

On-sky results for adaptive optics control with data-driven models on low-order modes

Baptiste Sinquin¹,¹★ Léonard Prengère,^{1,2} Caroline Kulcsár,¹★ Henri-François Raynaud,¹ Eric Gendron,³ James Osborn⁴,⁴ Alastair Basden,⁵ Jean-Marc Conan,² Nazim Bharmal,⁴ Lisa Bardou,⁴ Lazar Staykov,⁴ Tim Morris,⁴ Tristan Buey,³ Fanny Chemla⁶ and Matthieu Cohen⁶

¹Laboratoire Charles Fabry, Institut d'Optique Graduate School, CNRS, Université Paris-Saclay, 91120 Palaiseau, France

²ONERA, DOTA, Université Paris-Saclay, F-92322 Châtillon, France

³Laboratoire d'Etudes Spatiales et d'Instrumentation en Astrophysique, Observatoire de Paris, 92190 Meudon, France

⁴Centre for Advanced Instrumentation, Durham University, South Road, DH1 3LG Durham, UK

⁵Institute for Computational Cosmology, Durham University, South Road, DH1 3LG Durham, UK

⁶Galaxies Etoiles Physique et Instrumentation, Observatoire de Paris, 92190 Meudon, France

Accepted 2020 August 18. Received 2020 August 18; in original form 2020 January 29

ABSTRACT

Dedicated tip–tilt loops are commonly implemented on adaptive optics (AO) systems. In addition, a number of recent high-performance systems feature tip–tilt controllers that are more efficient than the integral action controller. In this context, linear–quadratic–Gaussian (LQG) tip–tilt regulators based on stochastic models identified from AO telemetry have demonstrated their capacity to effectively compensate for the cumulated effects of atmospheric disturbance, windshake and vibrations. These tip–tilt LQG regulators can also be periodically retuned during AO operations, thus allowing to track changes in the disturbances' temporal dynamics. This paper investigates the potential benefit of extending the number of low-order modes to be controlled using models identified from AO telemetry. The global stochastic dynamical model of a chosen number of turbulent low-order modes is identified through data-driven modelling from wavefront sensor measurements. The remaining higher modes are modelled using priors with autoregressive models of order 2. The loop is then globally controlled using the optimal LQG regulator build from all these models. Our control strategy allows for combining a dedicated tip–tilt loop with a deformable mirror that corrects for the remaining low-order modes and for the higher orders altogether, without resorting to mode decoupling. Performance results are obtained through evaluation of the Strehl ratio computed on *H*-band images from the scientific camera, or in replay mode using on-sky AO telemetry recorded in 2019 July on the CANARY instrument.

Key words: turbulence – atmospheric effects – instrumentation: adaptive optics – methods: data analysis – methods: observational – telescopes.

1 INTRODUCTION

Adaptive optics (AO) allows the real-time correction of wavefront aberrations induced by a time-varying perturbation using a deformable mirror. Such systems have been implemented in some ground-based telescopes to strive for diffraction-limited images (Rousset et al. 1990). The standard control method is the integrator, whose main assets are the high rejection at low temporal frequencies and the little tuning required (one or two scalar gains for tip–tilt and higher orders, or Optimal Modal Gain Integrator as in Gendron & Léna 1994, Rousset et al. 2003). As wavefront perturbations induced by windshake and telescope vibrations may significantly modify the spectrum of the disturbance wavefront, the performance of the integrator will be altered as vibrations frequently appear at frequencies that correspond to its rejection overshoot (i.e. the

domain where the controller amplifies the disturbance). A model-based alternative is the linear–quadratic–Gaussian (LQG) controller (Le Roux et al. 2004) that relies on a quadratic cost function in order to minimize the variance of the residual wavefront. A dynamical model of the stochastic disturbance is then required and is of prime importance to the closed-loop performance. It can incorporate priors about the turbulence such as wind speed and seeing, but may also include data-driven models.

It has been shown first in laboratory experiments (Petit et al. 2008), then on-sky at the William Herschel Telescope (WHT; Sivo et al. 2014), that a full LQG controller with a data-driven vibration and turbulence model on the tip–tilt modes (Meimon et al. 2010) combined with a model based on priors for higher order modes is able to efficiently compensate these additional perturbations. A data-driven controller for the tip–tilt combined with an integrator was further demonstrated for imaging the Sun at the McMath–Pierce solar telescope (Doelman, Fraanje & Breeje 2011). Tesch et al. (2015) used also a data-driven model on low-order Karhunen–Loève (KL) modes and a leaky integrator on sky at Palomar Observatory.

* E-mail: baptiste.sinquin@institutoptique.fr (BS); caroline.kulcsar@institutoptique.fr (CK)

This strategy of embedding a model based on recent telemetry data into a low-order modes LQG controller is widespread and has been further used for instruments such as the Spectro-Polarimetric High-contrast Exoplanet REsearch (SPHERE; Petit et al. 2009) and the Gemini Planet Imager (Poyneer et al. 2016), and for the Subaru Coronagraphic Extreme Adaptive Optics (SCExAO) systems (Lozi et al. 2016).

The identification of the model for the low orders impacts the closed-loop performance all the more as these have the biggest contribution to the temporal error. A number of model identification strategies have been investigated. Subspace identification algorithms (Overschee & Moor 1994) have been used for modelling tip-tilt, as in Doelman et al. (2011), or low-order KL, as in Tesch et al. (2015). They can also be used as an initialization step for prediction error methods, see e.g. Kulcsár et al. (2012b) for a cascaded version that sparsifies the model matrices. Grey-box methods (Meimon et al. 2010; Yang et al. 2018) describe each vibration with a mass-spring-damper system. The scope of this contribution is however not to determine the best of these, and we refer to Juvénal et al. (2015) for tests using on-sky data in this direction. Further research include Agapito et al. (2012) and Guesalaga et al. (2013) for \mathcal{H}_∞ -based control strategies, and Correia, Véran & Herriot (2012) for a multirate sampling approach of the LQG.

A number of questions nevertheless arise such as the ability of models identified from wavefront sensor data to capture the dynamics of the disturbance, including dome turbulence. Their spatiotemporal statistics are indeed likely to be very different from the ones of the atmospheric layers at ground level and in altitude. Moreover, vibrations on higher order modes than the tip and tilt were observed on the Nasmyth Adaptive Optics System (NAOS; Rousset et al. 2003). Would a rejection using a tailored controller further improve the performance? Along this line, Is the model postulated in Sivo et al. (2014) suited for other behaviours than pure boiling-type atmosphere? Frozen-flow behaviours may indeed be partly missed as the model is independent from the wind direction. This is both a strength and a drawback, because on the one hand it is more robust to temporally varying conditions as fewer priors are required, but on the other hand further gain in performance may be expected from a model embedding a frozen-flow behaviour of the atmosphere (Pregère, Kulcsár & Raynaud 2020). All these arguments provide a motivation to further investigate how data-driven approaches can help in accurately modelling the disturbance, and how much performance can be gained by identifying the temporal dynamics of more modes than only the tip and tilt.

These questions are investigated in this paper through on-sky experiments conducted on the demonstrator CANARY at the WHT during an observing run in 2019 July. We tested controllers that include a data-driven model of an increased number of Zernike modes, up to 14. Further questions to investigate the impact of the turbulence conditions measured with the SCIntillation Detection And Ranging (SCIDAR) in Osborn (2018), such as the wind speed, are answered by replaying the wavefront sensor data in simulation. We analyse whether this modifies the stability margins and to what extent. These insights raise new questions for the development of the next generation of LQG controllers for extremely large telescopes, including ones merging data-driven dynamical models for a few modes, for example in a Zernike or KL basis, into a scalable control strategy.

This paper is organized as follows. Section 2 describes the single-conjugate adaptive optics (SCAO) mode of the demonstrator CANARY at the WHT. Section 3 summarizes the control algorithm and especially the dynamical models of the wavefront disturbance.

