Challenges in simulating advanced control methods for AO

Piscaer, Pieter; Soloviev, Oleg; Verhaegen, Michel

DOI

[10.1117/12.2577241](https://doi.org/10.1117/12.2577241)

Publication date

2020

Document Version

Final published version

Published in

Adaptive Optics Systems VII

Citation (APA)

Piscaer, P., Soloviev, O., & Verhaegen, M. (2020). Challenges in simulating advanced control methods for AO. In L. Schreiber, D. Schmidt, & E. Vernet (Eds.), *Adaptive Optics Systems VII* [114480O] (Proceedings of SPIE - The International Society for Optical Engineering; Vol. 11448). SPIE.
<https://doi.org/10.1117/12.2577241>

Important note

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

Copyright

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

Takedown policy

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

PROCEEDINGS OF SPIE

SPIDigitalLibrary.org/conference-proceedings-of-spie

Challenges in simulating advanced control methods for AO

Piscaer, Pieter, Soloviev, Oleg, Verhaegen, Michel

Pieter Piscaer, Oleg Soloviev, Michel Verhaegen, "Challenges in simulating advanced control methods for AO," Proc. SPIE 11448, Adaptive Optics Systems VII, 114480O (13 December 2020); doi: 10.1117/12.2577241

SPIE.

Event: SPIE Astronomical Telescopes + Instrumentation, 2020, Online Only

Challenges in simulating advanced control methods for AO

Pieter Piscaer^a, Oleg Soloviev^a, and Michel Verhaegen^a

^aDelft Center for Systems and Control, Delft University of Technology, Mekelweg 2, 2628 CD, The Netherlands

ABSTRACT

This paper discusses various practical problems arising in the design and simulation of predictive control methods for adaptive optics. Although there has been increased attention towards optimal prediction and control methods for AO systems, they are often tested in simplified simulation environments. The use of advanced AO simulators however, is a valuable alternative to the use of real data or laboratory experiments, as they provide both a flexible environment which is ideal for testing a new algorithm and are more accessible to non-experts. Topics that are often not explicitly discussed, such as the identification of a turbulence dynamics model from data, the use of matrix structures in AO systems to decrease the computational complexity and the implementation of Kalman filters to optimally deal with realistic noise conditions are examined. All topics discussed are illustrated by an accompanying Matlab code, which is based on the existing Matlab AO toolbox OOMAO.

Keywords: Turbulence modelling, System identification, Wavefront prediction, Optimal control, Adaptive Optics Simulation.

1. INTRODUCTION

Phase aberration in optical systems cause a deterioration of the image quality. Adaptive optics aims to compensate for the phase aberrations to obtain a sharper image. As reconstructing the phase aberrations from focal-plane images alone is challenging, a part of the light is directed towards a dedicated wavefront sensor (WFS). Using the WFS signal, a controller regulates the shape of a deformable mirror (DM) which aims to flatten the wavefront. In many cases, including astronomy, the phase aberrations are continuously evolving over time. Due to delays in obtaining the WFS image, computing the control action and applying this action to the DM, prediction of the wavefront is an important step in improving the image quality.

Several algorithms have been proposed that address the AO control problem using predictive optimal control algorithms. Most of these algorithms assume presence of a dedicated WFS to regulate the DM,¹⁻⁴ but other methods have been proposed that solve the AO problem without the use of WFSs.⁵⁻⁷ Often, the temporal evolution of the wavefront is either ignored or a dynamic model is obtained from first principles, but since the prior knowledge necessary to generate the model from first principles is generally not available, methods that obtain a model from data have gained attention.⁸⁻¹⁰

The computational complexity of data-driven predictive control problem forms a complication in the development of real-time algorithms for high-resolution AO. This has caused an increased attention in structured models for AO systems.⁹⁻¹² An additional complication in high-resolution AO is that, by increasing the resolution of the sensors, the relative brightness per observed pixel is decreased, which corresponds to a decrease of the signal to noise ratio (SNR). For the development of control methods that can function with low SNRs, not only a model of the aberration dynamics, but also an accurate measurement noise model is necessary.

This paper discusses challenges in the design of predictive and noise-robust AO algorithms, with a particular focus on the use of computationally efficient matrix structures and linear algebra operations to overcome computational complexity issues. These structured approaches do, however, need to make careful assumptions and approximations. As most of these methods are designed and validated in simulation studies, the exact simulation conditions can highly affect their accuracy. An accurate simulation test-bench, which

Further author information: (Send correspondence to P.P.)

