) angular constellation, (ii) altitude, and (iii) sub-aperture geometry as the actual LGSs, but is capable of measuring both TT and higher-order slopes. In this case, TTF shall be applied to the actual LGS slopes as well as the covariance matrices, so Eq. (1) becomes:

$$\hat{s}_v = \underbrace{\langle s_v, s_m^T \rangle \mathcal{Q}^T (\mathcal{Q} \langle s_m, s_m^T \rangle \mathcal{Q}^T)^{-1} \mathcal{Q}}_{W_{TT}} s_m^{POL}, \quad (6)$$

where the noise covariance term is neglected, s_m^{POL} is the POL slopes (the POL implementation will be explained in more details in the following sub-section) collected from actual WFSs. As shown by Eq. (6), the TT measurements from LGSs, if any, do not propagate through the reconstruction process, therefore the reconstructed TT in \hat{s}_v only comes from the NGS slopes. Note that $\langle s_v, s_m^T \rangle$ is almost identical to the non-filtered $\langle s_m, s_m^T \rangle$, except that the NGS rows are removed.

Given the system geometry and the atmospheric turbulence profile, the overall projection matrix W_{TT} can be pre-calculated and combined to the POL slopes. The reconstructed \hat{s}_v is then used to populate the TT-included numerical covariance matrix, so that SLM can be directly applied without any TTF during the pLearn algorithm.

2.2 Pseudo-open-loop control

Previous discussions are mainly based on the statistical properties of the open-loop slope. However, the statistical properties of the open-loop slope are different from those of the closed-loop slope. This discrepancy may lead to degraded performance and, in extreme cases, for the AO system to become unstable. The POL method described in [8, 3] provides a computationally efficient yet robust cure.

If we assume that perfect *a priori* knowledge of the DM actuators to WFSs slopes interaction matrix is available, so that the open-loop equations above can be extended to the closed-loop case by recovering s_m^{OL} from the closed-loop slopes:

$$s_{m,k}^{POL} = s_{m,k}^{CL} + \mathcal{D} [(1 - \Delta)u_{k-\overline{\delta k}+1} + \Delta u_{k-\overline{\delta k}}], \quad (7)$$

where $s_{m,k}^{CL}$ is the closed-loop slope of actual WFSs at the k th iteration, u is the DM command vector, δk accounts for the delay and is assumed to be exactly 2 for the MAVIS configuration, Δ is defined as $(\delta k \bmod 1)$, $\overline{\delta k}$ is defined as $(\text{ceil } \delta k)$. $s_{m,k}^{POL}$ is only an approximation of s_m^{OL} at the k th iteration based on the interaction matrix \mathcal{D} and applied commands u , and as such the following terms are not (and cannot be) fully considered:

1. WFSs uncertainties, including the device noise, mis-calibrations, mis-alignments, as well as the non-linearity from elongation and centroiding algorithms;
2. DM uncertainties caused by the electrical and mechanical noises, and the un-modelled dynamics of the DMs;
3. Discrepancy in the interaction matrix caused by the mis-registration between DMs and WFSs.

Despite these error sources, [23] and [14] have demonstrated remarkable stability, efficiency, and performance of the POL method in MCAO systems via theoretical and numerical approaches.

2.3 Mitigation of the tilt anisoplanatism

The tilt anisoplanatism is closely related to the TT indetermination: due to the TTF process, the TT measurements are completely filtered out from every LGS, these include the global and field dependent TT. The latter can be modelled and visualised by three quadratic modes (one defocus and two astigmatisms) on an altitude layer, the readers may refer to [33] for a detailed discussion.

It is shown in [7] that these LGS-insensitive modes are well measured by at least three 1×1 NGSs with a reasonable asterism (e.g., not located in a line), as is the case in MAVIS. At least two DMs, one conjugated at the ground layer $h_0 = 0$, and the other conjugated at altitude h_d , are then used for compensations, as detailed below.

In the general case where the turbulence phase φ_a on layer a is given by a combination of quadratic modes, we have:

$$\varphi_a(x_a, y_a) = \alpha_{a1}x_a^2 + \alpha_{a2}x_a y_a + \alpha_{a3}y_a^2, \quad (8)$$

where $\alpha_{aj}, j \in \{1, 2, 3\}$ are scalar coefficients.