On-sky results are presented in Section 4 and the performance is further analysed in Section 5 from the rejection transfer functions. An error budget assessment for the proposed regulators used on sky on CANARY is proposed in Section 6. We discuss data-driven controllers in Section 7 and conclude in Section 8.

2 SCAO ON CANARY

The AO bench CANARY is set up on a Nasmyth platform of the 4.2-m WHT on La Palma, Canary Islands, Spain (Myers et al. 2008). Its SCAO mode features an on-axis Shack–Hartmann sensor with 14×14 lenslets, for which the image scale is $0.428 \text{ arcsec pixel}^{-1}$. There are 144 active lenslets leading to 288 measurements. Both a tip-tilt mirror and an ALPAO deformable mirror (coupling factor of 0.45) with 243 active actuators are used to correct the incoming wavefront disturbance. The near-infrared (IR) camera CAMICAZ, developed by Observatoire de Paris, has a Near Infrared Camera and Multi-Object Spectrometer-3 (NICMOS-3) array with 256×256 pixels and the observations were restricted to the H spectral band with a central wavelength, λ_H , of $1.667 \text{ }\mu\text{m}$. The real-time computer is Durham AO real-time controller (DARC; Basden et al. 2010). During this run, the control frequency was set to 200 Hz. A loop delay of 2.25 frames was evaluated by poking the mirror and measuring the time delay before observing non-zero slopes on the Shack–Hartmann.

3 DESCRIBING THE LQG REGULATOR

3.1 Performance criterion and optimal controller

The control objective is to minimize the variance of the residual phase $\phi_k^{\text{res}} = \phi_k + \phi_k^{\text{cor}}$, where ϕ_k and ϕ_k^{cor} are vectors representing, respectively, the turbulent wavefront and the correction wavefront generated by the deformable and tip-tilt mirrors at discrete time index k . Assuming an infinitely fast mirror response, the correction wavefront is expressed as a function of the actuators commands as $\phi_k^{\text{cor}} = Nu_{k-1}$. The optimal solution in terms of control inputs u_k minimizes the cost function J ,

$$J(u) = \min_u \lim_{K \rightarrow \infty} \frac{1}{K} \sum_{k=1}^K \|\phi_{k+1} + Nu_k\|_2^2. \quad (1)$$

This minimization problem is solved using the separation principle. It states that the spatiotemporal prediction of the disturbance wavefront can be carried out independently of the projection on to the mirror's space, leading to the optimal control,

$$u_k = -(N^T N)^{-1} N^T \hat{\phi}_{k+1|k}, \quad (2)$$

where $\hat{\phi}_{k+1|k}$ is the minimum variance estimate of the wavefront at time instant $k+1$ using all the data up to time k , see Kulcsár et al. (2012a) for a detailed description of minimum-variance control in AO. The experiments we carried out relied on a modification of the criterion (1), which rather minimizes the residual slopes so as to formulate the solution (2) using the interaction matrix M_{int} instead of the influence matrix N , as proposed in Sivo et al. (2014),

$$\bar{J}(u) = \min_u \lim_{K \rightarrow \infty} \frac{1}{K} \sum_{k=1}^K \|D\phi_{k+1} + M_{\text{int}}u_k\|_2^2, \quad (3)$$

where D is the wavefront sensor matrix. The solution is expressed as a function of the command matrix M_{com} (the pseudo-inverse of the interaction matrix),

$$u_k = -M_{\text{com}} D \hat{\phi}_{k+1|k}. \quad (4)$$

This modification of $J(u)$ bypasses the problem of estimating the transformation between the wavefront sensor and the influence matrix N , such as rotations or magnifications between the wavefront sensor and actuator grids. The methodology for wavefront prediction is unchanged, only the projection on the mirror space is modified, which therefore allows extension to multiconjugate AO schemes (Sivo et al. 2013).

We are now left with the problem of predicting the one-step ahead turbulent wavefront, $\hat{\phi}_{k+1|k}$.

3.2 Disturbance stochastic dynamical model

The stochastic dynamical model for the disturbance has a significant impact on the performance of the AO system. We first express the spatiotemporal dynamics of the global turbulence-induced wavefront disturbance with an autoregressive (AR) model of order two (Sivo et al. 2014),

$$\begin{cases} \phi_{k+1} = A_1\phi_k + A_2\phi_{k-1} + Q^{1/2}v_k, \\ y_k = D\phi_{k-1} + w_k, \end{cases} \quad (5)$$

where the wavefront ϕ_k is expressed in a Zernike basis with 495 modes. The signal $\{v_k\}$ is a Gaussian process noise with covariance matrix $Q \geq 0$, $\{w_k\}$ is a zero-mean white Gaussian measurement noise with diagonal covariance matrix Σ_w , and y_k are open-loop slopes. Matrices A_1 and A_2 are set as in Sivo et al. (2014) and depend on the Fried parameter r_0 and on the resultant wind speed norm.

We introduce a linear time-invariant stochastic model describing the low-order (LO) modes only,

$$\begin{cases} x_{k+1}^{\text{LO}} = A^{\text{LO}}x_k^{\text{LO}} + Q^{\text{LO}1/2}v_k, \\ \phi_k^{\text{LO}} = C^{\text{LO}}x_k^{\text{LO}} + \eta_k, \end{cases} \quad (6)$$

where x_k^{LO} is a state without particular physical interpretation and ϕ_k^{LO} is a vector containing the coefficients of the first n_{LO} modes. This state equation for the model (6) is concatenated with the one of equation (5) and the wavefront reaching the Shack–Hartmann is modelled as the sum of ϕ_k and ϕ_k^{LO} for the low orders and only ϕ_k for all the other ones, yielding

$$\begin{cases} \begin{bmatrix} \phi_{k+1} \\ \phi_k \\ x_{k+1}^{\text{LO}} \end{bmatrix} = \underbrace{\begin{bmatrix} A_1 & A_2 & 0 \\ I & 0 & 0 \\ 0 & 0 & A^{\text{LO}} \end{bmatrix}}_A \begin{bmatrix} \phi_k \\ \phi_{k-1} \\ x_k^{\text{LO}} \end{bmatrix} + \begin{bmatrix} Q^{1/2} \\ 0 \\ Q^{\text{LO}1/2} \end{bmatrix} v_k, \\ y_k = \underbrace{\begin{bmatrix} 0 & D & D^{\text{LO}} \end{bmatrix}}_C \begin{bmatrix} \phi_k \\ \phi_{k-1} \\ x_k^{\text{LO}} \end{bmatrix} + w_k, \end{cases} \quad (7)$$

where D^{LO} is the sensor matrix for the first n_{LO} modes. Let the covariance matrices of the process and measurement noise be defined as

$$S_v = \begin{bmatrix} Q & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & Q^{\text{LO}} \end{bmatrix}, \quad (8)$$

$$S_w = f\Sigma_w, \quad (9)$$

where $f > 0$ is a so-called fudge factor that is set slightly larger than 1 to cope with uncertainties in the modelling of the dynamic equation (7). For example, the Fried parameter and wind speed may vary over time and their estimation is prone to errors. Increasing this fudge factor leads to more stable controllers without much impact on the performance (Petit et al. 2009; Parisot et al. 2012).

The Kalman gain L_∞ associated with the matrices (A, C, S_v, S_w) is subsequently computed solving the associated discrete algebraic Riccati equation (DARE). Let the state of the controller \tilde{x}_k be partitioned as $[x_k^T \ x_{k-1}^T]^T$, where $x_k = [\phi_k^T \ \phi_{k-1}^T \ x_k^{\text{LO}T}]^T$. The controller is obtained in state-space form as

$$\begin{cases} \tilde{x}_{k+1} = \begin{bmatrix} A - L_\infty C & -L_\infty M_{\text{int}} P \\ I & 0 \end{bmatrix} \tilde{x}_k + \begin{bmatrix} L_\infty \\ 0 \end{bmatrix} y_k^{\text{res}}, \\ u_k = P[A - L_\infty C \quad -L_\infty M_{\text{int}} P] \tilde{x}_k + PL_\infty y_k^{\text{res}}, \end{cases} \quad (10)$$

where $P = -M_{\text{com}}[D \ 0 \ D^{\text{LO}}]$, and y_k^{res} are the closed-loop slopes.