P.P.: E-mail: P.J.Piscaer@tudelft.nl

can be used to objectively test different methods and their applicability under realistic conditions is therefore important in the further development of computationally efficient AO methods. Due to the steep learning curve of open-source accurate AO simulators, this paper proposes the development of a testbench that is user friendly for non-AO experts in order to bridge the gap between the optics and control communities. The theory discussed in this paper is implemented in an accompanying Matlab code, which can be downloaded from: <https://bitbucket.org/csi-dcsc/aotestbench>, based on the Matlab toolbox OOMAO.¹³

This paper is structured as follows. Section 2 discusses the conventional approach of modelling the AO system. The appearance of sparsity in the system matrices describing the turbulence dynamics is addressed in Section 3. Section 4 discusses the structures that can be used to efficiently model the sensor equations for both WFS-based and WFSless AO. The potential of the structured modelling approach is illustrated in Section 5, and challenges regarding the implementation and realistic simulation are addressed. Finally, Section 6 summarizes the main conclusions.

Notations: Vectors and matrices will be denoted by bold and capital symbols respectively, e.g. $\mathbf{x} \in \mathbb{R}^n$ and $X \in \mathbb{R}^{n \times n}$. The selection of rows with vectors ℓ_{row} and columns ℓ_{col} of a matrix X as $X^{(\ell_{row}, \ell_{col})}$. The pseudo-inverse of a matrix will be denoted by $X^\dagger$. Gaussian process $\mathbf{z}$ with mean μ and covariance matrix M will be denoted by $\mathbf{z} \sim \mathcal{N}(\mu, M)$.

2. MODELLING THE AO SYSTEM

Modelling the AO system consists of four important components: the deformable mirror (DM), the wavefront sensor (WFS), the focal-plane camera and the atmospheric turbulence.

2.1 The sensor and actuator models

When an imaging system observing a point source is exposed to a phase aberration, its influence on the image taken by a focal-plane camera is known as the point spread function (PSF), which can be formulated using the following equation:

$$\mathbf{y}(k) = |\mathcal{F}_{vec}(\exp(j\phi(k)))|^2 + \mathbf{v}_y(k), \quad (1)$$

where $\mathcal{F}_{vec}(\cdot)$ is the oversampled 2D-DFT and $\mathbf{v}_y(k)$ is the camera measurement noise. The aberration hence distorts the image in a non-linear sense, making it difficult to reconstruct the phase aberrations from the focal plane image alone. This problem is known as phase retrieval and has been widely studied in literature.¹⁴

Due to this highly non-linear relation, a WFS is often included to measure the phase aberration. The most common WFS is the Shack-Harmann (SH) sensor. The SH sensor signal, denoted by $\mathbf{s}(k)$, has the following linear relation with the wavefront:

$$\mathbf{s}(k) = G\phi(k) + \mathbf{v}_s(k), \quad (2)$$

where G is the SH *geometry matrix* and $\mathbf{v}_s(k)$ is the measurement noise, which is usually approximated by a zero-mean white Gaussian noise. The measurement matrix G is such that the signal $\mathbf{s}(k)$ represents local slopes of the the wavefront. Due to the linear relation between $\mathbf{s}(k)$ and $\phi(k)$, reconstructing the wavefront from the SH signal is easier than solving the phase retrieval problem.

The DM aims to compensate for the wavefront aberration. Since reading out the sensors and applying the control action to the DM takes time in practice, it is usually assumed that there is a delay in the system here, i.e.

$$\phi_m(k) = H\mathbf{u}(k-1), \quad (3)$$

where H is the DM *influence matrix* and $\mathbf{u}(k)$ is the control signal. In order to narrow the scope of this paper, the focus will be on wavefront reconstruction and prediction. Modelling the DM and the computation of the control signal will not be explicitly discussed.

Figure 1: Illustration of frozen flow's local influence. The top part displays the shift between consecutive time samples of a total distance $\mathbf{v}t_s$ per sampling time. Below, the phase aberration along the dashed line is plotted.

2.2 The turbulence dynamics model

The temporal dynamics are often neglected when considering the design of a controller, leading to an increased *temporal error* in AO systems. In order to decrease the temporal error, the sampling frequency could be increased, but this decreases the number of photons collected by the focal-plane and WFS each sampling time. An alternative way to decrease the temporal error is to include a model the temporal evolution of wavefront within the control algorithm, and compensate for the predicted wavefront instead.

A common way to represent the turbulence is by modelling it as a stack of *phase screens*, each moving in a different direction with a different speed. This assumption is known as Taylor's *frozen flow* hypothesis and can be formulated as follows:

$$\phi(\boldsymbol{\zeta}, k+1) = \phi(\boldsymbol{\zeta} - \mathbf{v}t_s, k), \quad (4)$$

where $\boldsymbol{\zeta} \in \mathbb{R}^2$ denotes the spatial coordinates, t_s a sampling time and $\mathbf{v} \in \mathbb{R}^2$ the direction and speed of translation of the phase screen. In Fig. 1, the frozen flow a single phase screen is illustrated. Based on the frozen flow assumption, a common turbulence model in literature is the Vector Auto-Regressive (VAR) model:

$$\boldsymbol{\phi}(k+1) = \sum_{i=0}^{q-1} A_{i+1} \boldsymbol{\phi}(k-i) + \mathbf{w}(k), \quad (5)$$

where $\boldsymbol{\phi} \in \mathbb{R}^n$ is the vectorized wavefront, A_{i+1} are the VAR coefficient matrices and $\mathbf{w}(k)$ is a Gaussian white noise signal with covariance matrix Q , denoted by $\mathbf{w}(k) \sim \mathcal{N}(0, Q)$.