The overall projected phase on the pupil plane from all layers, as seen by a NGS or a science target at infinite altitude (i.e., no cone effect), is given by:

$$\varphi_0(x_0, y_0) = \sum_{a=1}^{n_L} \varphi_a(x_0 + h_a \theta_x, y_0 + h_a \theta_y), \quad (9)$$

where (θ_x, θ_y) is the angular position of the NGS or the target, n_L is the number of turbulent layers.

Our goal is to find a set of DM commands, such that for any NGS or target in the FoV: (i) the DMs do not introduce any second-order phase, and (ii) the overall plate scale modes are cancelled out. Similarly, we define the ground layer DM phase φ_0 , and the altitude layer DM phase φ_d as:

$$\begin{aligned}\varphi_0(x_0, y_0) &= \alpha_{01}x_0^2 + \alpha_{02}x_0y_0 + \alpha_{03}y_0^2 \\ \varphi_d(x_d, y_d) &= \alpha_{d1}x_d^2 + \alpha_{d2}x_dy_d + \alpha_{d3}y_d^2.\end{aligned}\quad (10)$$

The solution is found when the projected result of Eq. (10) cancels the turbulence phase in Eq. (9), which yields:

$$\begin{aligned}\alpha_{0j} &= -\alpha_{dj} \\ \alpha_{dj} &= -\frac{1}{h_d} \sum_{a=1}^{n_L} h_a \alpha_{aj},\end{aligned}\quad (11)$$

where $j \in \{1, 2, 3\}$.

Due to the limited number of NGS sub-apertures, it is generally not possible to resolve the altitude distribution of TA modes. But this does not impact the conclusion in Eq. (11). As discussed in [24], the projected phase of any Zernike mode on a sub-region of a given altitude can be written as a combination of Zernike modes up to the same order, therefore the TA modes at different altitudes can be projected to a single layer, e.g., at h_d , with a combination of TT and TA coefficients (piston removed). Consequently, the second line in Eq. (11) becomes $\alpha_{dj} = -\alpha_{aj}$, as $h_a = h_d$.

Various methods have been investigated to compensate for the TA modes in closed-loop, such as split tomography [15], the slope-merging method [33], and the generalised model presented in this paper. As discussed in Sec. 2.1, the generalised model allows estimating the cross-covariance between HO and LO measurements. Consider a set of VGSs as shown in Fig. 1, the pL&A reconstructor takes the same format as in Eq. (6), except that the LLMSE estimator here optimises the residual slope for the whole FoV.

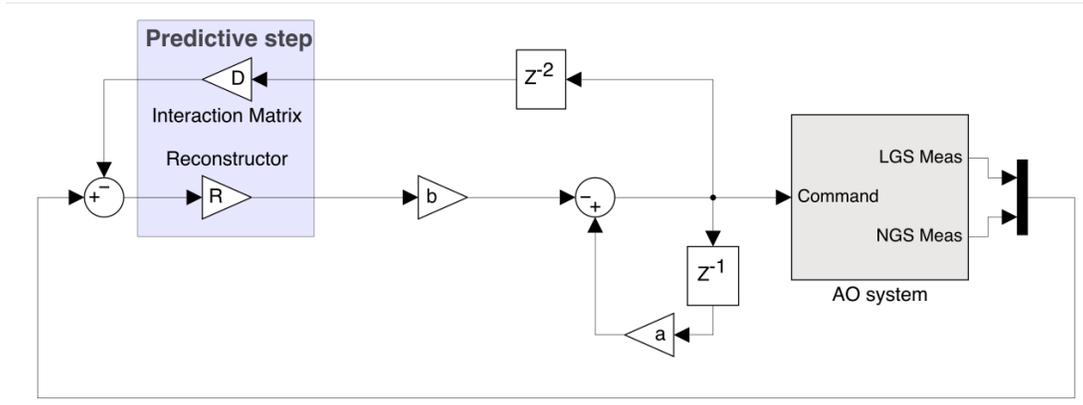


Figure 3. Integrated controller block diagram, where the LGS and NGS measurements are combined together before the predictive reconstruction step.

An integrated controller (IC) is used to control both the LO and HO sub-systems simultaneously. The corresponding block diagram is shown in Fig. 3, where the LO and HO measurements are combined together using the cross-covariance matrices during the tomographic reconstruction. Compared to conventional split tomography, IC has the advantage of simplicity, thus reducing the complexity of the soft real-time cluster (SRTC) design. Additionally, the predictive tomographic reconstruction can now be applied to both sub-systems, so an improved MCAO performance is expected especially for the LO aberrations.