The control methodology we propose mainly differs from Sivo et al. (2014) on two counts: first, the tip and tilt are not computed from the mean slopes in the horizontal and vertical directions, but are estimated by projection on the Zernike basis using a minimum-variance reconstructor; second, the identification is not limited to two Zernike modes. This has the advantage of modelling not only the vibrations but also the dome turbulence, and possibly wind-induced frozen-flow behaviours. The model relying on priors still represents the spatiotemporal evolution of the 495 modes, even though the data-driven model already describes the behaviour of the first ones. Coupling between the two parts is not ignored since the Kalman gain is computed globally for the whole model (7).

Note that the controller (10) is designed for a delay of 2 frames instead of 2.25, which slightly degrades the performance with respect to the case when this true delay is taken into account (Sivo et al. 2014).

3.3 Identifying the state-space matrices for the low-order modes

As already mentioned, the matrices A_1, A_2 , and Q in equation (5) are defined as in Sivo et al. (2014). We now show how the matrices that describe the behaviour of the low-order modes in equation (6) are estimated using the subspace identification method in Overschee & Moor (1994). Let us assume that a sequence with N_t temporal samples of open-loop (or pseudo open-loop) slopes y_k has been collected. We first reconstruct the wavefront on a Zernike basis with the minimum-variance static estimator,

$$Z_k = C_\phi D^T (DC_\phi D^T + \Sigma_w)^{-1} y_k, \quad (11)$$

where C_ϕ is the theoretical spatial covariance of the wavefront. Although Z_k contains 495 coefficients to avoid aliasing on the lower orders in the wavefront reconstruction, only the n_{LO} first ones are selected in the forthcoming algorithm. The tuple $Z_k(1:n_{\text{LO}})$ is denoted with z_k .

We strive for more generality than the approaches for tip-tilt mode identification described in Section 1. This approach should hold for any value of n_{LO} below few hundreds, and subspace identification methods are therefore good candidates. These are cost-function free black-box approaches that are more scalable than prediction error methods none the less yielding consistent estimates of the matrices with increasing data set length, N_t . The identification algorithm numerical subspace state space system identification (N4SID; Overschee & Moor 1994) is described in more details in Appendix A. It relies on the model (6) written in innovation form

$$\begin{cases} \hat{x}_{k+1|k}^{\text{LO}} = \tilde{A}^{\text{LO}} \hat{x}_{k|k-1}^{\text{LO}} + L_\infty^{\text{LO}} z_k, \\ \hat{\phi}_k^{\text{LO}} = C^{\text{LO}} \hat{x}_{k|k-1}^{\text{LO}}, \end{cases} \quad (12)$$

where $\tilde{A}^{\text{LO}} = A^{\text{LO}} - L_\infty^{\text{LO}} C^{\text{LO}}$. It exploits a low-rank property of a matrix built from data, which allows to estimate both the extended observability matrix and a state sequence $X_{s,N}$. This is subsequently

used to estimate the matrices \tilde{A}_{LO} and C_{LO} by first solving the least squares,

$$\min_{\tilde{A}_{LO}, C_{LO}} \left\| \begin{bmatrix} \hat{X}_{s+1,N} \\ \Phi_{s,N}^{LO} \end{bmatrix} - \begin{bmatrix} \tilde{A}_{LO} \\ C_{LO} \end{bmatrix} \hat{X}_{s,N} \right\|_F^2, \quad (13)$$

and then estimating the covariance matrix Q^{LO} from the least-squares residual. We refer to Appendix A for the detailed algorithm.

We note that this additional computational complexity for identifying a model is not an issue for a few hundreds of modes and $n_{LO} = 24$ was the maximum value for our on-sky experiments. The system order for the model in equation (7) is equal to $2 \times 495 + n$, where n was ranging from 30 to 100 in these examples, implying no significant increase of the cost for solving the DARE associated with equation (7) nor of the online computational cost for the controller (10).

4 ON-SKY RESULTS

4.1 The methodology

4.1.1 On-sky data acquisition procedure and performance evaluation

The loop is first closed on the tip and tilt using an integrator with a low gain of 0.3 on these modes only (the gain on the higher orders is set to 0). The centroids then remain within the dedicated boxes in the Shack–Hartmann sensor. The models for describing the temporal dynamics of the low orders are identified using the first 4000 samples and the 1000 remaining ones are used to evaluate the control performance, in order to check that the prediction error for the particular model (12) is indeed small. These values are mostly indicative and may vary depending on the size of the buffers and of the wavefront sensors. Once the matrices in the disturbance model (7) are known, the DARE is solved and the controller (10) computed, and eventually loaded into DARC. Because of the large variability of atmospheric conditions, scripts with two or three controllers interleaved were executed for comparison: approximately 20 s (4096 samples) for each controller repeated 10 times.

Performance is evaluated in terms of the Strehl ratio, computed from images of the scientific camera. Strehl ratios are plotted as a function of time or of the Fried parameter. The Fried parameter r_0 is estimated from pseudo-open-loop slopes by fitting the theoretical variances of the Zernike decomposition of the reconstructed wavefront to the Kolmogorov spectrum. The least-squares fit involves only modes 3–135. Under the assumption of frozen-flow turbulence with a constant wind speed, the residual phase variance σ_{res}^2 for a controller operating at a frequency f_s satisfies an affine relationship with $r_0^{-5/3}$ (Roddier 1999):

$$\sigma_{res}^2 \propto \alpha + \beta r_0^{-5/3}. \quad (14)$$

The residual phase variance can be deduced from the Strehl ratio when large enough (>0.3) using the Maréchal approximation for small aberrations.

No vibration was observed in the spectra for the modes higher than the tip and the tilt in any of the data sets.

4.1.2 Replaying on-sky data

Closed-loop data from the on-sky experiments are reused a posteriori to evaluate different controllers and enable further comparisons. In replay mode, the first 50 s of pseudo-open-loop data are used for

Table 1. (Simulation) Residual phase variance σ_{res}^2 (in rad^2) at the wavelength of the science camera using OMAO. The LQG method tested here follows equation (10) without data-driven model ($n_{LO} = 0$).

	Integrator	LqgZer	Rel. imprvt
Wind = 10 m s ⁻¹ and $r_0 = 20$ cm			
True σ_{res}^2	0.19	0.13	−32%
σ_{res}^2 from the closed-loop slopes	0.13	0.09	−31%
Wind = 10 m s ⁻¹ and $r_0 = 10$ cm			
True σ_{res}^2	0.58	0.38	−34%
σ_{res}^2 from the closed-loop slopes	0.33	0.20	−39%

identifying the models. These data sets are discarded for closed-loop simulations. We highlight that this is not what was tested on-sky: the data used for model identification may have been collected a few hours before the controller is tested (the exact times are mentioned in the next section). Using the on-sky data, we compare the proposed approach to the integrator and to the strategy in Sivo et al. (2014). After reconstructing the pseudo-open-loop slopes y_k^{POL} ,

$$y_k^{\text{POL}} = y_k^{\text{res}} - (0.75z^{-2} + 0.25z^{-3})M_{\text{int}}u_k \quad (15)$$

(taking into account the true system delay), we test different controllers and assess their closed-loop performance. To do so, the residual wavefront is reconstructed with a weighted least squares,

$$\hat{\phi}_k^{\text{res}} = (D^T D + \Sigma_w)^{-1} D^T y_k^{\text{res}}, \quad (16)$$

and the performance is evaluated as the sum over the variances of the first 197 modes (these correspond to the mirror modes) at the wavelength of the imaging camera. This criterion is such that fitting and aliasing errors are underevaluated. The fitting error is the same for all controllers, whereas the aliasing error has a much lower contribution than the temporal error for the current configuration of CANARY, thanks to atmosphere sampling with a 14×14 Shack–Hartmann sensor for a 4.2-m telescope. The temporal error is predominant in the performance criterion used. For a discussion of the pitfalls in replay mode, we refer to Kulcsár et al. (2018).