2.4 MCAO supervisor pipeline

Finally, we implement the pL&A scheme as a MCAO supervisory solution. As one of the most critical components of the SRTC, the MCAO supervisor is responsible for providing the hard real-time computer (HRTC) with control matrices, and updating the matrices based on the statistics of atmospheric turbulence at a regular rate (depending on the control strategy and the turbulence condition). It splits the MCAO observation time into multiple rounds (see Fig. 4): an initial data collection round, i.e., “zeroth round”, and the following pL&A rounds. Every round consists of two parts: closed-loop MCAO operation and telemetry data collection, and the reconstructor calculation. Note that under observational conditions where the turbulence statistics are constantly evolving, the time taken by the algorithm execution has to be considered: in each round, only a part of the telemetry buffer (green blocks) can be used in the current pLearn algorithm, whereas the rest is sent to the next round.

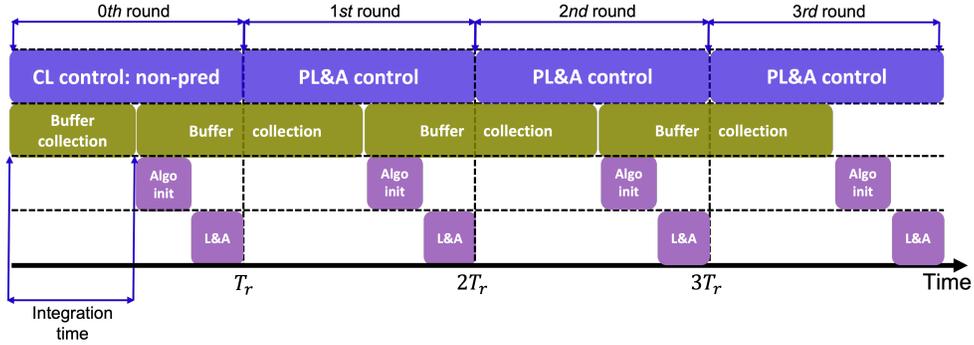


Figure 4. Timing diagram for the pL&A MCAO supervisor, where T_r is the overall round duration. The size of the blocks is scaled for better visualisation, not necessarily aligned with the actual operation time.

The main goal of the zeroth round is to collect the WFS measurements and DM commands buffer for the pL&A tomography. Real-time information of the atmospheric turbulence profile is not available at the beginning of the zeroth round, therefore predictive tomography cannot be performed. However, the MCAO loop shall be closed to avoid significant truncation in the LGS WFS spots, and maintain a reasonable signal-to-noise ratio. In this paper, the zeroth round is controlled by an integrated controller based on non-predictive L&A tomography.

As discussed in Sec. 2.1 and Sec. 2.2, the reconstruction of POL slopes uses a TT-included covariance matrix between the WFS measurements and themselves. This covariance matrix has to be populated analytically for a better estimation of the low-order information. In this case, an initial estimation of the atmospheric turbulence profile is required, including the overall r_0 , the altitude for each layer, and the fractional C_n^2 profile. The wind profile, on the other hand, is not necessary during the reconstruction, as the covariance matrix used here does not include any time delay. Experience data or the off-line atmospheric turbulence profiling results (e.g., MASS, SCIDAR, or SLODAR, see in [19, 28, 22]) can be used.

After calculating the reconstructor for each pL&A round, the pseudo-analytical, TT-included covariance matrix is saved for the next POL reconstruction process. This introduces the risk of instability if there are any unseen modes growing through different rounds. When instability happens, i.e., when the round-wise LESR keeps reducing in multiple continuous rounds, the supervisor shall be reset with new initial profiles and geometric POL projection matrices to avoid further loss of performance. Note that we have not witnessed any of these instabilities in our simulations, and the reset feature is only provided as a precautionary measure.

At this stage, we do not take the real-time evolution of the turbulence statistics into consideration. However, this study is quite important in optimising the observation time for each round, and will be carried out in the following phases of the MAVIS project.

3. NUMERICAL SIMULATION RESULTS

3.1 Simulation configurations

Detailed system configurations and parameters are presented as follows. The MCAO system configuration can be found in Table 1. The true atmospheric turbulence profile used in the simulations is adapted from the Paranal Stereo-SCIDAR median profiles release 2019B, see [22]. Note that the wind direction settings are not directly related to the AO performance, and are only selected to increase the complexity of the atmospheric turbulence profiling process to fully investigate its performance. Atmospheric turbulence parameters used for a *close init* (only used for the non-predictive control in the zeroth round) are presented in Table 2. The parameters used for the pL&A algorithm are based on the configurations presented in [34] and [3].

Table 1. MAVIS MCAO module configuration and simulation parameters

Sub-Module	Parameters	
General	Tel. Diam	8 m
	Central Obstruction	16% \emptyset
	Nominal target Wavelength	550 nm
	Frame-rate	1000 Hz
	Science FoV	30" \emptyset
Atmos	Num. of Layers	10
	r_0 at 500 nm	12.6 cm
	Altitude (min, max)	(30, 14000) m
	Wind Speed (min, max)	(4.5, 34.3) m/s
	Wind Dir.* (min, max)	(0°, 25°)
LGS	Num. of LGS	8
	Constellation	35" \emptyset
	Num. of Sub-aperture	40 \times 40
	Read-out Noise	0.2 e^-
	Flux/sub-ap/ms	75 photons
	Excess noise ratio	2.0
NGS	Num. of NGS	3
	Constellation	40" \emptyset
	Num. of Sub-aperture	1 \times 1
	Read-out Noise	0.5 e^-
HO DMs	H-Magnitude	down to 18.5
	Altitude	[0, 6, 13.5] km
	Actuator Pitch	[22, 25, 32] cm
Performance Requirements	Sky coverage**	$\geq 50\%$
	LESR***	$\geq 10\%$ (goal 15%) in V band

* This wind direction is chosen to test the controller performance.

** The sky coverage is defined at the galactic poles.

*** Here LESR is averaged over a circular field of diameter 30".

3.2 Simulation results

We performed benchmarking of the closed-loop MCAO supervisor on a DGX server, equipped with 8 NVIDIA V100 GPUs. At this stage, most of the tasks are carried out by a single GPU, including the SLM algorithm and the calculation of analytical covariance matrices. The rest of the tasks, such as the POL calculation and the sub-aperture filtering, are performed on CPUs. Comprehensive GPU-based computation and multi-GPU parallelisation are not implemented in the latest version of the SRTC software, and will be implemented as future work.

Table 2. Atmospheric turbulence profile used in the zeroth round

Name	Value
Close init	
Number of layers	4
r_0 at 500 nm	12.6 cm
Altitude	[30, 4500, 11000, 14000] m
C_n^2	[0.7, 0.1, 0.1, 0.1]
Wind speed	Not used
Wind direction	Not used
Outer scale	[25, 25, 25, 25] m

The overall observation time is set to 30 mins (on-sky time) in our simulations to demonstrate the long exposure stability of the MCAO supervisor. The averaged long-exposure Strehl ratio (LESR) from 15 targets distributed in the 30" \emptyset science FoV is used to investigate the overall performance, along with the corresponding (spatial) standard deviation. The LESR of every target is reset before the first pL&A round, so that the limited performance of the non-predictive zeroth round does not affect the final results.