We have validated this criterion evaluation procedure using the simulation software OMAO(Conan & Correia 2014) on a CANARY configuration by comparing with the true residual phase variance. In replay mode, the difference in residual phase variance between the integrator and the LQG method decreases in general, as shown in Table 1. The relative improvement between both is better preserved for large r_0 (−32 per cent versus −31 per cent) but stays close even for small r_0 (−34 per cent versus −39 per cent). Note that the smaller r_0 , the larger the temporal and aliasing errors, the latter being especially reduced with an LQG controller (Juvénal et al. 2018). Results in Table 1 show that the procedure is valid for control performance comparison.

4.2 Case 1: large wind speeds

We report on the experiments from the night of 2019 July 18, between 3h08min and 4h21min (local time). The guide star has magnitude 5.2 in H band. Main layers at 0, 2, and 8 km were measured with a wind speed of respectively 12, 15, and 13 m s⁻¹ using the stereo-SCIDAR (Osborn 2018). Wind profiles corresponding to this night are shown in Fig. 1.

We first present the results when there is a data-driven model for the tip and tilt only. This allows to confirm the results in Sivo et al. (2014) and check whether the identification algorithm has a significant impact on the performance. When such a model is computed from the




Figure 1. Altitude profile of the turbulence strength factor $C_n^2 dh$ as a function of time for the night from 2019 July 17 to 18. The length of the arrow in the upper left-hand side in the north-east compass of the figure at the bottom is equivalent to 10 m s^{-1} for scale.




Figure 2. (On-sky) Strehl ratio as a function of time for the experiments from 03h08min to 03h58min on 2019 July 18. LqgZer-TT (N4SID) is a controller relying on the structure (7) that identifies the tip and tilt from the mean slopes using N4SID. The identification method in LqgZer-TT (cPEM) is the cascaded PEM.

mean slopes (for sake of comparison), we denote the controller with LqgZer-TT. If computed from a minimum variance reconstructor, it is denoted with LqgZer-LO2. The cascaded Prediction Error Method (cPEM; Kulcsár et al. 2012b) and N4SID (Overschee & Moor 1994) were compared as a first test for identifying only the dynamics of the tip and tilt signals. The data set used for the identification was collected at 2h02min17s. Figs 2 and 3 display the Strehl ratio computed from the IR images as a function of the time and of the Fried parameter, respectively. Fig. 4 plots the residual phase variance computed using equation (14) as a function of $r_0^{-5/3}$ to fit affine trends. There is an important gap in performance between the integrator and either of the LQG controllers, as already shown in Sivo et al. (2014). In our experiments, an average of seven points of Strehl ratio performance improvement has been measured. Increasing the gain of the integrator to larger values than 0.4 caused many instabilities




Figure 3. (On-sky) Strehl ratio as a function of r_0 for the experiments from 03h08min to 03h58min on 2019 July 18.




Figure 4. (On-sky) Residual phase variance $\sigma_{\text{res}}^2 = -\ln(\text{SR})$ as a function of $r_0^{-5/3}$ for the experiments from 03h08min to 03h58min on 2019 July 18. Results corresponding to Fig. 3. The residual least-squares error when fitting linear models is, respectively, 0.051, 0.066, and 0.067. The two points with $r_0^{-5/3} > 55$ were discarded for fitting the models.

Table 2. (Replay) Variance of residual wavefront in nm rms at the science camera wavelength in replay mode. LqgZer-LO n_{LO} is a controller (7) with n_{LO} modes relying on a data-driven temporal model.

Start time	3h08min35s	4h06m13s
Data length (s)	800	450
Integrator	227.79	232.58
LqgZer-TT (N4SID)	193.76	196.04
LqgZer-LO2	189.72	193.65
LqgZer-LO5	185.64	193.14
LqgZer-LO9	185.00	192.02
LqgZer-LO14	183.37	191.16
LqgZer-LO20	182.16	190.55
LqgZer-LO35	179.18	188.79
LqgZer-LO65	175.52	187.30

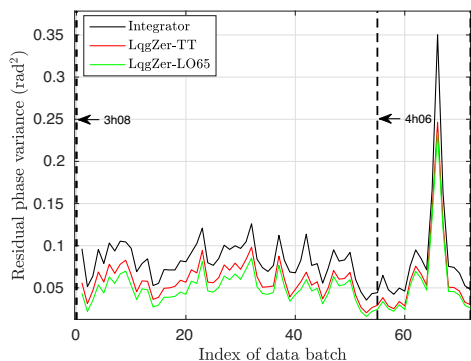


Figure 5. (Replay) Residual phase variance as a function of the time for three controllers when replaying on-sky data. Each point in the x -axis corresponds to the residual variance for a data batch of 20 s collected on 2019 July 17. For each black dashed line, a new data-driven model is identified leading to an updated controller.

on-sky. An integrator with optimized modal gains (Gendron & Léna 1994) was not tested because the star has a low magnitude for AO. The Strehl ratio is slightly larger using cPEM rather than N4SID. Such little differences are an incentive to investigate whether identifying the dynamics of higher order modes makes a further gap.

We have then evaluated the performance with low-order modelling for different values of n_{LO} replaying on-sky data. Table 2 displays the results obtained in replay mode, and Fig. 5 shows the residual phase variance as a function of time for three controllers. The more modes with a data-driven model for their temporal dynamics, the lower the variance of the residual wavefront. In case of large wind speed, the improvements are not limited to the first few radial orders, although they are smaller than between integrator and tip-tilt only rejection with a LQG controller. The frozen-flow behaviour is thus captured by the data-driven model and complements the boiling type of model in the AR2 built from priors (5).

In conclusion of these large wind speed results, we can say that as expected, the most important gap in improvement is when switching from an integrator to a tip-tilt LQG. The gain is smaller but still present when identifying more low orders and would increase for larger wind speed values.

Subspace identification does not rely on a cost function but on structural properties of equations built from data. It is therefore not possible to identify the matrices setting a prior bound on the accuracy. The accuracy is checked a posteriori and computed on a validation data set that is different from the one that is used for the identification. We show in Fig. 6 the variance of the residual wavefront as a function

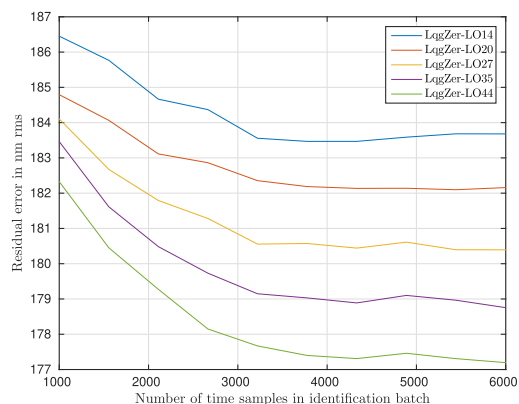


Figure 6. (Replay) Standard deviation of residual error (nm rms) evaluated with 197 Zernike modes as a function of the number of time samples used in the identification batch replaying the data sets collected on sky in case of large wind speed.

of the number of time samples used for the identification that provides guidelines on how to choose the length of the data set. The more modes to identify, the larger the data set. In this example, identifying the dynamics with more than 3000 temporal samples (15 s) for 14 modes is enough, whereas 4000 samples (20 s) are necessary when identifying dynamics for 44 modes.