3.3 Overall MCAO supervisor performance

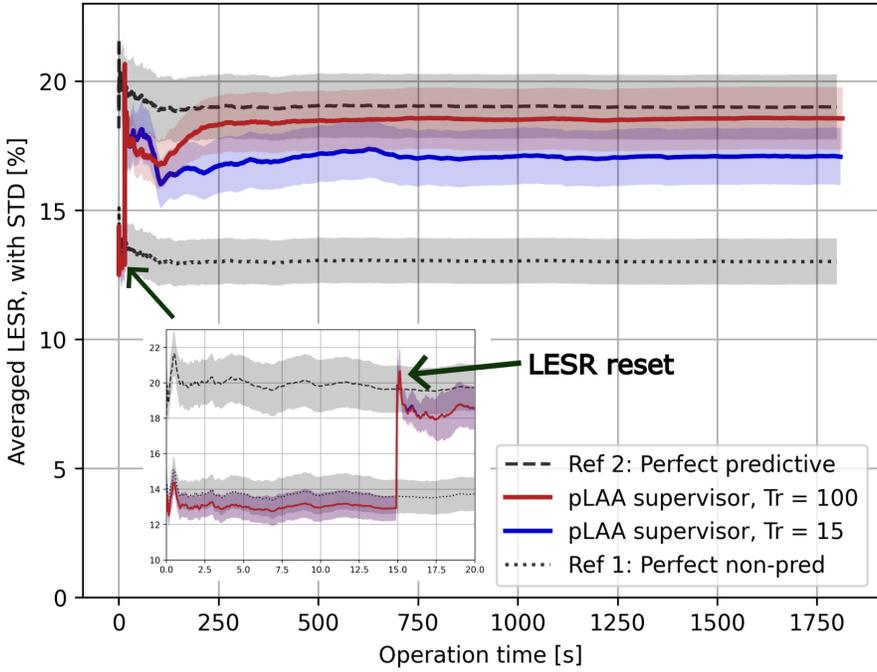


Figure 5. Overall performance of the MCAO supervisor. Solid curves: cumulative LESR acquired with 100s or 15s round duration T_r ; Shaded area: standard deviation of the LESR from all targets. Zoom-in area: performance of the zeroth round.

Fig. 5 presents the long exposure performance of the MCAO supervisor. Field averaged LESR acquired from perfect predictive (i.e., given the exact atmospheric turbulence and wind profile) and non-predictive (i.e., given perfect knowledge of r_0 , C_n^2 and altitude distribution, without wind profile) L&A control scheme are provided as reference. After 30 mins exposure time, perfect predictive controller gives $19.00 \pm 1.25\%$ LESR over the science FoV, whereas perfect non-predictive controller gives $13.01 \pm 0.89\%$.

In the simulations, the zeroth round takes 15 secs, as shown in the zoom-in area. The MCAO performance of the zeroth round is worse than the non-predictive curve due to the close init profile used for the reconstructor calculation. The LESR is reset before the first pL&A round, and a significant improvement in MCAO performance can be seen in both cases. Note that a spike can be seen after resetting LESR, which is due to the settling time of the long exposure performance.

Round durations of 100 secs and 15 secs are considered, the LESR yield $18.56 \pm 1.21\%$ and $17.07 \pm 1.10\%$, respectively. Compared to the non-predictive curve, the improvements in terms of RMS wavefront error are 52.18 nm ($T_r = 100$) and 45.62 nm ($T_r = 15$). In fact, the LESR acquired with $T_r = 100$ secs is quite close to the perfect predictive curve, the difference corresponds to only 13.40 nm RMS error.

The results also demonstrate the stability of the MCAO supervisor: in both cases, the LESR converges after a few rounds, with no significant fluctuations observed throughout the 30 mins exposure time. The stability is improved with longer T_r , because of the lower convergence noise in numerical covariance matrices.

3.4 Robustness study: cumulative v.s. round-wise LESR

In Fig. 6, we show the round-wise field LESR side by side with the cumulative curves. The round-wise scatters are acquired by resetting the LESR before every pL&A rounds, while all other settings, including the random seed for phase screen generation, are identical to the red and blue curves. Due to the time required for LESR settling, the performance at the beginning of each round might be different from the converged LESR, i.e., the *spikes* in the round-wise scatters. These results are not directly related to the science cases, and typically not available in real applications. Nevertheless, they provide useful information to study the robustness of the MCAO supervisor given the stochastic nature of atmospheric turbulence.