4.3 Case 2: low wind speed

This section expands on the experiments from the night of July 21, from 4h31min to 6h. The guide star has magnitude 5.8 in H band. The wind speed at ground level was no more than 3 m s^{-1} . Wind profiles corresponding to this night are shown in Fig. 7. These are challenging conditions to show significant improvements using LQG controllers, which are known to perform better when the wind speed at the ground and in altitude is large.

Figs 8 and 9 compare the integrator (gain of 0.4) with the proposed LQG controller (7) with nine modes identified from a pseudo-open-loop data set of slopes collected at 02h13m46s. LqgZer-LO9 improves the Strehl ratio over the integrator by 6 per cent on average. A comparison with the LQG that only identifies the tip and tilt temporal dynamics is shown in Fig. 10 in replay mode. The larger r_0 , the lower the temporal error, the less to be gained (in absolute values) identifying further dynamics and increasing n_{LO} . This is expected as the temporal error is a major contribution in the residual error. The systematic improvement brought by LqgZer-LO9 over LqgZer-TT is smaller than in Fig. 2 because the wind speed is smaller. However, the behaviour in terms of stability was far better with the former. This is studied further in Section 5.

Figs 11 and 12 from on-sky results show that increasing the number of modes n_{LO} increases the Strehl ratio. The spread of the LqgZer-LO9 is less than the one of LqgZer-LO14. The trend that nine modes would be sufficient is confirmed using linear fittings in Fig. 12, where the relative improvement is more significant from 5 to 9 than from 9 to 14.

Table 3 and Fig. 13 further show the superiority of the controller (10) for all replayed data. Similarly to the on-sky results in Fig. 11, increasing n_{LO} marginally increases the performance. This is due to the fact that the turbulence is essentially a boiling type for which the model built on priors is very much suited. The variance of the residual wavefront with the controller (10) is all the lower compared to the other ones shown in Table 2 when the wind speed is large.




Figure 7. Altitude profile of the turbulence strength factor $C_n^2 dh$ as a function of time for the night from 2019 July 20 to 21. The length of the arrow in the upper left-hand side in the north-east compass of the figure at the bottom is equivalent to 10 m s^{-1} for scale.




Figure 8. (On-sky) Strehl ratio as a function of time for the experiments from 04h31m56s to 04h43m35s on 2019 July 21. LqgZer-LO9 stands for the controller (7) with nine modes identified.




Figure 9. (On-sky) Strehl ratio as a function of the Fried parameter for the experiments from 04h31m56s to 04h43m35s on 2019 July 21.




Figure 10. (Replay) Residual phase variance σ_{res}^2 in rad^2 at the science camera wavelength as a function of $r_0^{-5/3}$ for the experiments from 04h31m56s to 04h43m35s on 2019 July 21, replaying the data with simulations. The values are small because they come from simulation replaying the on-sky data, and they are not evaluated using the Maréchal approximation.




Figure 11. (On-sky) Strehl ratio as a function of the time occurrence for the experiments from 04h51m36s to 05h09m11s on 2019 July 21.




Figure 12. (On-sky) Residual phase variance σ_{res}^2 in rad^2 at the science camera wavelength as a function of $r_0^{-5/3}$ for the experiments from 04h51m36s to 05h09m11s on 2019 July 21. Results corresponding to Fig. 11. These are obtained from the Strehl ratio and using the Maréchal approximation. The residual least-squares error when fitting linear models is, respectively, 0.025, 0.016, and 0.038.

Table 3. (Replay) Variance of residual wavefront in nm rms at the science camera wavelength in replay mode.

Start time	4h31m54s	4h51m35s	5h14m01s	5h32m46s
Data length (s)	450	700	700	450
Integrator	167.30	184.00	197.15	152.06
LqgZer-TT (N4SID)	154.23	170.78	187.74	146.81
LqgZer-LO2	151.38	169.30	183.11	139.00
LqgZer-LO5	150.67	168.51	182.80	138.56
LqgZer-LO9	150.17	167.46	181.29	137.79
LqgZer-LO14	149.51	167.26	180.78	137.44




Figure 13. (Replay) Residual phase variance as a function of the time for three controllers when replaying on-sky data. Each point in the x-axis corresponds to the residual variance for a data batch of 20 s collected on 2019 July 21. For each black dashed line, a new data-driven model is identified leading to an updated controller.

5 CONFIRMING ON-SKY RESULTS WITH NUMERICAL SIMULATIONS

We present further simulations using the numerical software OMAO (Conan & Correia 2014) to analyse the impact of both the wind speed and the variance of the measurement noise. No vibration is added at all. Three types of regulators are compared: the integrator (with a gain ranging between 0.4 and 0.7 and the best Strehl ratio is kept); the LQG regulator built only from priors; and the LQG regulator with both priors and a data-driven model on a given number of Zernike




Figure 14. (Numerical simulation OMAO) Strehl ratio as a function of the wind speed for a single atmosphere layer with Fried parameter of 10 cm and variance of the measurement noise of 0.2 rad^2 .




Figure 15. (Numerical simulation OMAO) Strehl ratio as a function of the variance of the measurement noise for a single atmosphere layer with Fried parameter of 10 cm and wind speed of 2 m s^{-1} .

modes. A single atmosphere layer on the ground is considered with a Fried parameter of 15 cm.

As shown in Figs 14 and 15, if both the variance of the measurement noise is low (e.g. in case of a bright guide star) and the wind speed norm is below 5 m s^{-1} , then there is hardly any difference between the LQG and the integrator. This was already shown in Sivo et al. (2014): when the temporal error is small in low wind speed scenario, the main advantage of the LQG framework is to embed a data-driven model for the vibrations and reach better performance in case of faint guide stars. Adding a data-driven model on Zernike modes not limited to the tip-tilt improves the performance of the LQG all the more as the wind speed is large, which confirms what has been observed on-sky and replaying the data. The gap in performance by adding further modes in the data-driven model widens for increasing wind speed.

6 TRADE-OFF BETWEEN ROBUSTNESS AND PERFORMANCE

As noted above, the on-sky performance of the regulators in terms of robustness in stability was different and is studied together with the analysis of the transfer functions relative to the controllers used on-sky for an atmosphere with *low wind speed*, which reveals further insights. The rejection transfer function is defined as

$$T_{\text{rej}}(z) = D^\dagger (I - \delta(z) M_{\text{int}} G(z))^{-1} D, \quad (17)$$




Figure 16. Modulus of the rejection transfer function as a function of the frequency for the tip. The normalized PSD of pseudo-open-loop data displays the vibration peaks.




Figure 17. Modulus of the rejection transfer function as a function of the frequency for the tilt. The normalized PSD of pseudo-open-loop data displays the vibration peaks.

where G is the transfer function from the slopes to the commands (associated with equation 7), and $\delta(z) = 0.75z^{-2} + 0.25z^{-3}$ is the delay transfer function. These are shown in modulus for the tip, tilt, focus, and a trefoil in Figs 16, 17, 18, and 19 along with the normalized power spectral density (PSD) of POL data. It can be seen that there are only vibrations peaks for the tip and tilt. The rejection transfer functions differ when the dynamics for the tip and tilt are identified from the mean slopes or with a maximum a posteriori reconstruction: the rejection of the vibration peaks is larger for the former, at the expense of a worse attenuation between 1 and 10 Hz. Increasing n_{LO} reshapes the rejection so that the maximum of attenuation is shifted by about 10 Hz compared to the controller with a data-driven tip-tilt model only, as highlighted in Fig. 20. These trends are also to be observed when the wind speed is large, and the rejection in low frequencies is also modified using models identified from data.

The spectral radius of the matrix $A-L_{\infty}C$ in equation (7) describes how much the Kalman filter remembers previous data: the larger, the less responsive the filter is to fluctuations in the output although the more stable in terms of e.g. delay margins. The trend is that increasing




Figure 18. Modulus of the rejection transfer function as a function of the frequency for the focus.




Figure 19. Modulus of the rejection transfer function as a function of the frequency for a trefoil.




Figure 20. Frequency of the maximum of rejection of the controllers as a function of the Zernike mode.

the number of modes identified decreases the spectral radius, from 0.9967 for LqgZer-LO2 to 0.9954 for LqgZer-LO20.

The Nyquist diagram allows to determine the gain and phase margins from the open-loop transfer function, i.e. respectively the additional gain at 180° and additional phase shift at unit gain to make the closed-loop system unstable. The stability margins slightly




Figure 21. Nyquist diagram for the focus.




Figure 22. Nyquist diagram for one trefoil.

decrease for each mode for the controllers (7), see the Nyquist diagrams in Figs 21 and 22. The phase margin for each of the first 14 modes tip-tilt excluded is approximately equal to 170° . It decreases by no more than 5° when a data-driven model is used on the mode of interest. The delay margins (i.e. largest additional delay for which the loop remains stable) for the tip and tilt are shown in Fig. 23. These still remain above one frame, which is already a very good delay margin. If a mode relies on both a model built on priors and data, then the delay margins reduce but all the regulators show good stability margins.

7 ERROR BUDGET

To further understand the performance of the proposed controllers and the potential to improve over these solutions, we now detail an error budget based on formulas in Juvénal et al. (2018), using simulated disturbance sequences.

The disturbance is generated using for the first 20 modes a stochastic LTI model identified from on-sky data, and for the following 475 modes an AR2 model (5) with the same Fried parameter. Let v be a realization of a multivariable white Gaussian noise with identity




Figure 23. Delay margins associated with the controller in Sivo et al. (2014) (indicated with a circle) and the ones of the controllers (7) as a function of the radial order of the data-driven model.

covariance matrix, then it is possible to generate a trajectory of the turbulent wavefront with a linear shaping filter $H(z)$ such that $\phi = H(z)v$. The transfer from ϕ to ϕ^{res} is

$$T_{\text{rej}}(z) = (I - \delta(z)NG(z)D)^{-1}. \quad (18)$$

The residual phase variance σ_{res}^2 is the sum of the contribution of the residual induced by the turbulent wavefront σ_{rej}^2 and of the one induced by the measurement noise propagated in the loop, σ_{noise}^2 , with

$$\sigma_{\text{rej}}^2 = \text{trace} \frac{1}{\pi} \int_0^\pi T_{\text{rej}} \mathcal{S}_\phi T_{\text{rej}}^H dw, \quad (19)$$

$$\sigma_{\text{noise}}^2 = \text{trace} \frac{1}{\pi} \int_0^\pi T_{\text{noise}} \Sigma_w T_{\text{noise}}^H dw. \quad (20)$$

The term σ_{rej}^2 quantifies how much the system rejects the turbulent wavefront and can be broken down into three terms: the temporal error σ_{temp}^2 , the aliasing error σ_{alias}^2 due to the WFS finite sampling, and the fitting error σ_{fit}^2 (which is independent of the controller). As in Juvénal et al. (2018), we introduce the projectors on the subspaces $\text{Im}(N)$ and $\text{Im}(N)^\perp$,

$$\mathcal{P}^{\parallel} = NN^\dagger, \quad (21)$$

$$\mathcal{P}^\perp = I - NN^\dagger, \quad (22)$$

in order to separate the two control-dependent contributions σ_{temp}^2 and σ_{alias}^2 and the control-independent term σ_{fit}^2 . The temporal and fitting errors are evaluated with equation (19) replacing T_{rej} by, respectively, $\mathcal{P}^{\parallel} T_{\text{rej}} \mathcal{P}^{\parallel}$ and $\mathcal{P}^\perp T_{\text{rej}} \mathcal{P}^\perp$. The aliasing error is then computed as

$$\sigma_{\text{alias}}^2 = \sigma_{\text{rej}}^2 - \sigma_{\text{temp}}^2 - \sigma_{\text{fit}}^2. \quad (23)$$

Table 4 summarizes the budgets, which highlights that the proposed controllers essentially reduce the temporal and aliasing error, while the noise propagation error increases.

The temporal error has the largest contribution, and decreases the most using controllers (7). The aliasing error decreases using an LQG controller, while it slightly fluctuates with the number of modes identified. This may be due to aliasing when reconstructing the wavefront on a Zernike basis. Further, the noise propagation error increases with LQG controllers whose state-space model order is larger as noticed in Juvénal et al. (2018).

Table 4. Error budget for CANARY for different controllers in nm rms at the science camera wavelength.

	σ_{temp}^2	σ_{fit}^2	σ_{alias}^2	σ_{noise}^2	σ_{res}^2
Integrator	214	116	58.1	20.6	271
LqgZer-TT (N4SID)	196	116	55.8	20.4	256
LqgZer-LO2	166	116	55.1	21.8	232
LqgZer-LO5	165	116	55.3	22.2	232
LqgZer-LO9	163	116	55.2	22.3	230
LqgZer-LO35	159	116	55.6	22.6	227
LqgZer-LO65	156	116	55.3	22.9	225

Table 5. Variance of residual wavefront in nm rms at the science camera wavelength in replay mode. This is the same data set as for Table 2.

Start time	3h08min35s	4h06m13s
Data length (s)	800	450
Data-driven controller	187.11	194.48

8 COMPARISON WITH A FULL DATA-DRIVEN CONTROLLER

It has been shown throughout this paper that a data-driven model on the low-order turbulence modes improves the performance all the more as the number of identified modes increases for large wind speed. This stems from the fact that first, the larger the wind speed, the larger the temporal error and thus the more LQG controllers improve with respect to the integrator; second, the model built from priors represents only a boiling behaviour of the turbulent wavefront, whereas the data-driven captures all dynamics including frozen flow. We suggest to evaluate the performance of full data-driven controllers replaying the on-sky data. For a stochastic model of the slopes in innovation form given with

$$\begin{cases} x_{k+1|k} = Ax_{k|k-1} + Ke_k, \\ y_k = Cx_{k|k-1} + e_k, \end{cases} \quad (24)$$

where e_k is the innovation, and for a matrix M that relates the wavefront and the estimated state with

$$\phi_k = Mx_{k|k-1}, \quad (25)$$

the LQG controller becomes

$$u_k = -M_{\text{com}}DMx_{k+1|k}, \quad (26)$$

in a very similar manner as in Hinnen, Verhaegen & Doelman (2007). The state-space matrices in equation (24) are first identified with N4SID on a wavefront sensor data batch containing 15×10^3 points; the closed-loop performance is then evaluated using both the data sets corresponding to low and large wind speeds. Tables 5 and 6 display the results (these have to be compared with, respectively, Tables 2 and 3: the same data set is used). When the wind speed is large, the full data-driven controller performs better than LqgZer-LO2 and worse than LqgZer-LO5. When the wind speed is small, it is similar to LqgZer-TT. The reason why such a controller performs

relatively better for the former case is because it is not restricted to boiling: it captures both the wind speed norm and direction. The state in equation (24) has no physical meaning, and is not associated with Zernike modes.

The state-space model identified with N4SID represents all the better the statistics of the disturbance as the number of temporal samples increases. There is no matrix structure to parametrize the state-space matrices, thus the model is able to capture boiling, frozen flow, dome turbulence, and vibrations simultaneously. The larger the number of outputs, the more parameters to estimate and the less accurate these are for a given number of samples. A drawback is however the computational complexity and the length of the data set required for obtaining such estimates. One way to address both of these points is to recursively update such models, which would also allow to cope with temporally varying atmospheres, as in Houtzager, van Wingerden & Verhaegen (2012). We outline the modularity of equation (7) that relies on a dedicated identification of the most energetic modes and the very moderate computational requirements of the controller (10) as assets.

9 CONCLUSIONS

We have investigated the impact of identifying the temporal dynamics of more modes than only the tip and tilt from data. There are unmodelled dynamics when only using priors in a model that represents essentially the boiling, for example a frozen flow including its direction and dome turbulence (and which is frequently ignored by lack of prior knowledge and specificity for each AO system). We have formulated a disturbance model that embeds both a data-driven model for the low orders and a model relying on priors for the higher orders. The performance was evaluated on-sky at the WHT using the CANARY bench and then by replaying the data for two kinds of atmospheric conditions.