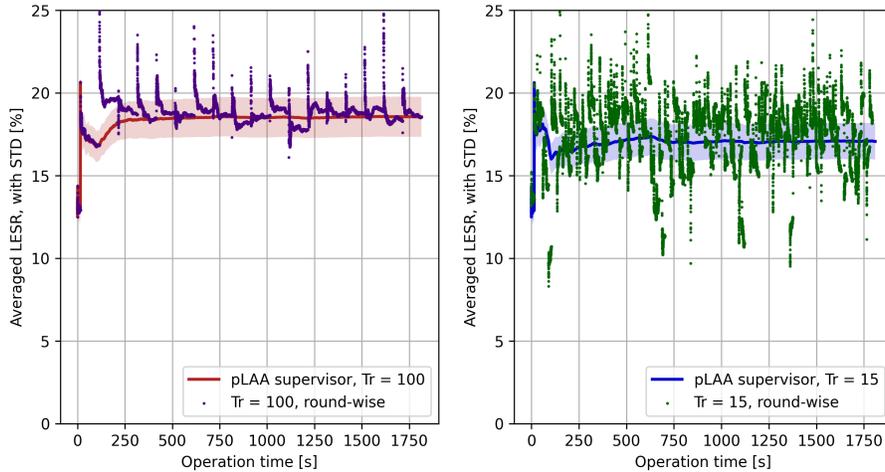


Figure 6. Cumulative v.s. round-wise LESR. Left: $T_r = 100$; Right: $T_r = 15$. Scatter plots: round-wise field averaged LESR. Round-wise scatters: acquired by resetting the LESR for all targets before every pL&A rounds. Note that the *spikes* are also caused by the LESR resetting.

Highly stable field averaged LESR is observed when $T_r = 100$. Throughout the exposure time, the worst performance is seen at the first round, where the field averaged LESR is 16.85%, 1.7% lower than the overall performance in the red curve. This is mainly constrained by the duration of the zeroth round, i.e., the convergence noise caused by limited samples. The performance of all other pL&A rounds are fairly close to the red curve, the maximum difference in LESR is within $\pm 1\%$.

The round-wise performance when $T_r = 15$, on the other hand, shows higher fluctuation. Among the 119 pL&A rounds in the green scatter, the LESR of 24 rounds are higher than 18.5%. The LESR of 5 rounds are lower than 13%, with the worst being the 6th and the 74th, whose LESR are 10.71% and 10.91%, respectively. At these rounds, the predictive controller are doing even worse than the perfect non-predictive controller (reference curve 1 in Fig. 5). Although this comparison is not really fair – the non-predictive controller requires exact knowledge of

the C_n^2 distribution which is not available in real applications, the results still indicate strong convergence noise in the numerical buffers.

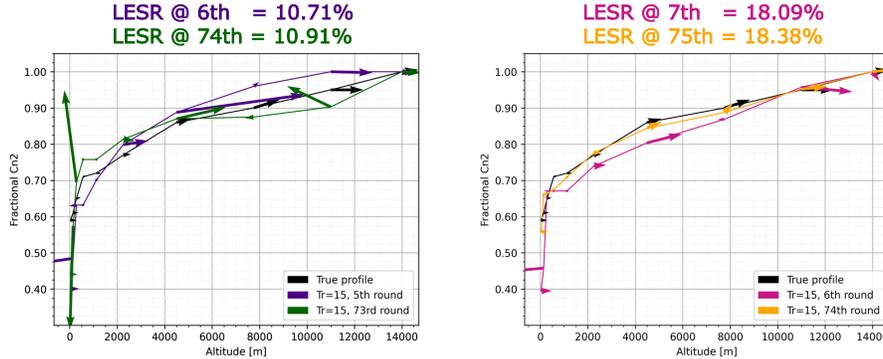


Figure 7. pLearn outputs from different rounds. Left: pLearn outputs **used** in the worst rounds; Right: pLearn outputs **generated** in the worst rounds, and **used** in the following rounds. Black curve: True profile.

The learnt profile from different rounds provide an efficient way to investigate the level of convergence noise, and the propagation of convergence noise through rounds. This is particularly relevant to the robustness of the MCAO supervisor. Note that in our implementation, a pL&A round is controlled by the reconstructor calculated from the previous pLearn output, and the numerical buffer collected in this round is used to control the next one. In this case, we present in Fig. 7 the learnt profiles generated from the 5th and the 73rd rounds, side by side with the profiles from the 6th and the 74th rounds.

Purple and green curves in the left subplot show strong convergence noise with the existence of undesired peaks/valleys: pLearn tends to introduce extremely fast layers to fit those values, which severely reduces the precision of the learnt profile. However, such error does not propagate to the next round. As shown by the pink and orange curves in the right subplot, pLearn in the 6th and the 74th rounds perform well in spite of their poor round-wise LESR. Consequently, the round-wise LESR climbs back immediately after the worst rounds, and the supervisory loop remains stable throughout the 30 mins exposure time.

4. DISCUSSIONS AND FUTURE WORK

In this paper we implemented the predictive L&A scheme in a closed MCAO supervisory loop. The supervisory loop can be divided into multiple rounds: a zeroth round controlled by a non-predictive scheme, and the following L&A rounds with a predictive, integrated controller. We analysed the major challenges: the TT indetermination, the pseudo-open-loop implementation, the mitigation of tilt anisoplanatism, and the MCAO supervising strategy. End-to-end numerical simulations are performed to demonstrate the overall performance and robustness of the MCAO supervisor. The RMS wavefront error is reduced by 45.62 nm when using 15 secs per round, or 52.18 nm when using 100 secs per round (both compared to perfect non-predictive controller). Due to the stochastic nature of the turbulence, some of the rounds suffer from high convergence noise, which reduces the MCAO performance, but this error does not propagate to the next round, so the supervisory loop remains stable.

The predictive L&A scheme is an effective supervisory solution as it does not require extra hardware or extra computations in hard real-time, and the increase in soft real-time computational load is affordable. Future work will focus on optimising the performance of the MCAO supervisor. As an example, the off-line atmospheric turbulence profiling data (e.g., the latest stereo-SCIDAR data release) acquired at the same observational site can be analysed to derive the evolution rate of the turbulence statistics, which is helpful to balance the convergence noise caused by limited time of data collection, and the capability of capturing the instantaneous turbulence profile. Experimental validations will also be performed by the MAVIS consortium. Before the actual implementation on the Very Large Telescope, key components of the AO module will be prototyped and tested in-lab, along with the control schemes.

References

- [1] Guido Agapito et al. “MAVIS: performance estimation of the adaptive optics module”. In: *Adaptive Optics Systems VIII*. Vol. 12185. SPIE. 2022, pp. 1094–1106.
- [2] Guido Agapito et al. “MAVIS: system modelling and performance prediction”. In: *Adaptive Optics Systems VII*. Vol. 11448. SPIE. 2020, pp. 763–774.
- [3] Jesse Cranney et al. “Predictive learn and apply: MAVIS application-apply”. In: *Adaptive Optics Systems VII*. Vol. 11448. SPIE. 2020, pp. 552–559.
- [4] Marcos A van Dam, David Le Mignant, and Bruce A Macintosh. “Performance of the Keck Observatory adaptive-optics system”. In: *Applied Optics* 43.29 (2004), pp. 5458–5467.
- [5] Nicolas Doucet et al. “Scalable soft real-time supervisor for tomographic AO”. In: *Adaptive Optics Systems VI*. Vol. 10703. International Society for Optics and Photonics. 2018, p. 107034L.
- [6] Brent L Ellerbroek. “Efficient computation of minimum-variance wave-front reconstructors with sparse matrix techniques”. In: *JOSA A* 19.9 (2002), pp. 1803–1816.
- [7] Brent L Ellerbroek and François Rigaut. “Methods for correcting tilt anisoplanatism in laser-guide-star-based multiconjugate adaptive optics”. In: *JOSA A* 18.10 (2001), pp. 2539–2547.
- [8] Brent L Ellerbroek and Curtis R Vogel. “Simulations of closed-loop wavefront reconstruction for multiconjugate adaptive optics on giant telescopes”. In: *Astronomical Adaptive Optics Systems and Applications*. Vol. 5169. SPIE. 2003, pp. 206–217.
- [9] Florian Ferreira et al. “Numerical estimation of wavefront error breakdown in adaptive optics”. In: *Astronomy & Astrophysics* 616 (2018), A102.
- [10] Thierry Fusco et al. “Optimal wave-front reconstruction strategies for multiconjugate adaptive optics”. In: *JOSA A* 18.10 (2001), pp. 2527–2538.
- [11] Daniele Gallieni and Roberto Biasi. “The new VLT-DSM M2 unit: construction and electromechanical testing”. In: *Proceedings of the Third AO4ELT Conference*. 2013, p. 122.
- [12] E Gendron et al. “A novel fast and accurate pseudo-analytical simulation approach for MOAO”. In: *Adaptive Optics Systems IV*. Vol. 9148. International Society for Optics and Photonics. 2014, p. 91486L.
- [13] E Gendron et al. “MOAO first on-sky demonstration with CANARY”. In: *Astronomy & Astrophysics* 529 (2011), p. L2.
- [14] Luc Gilles. “Closed-loop stability and performance analysis of least-squares and minimum-variance control algorithms for multiconjugate adaptive optics”. In: *Applied Optics* 44.6 (2005), pp. 993–1002.
- [15] Luc Gilles and Brent L Ellerbroek. “Split atmospheric tomography using laser and natural guide stars”. In: *JOSA A* 25.10 (2008), pp. 2427–2435.
- [16] W Hackenberg et al. “The four-laser guide star facility (4LGSF) for the ESO VLT adaptive optics facility (AOF)”. In: *Second International Conference on Adaptive Optics for Extremely Large Telescopes*. Online at; A href=” <http://ao4elt2.lesia.obspm.fr>” & <http://ao4elt2.lesia.obspm.fr/j/A>. 2011, p. 56.
- [17] Yuxi Hong et al. “Outsmarting the Atmospheric Turbulence for Ground-Based Telescopes Using the Stochastic Levenberg-Marquardt Method”. In: *European Conference on Parallel Processing*. Springer. 2021, pp. 565–579.
- [18] Thomas Kailath, Ali H Sayed, and Babak Hassibi. *Linear estimation*. BOOK. Prentice Hall, 2000.
- [19] Victor Kornilov et al. “MASS: a monitor of the vertical turbulence distribution”. In: *Adaptive Optical System Technologies II*. Vol. 4839. SPIE. 2003, pp. 837–845.
- [20] Brice Le Roux et al. “Optimal control law for classical and multiconjugate adaptive optics”. In: *JOSA A* 21.7 (2004), pp. 1261–1276.
- [21] Olivier Martin et al. “Temporal convergence of phase spatial covariance matrix measurements in tomographic adaptive optics”. In: *Adaptive Optics Systems III*. Vol. 8447. International Society for Optics and Photonics. 2012, 84472A.