We have shown that there is a gain in performance when identifying the temporal dynamics of more modes than only the tip and tilt even when no vibrations are present. When using subspace identification methods, the tip and tilt modes should be computed with a minimum variance reconstructor rather than with the mean slopes. The data-driven model for the low-order modes takes into account for example the dynamics of the dome turbulence, and all other ones neglected in the model (5). Increasing the number of modes to few dozen does not require much more computational resources, and it was shown on-sky and in replay that this improves performance. If the wind speed is large on the ground, it was shown that identifying a model for up to 65 Zernike modes reduces the residual phase variance, whereas for lower wind speeds, nine modes were shown to be sufficient for the data-driven model. These improvements in performance are however achieved at the expense of a reduction of the delay margins, which none the less remain well above one frame for the first modes. Other stability measures show very good stability that is coherent with what was observed on sky. The rejection transfer functions on the tip and tilt are identical for all the controllers LqgZer-LO and better Strehl ratios are thus achieved through a better modelling of the low orders, tip-tilt excluded. The budget error analysis has revealed that

Table 6. Variance of residual wavefront in nm rms at the science camera wavelength in replay mode. This is the same data set as for Table 3.

Start time	4h31m54s	4h51m35s	5h14m01s	5h32m46s
Data length (s)	450	700	700	450
Data-driven controller	152.21	173.60	187.00	143.09

the aliasing error has a minor impact in the current configuration of CANARY, and that most of the variance reduction stemmed from cutting down on the temporal error.

We highlight that the controller (7), for which the temporal dynamics of nine modes are identified from data, was readily used during the same run at the WHT to enhance the performance of an integral field spectrograph (Haffert et al. 2020), and exhibited excellent robustness.

During the on-sky experiments, we have used controllers that yielded good performance even though the data set was collected a few hours before and therefore, the statistics might have changed. However, we of course recommend for operational AO systems to regularly update the data sets and to identify new models (and hence update the controllers) to better cope with gradual changes of the disturbance. As an extension to the controllers proposed in this paper, frequently updating the data-driven models is the first step to cope with a slowly temporally varying atmosphere. The wind profiles in Figs 1 and 7 indeed demonstrate that the turbulence conditions evolve over the night, let alone the apparent pupil viewed by the sensor (which may feature derotations, etc.), whose orientation contributes to outdated data-driven models.

As for tomography, we can mention that data-driven temporal models for the tip and tilt are required. For all the other modes, the subspace-identification approach does not reconstruct the whole volume of the turbulence. It would be nevertheless possible to use a static reconstruction of the wavefront on each layer over a given temporal horizon and then compute models accordingly.

This contribution should be accounted for in the design of future AO regulators. Especially for the extremely large telescopes (ELTs), it clearly demonstrates that only controlling the tip and tilt with a data-driven model is not enough for optimal performance. How many modes should rely on a data-driven state-space model in ELTs? Following the guidelines in Conan et al. (2011), the radial order is multiplied by the ratio of the diameters (i.e. 10) to maintain a similar highest spatial frequency, which translates into 495 modes. A constraint on the number of modes to be identified is none the less set by the computational complexity for identifying the models from data. Keeping in mind that the complexity of subspace identification grows cubically with the number of outputs (here, number of modes in any basis, e.g. Zernike or KL), the number of modes may not be larger than a few hundreds.

The method shown in this paper is not restricted to Zernike modes and extends without changes to any other basis such as the KL one, which is used for example in Tesch et al. (2015). When combined with an integrator, the decoupling has to be done carefully, see e.g. Petit et al. (2010) in a Very Large Telescope (VLT) extreme adaptive optics (XAO) context. Controllers for ELTs AO systems that include data-driven models for more modes than only the tip and tilt are completely feasible and offer interesting perspectives in terms of performance and stability. Combining these data-driven models with either a prior-based dynamical model or with an integrator should therefore be investigated.

ACKNOWLEDGEMENTS

LQG tools for CANARY have been supported by the Agence Nationale de la Recherche (ANR) French program CHAPERSONA ANR-09-BLAN-0162-01. CANARY was supported by Agence Nationale de la Recherche (ANR) program 06-BLAN-0191, CNRS/INSU, Observatoire de Paris, and Université Paris Diderot Paris 7 in France, Science and Technology Facilities Council (Grants PP/E007651/1, ST/I002871/1, ST/K003569/1, ST/M007669/1, and ST/N002660/1),

the Durham University in the UK, and European Commission Seventh Framework Programme (E-ELT Preparation Infrastructure Grant 211257 and OPTICON Research Infrastructures Grant 226604). This project has received funding from the European Union's Horizon 2020 Research and Innovation Programme under grant agreement number 730890, and from ANR-10-LABX-0039-PALM. The William Herschel and Isaac Newton Telescopes are operated on the island of La Palma by the Isaac Newton Group of Telescopes in the Spanish Observatorio del Roque de los Muchachos of the Instituto de Astrofísica de Canarias. This work of JO was supported by the Science and Technology Facilities Council (UK) (ST/P000541/1), UKRI Future Leaders Fellowship (UK) (MR/S035338/1), and Horizon 2020.

DATA AVAILABILITY

The data underlying this paper will be shared on reasonable request to the corresponding author.

REFERENCES

- Agapito G., Battistelli G., Mari D., Selvi D., Tesi A., Tesi P., 2012, *Opt. Express*, 20, 27108
- Basden A., Geng D., Myers R., Younger E., 2010, *Appl. Opt.*, 49, 6354
- Conan R., Correia C., 2014, *Proc. SPIE*, 9148, 91486C
- Conan J.-M., Raynaud H.-F., Kulcsár C., Meimon S., 2011, Second International Conference on Adaptive Optics for Extremely Large Telescopes (AO4ELT2), Are integral controllers adapted to the new era of ELT adaptive optics?
- Correia C., Véran J.-P., Herriot G., 2012, *J. Opt. Soc. Am. A*, 29, 185
- Doelman N., Fraanje R., Breeje R., 2011, Second International Conference on Adaptive Optics for Extremely Large Telescopes (AO4ELT2), Real-sky adaptive optics experiments on optimal control of tip-tilt modes
- Gendron E., Léna P., 1994, *A&A*, 291, 337
- Guesalaga A., Neichel B., O'Neal J., Guzman D., 2013, *Opt. Express*, 21, 10676
- Haffert S. et al., 2020, Adaptive Optics Systems VII, Multi-core fibre-fed integral-field unit (MCIFU): Overview and first-light. *SPIE*, Bellingham
- Hinnen K., Verhaegen M., Doelman N., 2007, *J. Opt. Soc. Am. A*, 24, 1714
- Houtzager I., van Wingerden J., Verhaegen M., 2012, *IEEE Trans. Control Syst. Technol.*, 20, 934
- Juvénal R., Kulcsár C., Raynaud H.-F., Conan J.-M., Petit C., Lebouilleux L., Sivo G., Garrel V., 2015, Adaptive Optics for Extremely Large Telescopes IV (AO4ELT4), E64. Online at <http://escholarship.org/uc/ao4elt4>
- Juvénal R., Kulcsár C., Raynaud H.-F., Conan J.-M., 2018, *J. Opt. Soc. Am. A*, 35, 1465
- Kulcsár C., Raynaud H.-F., Petit C., Conan J.-M., 2012a, *Automatica*, 48, 1939
- Kulcsár C., Massioni P., Sivo G., Raynaud H.-F. G., 2012b, *Proc. SPIE*, 8447, 84470Z
- Kulcsár C., Raynaud H.-F., Juvénal R., Conan J.-M., 2018, Imaging and Applied Optics 2018, Optical Society of America, Washington, DC
- Le Roux B., Conan J.-M., Kulcsár C., Raynaud H.-F., Mugnier L. M., Fusco T., 2004, *J. Opt. Soc. Am. A*, 21, 1261
- Lozi J., Guyon O., Jovanovic N., Singh G., Goebel S., Norris B., Okita H., 2016, *Proc. SPIE*, 9909, 99090J
- Meimon S., Petit C., Fusco T., Kulcsár C., 2010, *J. Opt. Soc. Am. A*, 27, A122
- Myers R. M. et al., 2008, *Proc. SPIE*, 7015, 70150E
- Osborn J., 2018, Imaging and Applied Optics 2018. Optical Society of America, Washington, DC
- Overschee P. V., Moor B. D., 1994, *Automatica*, 30, 75
- Parisot A., Petit C., Fusco T., Conan J.-M., 2012, *Proc. SPIE*, 8447, 84471V
- Petit C., Conan J.-M., Kulcsár C., Raynaud H.-F., Fusco T., 2008, *Opt. Express*, 16, 87

- Petit C., Conan J.-M., Kulcsár C., Raynaud H.-F., 2009, *J. Opt. Soc. Am. A*, 26, 1307
- Petit C., Meimon S., Fusco T., Kulcsár C., Raynaud H.-F., 2010, 2010 IEEE International Conference on Control Applications. IEEE, Piscataway, NJ, p. 878
- Poyneer L. A. et al., 2016, *Appl. Opt.*, 55, 323
- Prenère L., Kulcsár C., Raynaud H.-F., 2020, *J. Opt. Soc. Am. A*, 37, 1083
- Rodier F., 1999, *Adaptive Optics in Astronomy*. Cambridge Univ. Press, Cambridge
- Rousset G., Fontanella J. C., Kern P., Gigan P., Rigaut F., 1990, *A&A*, 230, L29
- Rousset G. et al., 2003, in Wizinowich P. L., Bonaccini D., eds, Proc. SPIE Vol. 4839, Adaptive Optical System Technologies II. SPIE, Bellingham, p. 140
- Sivo G. et al., 2013, Proc. Third AO4ELT Conference, First on-sky validation of full LQG control with vibration mitigation on the CANARY MOAO pathfinder, INAF - Osservatorio Astrofisico di Arcetri, Firenze
- Sivo G. et al., 2014, *Opt. Express*, 22, 23565
- Tesch J., Truong T., Burruss R., Gibson S., 2015, *Opt. Lett.*, 40, 1575
- Yang K. et al., 2018, *Appl. Opt.*, 57, 2820

APPENDIX A: N4SID

Let the observability matrix associated with the pair (\tilde{A}_{LO}, C_{LO}) be written as

$$\mathcal{O}_s = \begin{bmatrix} C_{LO}^{\tilde{A}_{LO}} \\ C_{LO}^{\tilde{A}_{LO}^2} \\ \vdots \\ C_{LO}^{\tilde{A}_{LO}^{s-1}} \end{bmatrix}. \quad (\text{A1})$$

The integer s should be such that $\tilde{A}_{LO}^s \approx 0$ (Overschee & Moor 1994) and was set to 15 in this run.

Only a part from the N_t available samples will be used for the identification. The signal with the coefficients of the L first low-order modes is written here with z_k rather than ϕ_k^{LO} . Let $Z_{0,s,N}$ denotes a Hankel matrix with s block-rows and N columns,

$$Z_{0,s,N} = \begin{bmatrix} z_0 & z_1 & \cdots & z_{N-1} \\ z_1 & z_2 & \cdots & z_N \\ z_2 & \cdots & \cdots & \cdots \\ \vdots & \cdots & \cdots & \cdots \\ z_{s-1} & \cdots & \cdots & z_{N+s-2} \end{bmatrix}, \quad (\text{A2})$$

where N is such that $N + s - 2 < N_t$. The temporal sequence $Z_{0,1,N}$ is abbreviated with $Z_{0,N}$. From the RQ factorization,

$$\begin{bmatrix} Z_{0,s,N} \\ Z_{s,s,N} \end{bmatrix} = \begin{bmatrix} R_{11} & 0 \\ R_{21} & R_{22} \end{bmatrix} \begin{bmatrix} Q_1 \\ Q_2 \end{bmatrix}, \quad (\text{A3})$$

we have for $N \rightarrow \infty$,

$$\mathcal{O}_s X_{s,N} \approx R_{21} R_{22}^{-1} Z_{0,s,N} \quad (\text{A4})$$

and

$$\text{rank}(R_{21} R_{22}^{-1} Z_{0,s,N}) = n. \quad (\text{A5})$$

A singular value decomposition (SVD) of the matrix $R_{21} R_{22}^{-1} Y_{0,s,N}$ allows to express both the column space of \mathcal{O}_s and the row space of $X_{s,N}$ as

$$R_{21} R_{22}^{-1} Y_{0,s,N} = U_n \Sigma_n V_n^T, \quad (\text{A6})$$

where $U_n \in \mathbb{R}^{Ls \times n}$, $\Sigma_n \in \mathbb{R}^{n \times n}$, $V_n \in \mathbb{R}^{n \times N}$. The model order n is set equal to the first index of the singular values vector whose corresponding value is larger than 99 per cent of the accumulated sum of the singular values. It yields

$$\hat{\mathcal{O}}_s = U_n \Sigma_n^{1/2}, \quad (\text{A7})$$

$$\hat{X}_{s,N} = \Sigma_n^{1/2} V_n^T. \quad (\text{A8})$$

The matrices \tilde{A}_{LO} and C_{LO} are then estimated solving the least squares,

$$\min_{\tilde{A}_{LO}, C_{LO}} \left\| \begin{bmatrix} \hat{X}_{s+1,N} \\ Z_{s,N} \end{bmatrix} - \begin{bmatrix} \tilde{A}_{LO} \\ C_{LO} \end{bmatrix} \hat{X}_{s,N} \right\|_F^2, \quad (\text{A9})$$

and the covariance matrix Q^{LO} from its residual:

$$\begin{bmatrix} \hat{W}_{s,N} \\ \hat{V}_{s,N} \end{bmatrix} = \begin{bmatrix} \hat{X}_{s+1,N} \\ Z_{s,N} \end{bmatrix} - \begin{bmatrix} \tilde{A}_{LO} \\ C_{LO} \end{bmatrix} \hat{X}_{s,N}, \quad (\text{A10})$$

$$\begin{bmatrix} \hat{Q}^{LO} & \hat{S}^{LO} \\ \hat{S}^{LO} & \hat{R}^{LO} \end{bmatrix} = \lim_{N \rightarrow \infty} \frac{1}{N} \begin{bmatrix} \hat{W}_{s,N} \\ \hat{V}_{s,N} \end{bmatrix} \begin{bmatrix} \hat{W}_{s,N}^T & \hat{V}_{s,N}^T \end{bmatrix}. \quad (\text{A11})$$

The DARE associated with the set of matrices $(\tilde{A}_{LO}, C_{LO}, Q^{LO}, S^{LO}, R^{LO})$ is subsequently solved that yields the Kalman gain L_∞^{LO} in order to compute $A^{LO} = \tilde{A}_{LO} + L_\infty^{LO} C_{LO}$.

From the above steps, it can be seen that the computational complexity essentially comprises the cost of computing the RQ factorization, the SVD, solving the least squares (A9), and solving the DARE. Table A1 evaluates these costs more precisely, which reveals that the overall cost scales cubically with the number of modes.

Table A1. Computational complexity of N4SID.

RQ	$\mathcal{O}(N(sL)^2 + (sL)^3)$
SVD	$\mathcal{O}(N(sL)^2 + (sL)^3)$
Least squares (A9)	$\mathcal{O}(n(n+L)N)$
DARE	$\mathcal{O}(n^3)$

This paper has been typeset from a $\text{\TeX}/\text{\LaTeX}$ file prepared by the author.