- [22] J Osborn et al. “Optical turbulence profiling with Stereo-SCIDAR for VLT and ELT”. In: *Monthly Notices of the Royal Astronomical Society* 478.1 (2018), pp. 825–834.
- [23] Piotr Piatrou and Luc Gilles. “Robustness study of the pseudo open-loop controller for multiconjugate adaptive optics”. In: *Applied optics* 44.6 (2005), pp. 1003–1010.
- [24] Roberto Ragazzoni, Enrico Marchetti, and Francois Rigaut. “Modal tomography for adaptive optics”. In: *Astronomy and Astrophysics* 342 (1999), pp. L53–L56.
- [25] Francois Rigaut and Eric Gendron. “Laser guide star in adaptive optics-The tilt determination problem”. In: *Astronomy and Astrophysics* 261 (1992), pp. 677–684.
- [26] François Rigaut et al. “MAVIS conceptual design”. In: *Ground-based and Airborne Instrumentation for Astronomy VIII*. Vol. 11447. SPIE. 2020, pp. 378–393.
- [27] Francois J Rigaut, Jean-Pierre Véran, and Olivier Lai. “Analytical model for Shack-Hartmann-based adaptive optics systems”. In: *Adaptive Optical System Technologies*. Vol. 3353. International Society for Optics and Photonics. 1998, pp. 1038–1048.
- [28] HW Shepherd et al. “Stereo-SCIDAR: optical turbulence profiling with high sensitivity using a modified SCIDAR instrument”. In: *Monthly Notices of the Royal Astronomical Society* 437.4 (2014), pp. 3568–3577.
- [29] Geoffrey Ingram Taylor. “The spectrum of turbulence”. In: *Proceedings of the Royal Society of London. Series A-Mathematical and Physical Sciences* 164.919 (1938), pp. 476–490.
- [30] Fabrice Vidal, Eric Gendron, and Gérard Rousset. “Tomography approach for multi-object adaptive optics”. In: *JOSA A* 27.11 (2010), A253–A264.
- [31] Valentina Viotto et al. “MAVIS: preliminary design of the adaptive optics module”. In: *Adaptive Optics Systems VIII*. Vol. 12185. SPIE. 2022, pp. 623–634.
- [32] DM Winker. “Effect of a finite outer scale on the Zernike decomposition of atmospheric optical turbulence”. In: *JOSA A* 8.10 (1991), pp. 1568–1573.
- [33] Hao Zhang et al. “Mitigating Tilt Anisoplanatism with the Slope Merging Method for Multi-Conjugate Adaptive Optics Systems”. In: *2021 Australian & New Zealand Control Conference (ANZCC)*. IEEE. 2021, pp. 195–199.
- [34] Hao Zhang et al. “Predictive learn and apply: MAVIS application-learn”. In: *Adaptive Optics Systems VII*. Vol. 11448. SPIE. 2020, pp. 509–516.