3. SPARSE MODELLING OF THE TURBULENCE DYNAMICS

Especially for large-scale AO systems, efficient modelling of the turbulence and AO components is crucial for real-time applications. This section will focus on sparsity structures arising in VAR modelling of the turbulence dynamics.

3.1 VAR-1 modelling of frozen flow turbulence

The 2-dimensional sampled wavefront can be represented as a matrix $\Phi \in \mathbb{R}^{\bar{n} \times \bar{n}}$. Assume now the case in which $\mathbf{v}$ and t_s are such that, between two consecutive measurements, the wavefront moves exactly one sampling distance from left to right over the aperture, i.e.: $\Phi^{(i,j)}(k+1) = \Phi^{(i,j-1)}(k)$. In this case, predicting the value

of the phase screen at spatial-temporal coordinates $(i, j, k + 1)$ simply equals the phase screen at $(i, j - 1, k)$. The *spatio-temporal* correlations can be modelled by a simple VAR model of order 1 (VAR-1). When using the vectorized wavefront representation, this becomes:

$$\phi(k + 1) = A_s \phi(k), \quad (6)$$

where $A_s^{(i,j)} = 1$ if $j = i + 1$ and zero otherwise. This *shift matrix* A_s forms the intuitive basis when modelling frozen-flow phase screens using VAR models.

Shift-structured VAR-1 modelling does, however, have its limitations that can lead to significant modelling errors. Due to the finite aperture size, the information $(i, j - 1)$ might no longer be available as it can fall outside the aperture. Consequently, these terms have to be predicted via some kind of extrapolation. Another significant limitation of the VAR-1 model appears when considering a stack of multiple phase screens. In this case, the shift-based model can no longer be found. An even more fundamental limitation is the inaccuracy of interpolation when the wind speed per sample time do not correspond to integer values. In this case, there is no exact shift matrix available and the entries in the matrix A_s should resemble an interpolation matrix instead.

3.2 Banded sparsity of VAR coefficient matrices

The previous ideas for modelling a unidirectional shift can be extended for the modelling of local movements within a certain ‘close neighbourhoods’. Even when considering multiple phase screens and finite size apertures, the prediction of the temporal evolution of a single pixel only requires the data of surrounding pixels within its ‘close neighbourhood’. This reasoning and its relation with sparse VAR models can be formalized using *graphical modelling* theory.¹⁵ In this more flexible case, the wind direction is assumed to be unknown, so a circular ‘close neighbourhood’ will be considered. It is also assumed that the true wind speed of each layer is unavailable, but that an overestimate of the *average* wind speed over all layers, $\bar{v}$, can be estimated.

For the definition of the ‘close neighbourhood’, the shift in terms of pixel widths δ per sampling time corresponding to the average wind speed is of great importance, that is:

$$\nu = |\bar{v}|t_s/\delta. \quad (7)$$

For example, $\nu = 2$ indicates that the average wind speed is two pixels per sampling time. With ν estimated, the local ‘close neighbourhood’ is defined as all pixels within a circular area with radius $r_{A_1} = \nu$ pixels (or $\nu\delta$ meter) around the pixel of interest. For general VAR- q models, the extra time delays corresponding to the other coefficient matrix A_i , $i = 2, \dots, q$ will need increasing radii. In general, it makes sense to take $r_{A_i} = i r_{A_1}$ for $i = 2, \dots, q$, and hence $r_{A_i} = i\nu$ as a rule-of-thumb. This shows that, although higher order VAR models might lead to more accurate predictions, the coefficient matrices gradually get less sparse. The sparsity patterns of the matrices A_i will be defined by the sets $\mathcal{S}_{A_i}$, which correspond to it areas with radius r_{A_i} defining the close neighbourhoods. An example is visualized in Figure 2.

Although determining the values of the non-zero entries in A_i from first principles or manual tuning is a possibility, this paper will focus a more systematical approach in which the the matrices A_i are identified from data. The identification of sparse VAR models from data can be done by solving the following constrained least-squares problem:

$$\begin{aligned} \min_{A_i} \|\Phi_f - [A_1 \quad \dots \quad A_q] \Phi_p\|_F^2 \\ \text{s.t. } A_i \in \mathcal{S}_{A_i}, \quad i = 1, \dots, q \end{aligned} \quad (8)$$

with Φ_f and Φ_p are defined as:

$$\Phi_f = [\phi(q+1) \quad \phi(q+2) \quad \dots \quad \phi(N)], \quad \Phi_p = \begin{bmatrix} \phi(q) & \phi(q+1) & \dots & \phi(N-1) \\ \phi(q-1) & \phi(q) & \dots & \phi(N-2) \\ \vdots & \vdots & \ddots & \vdots \\ \phi(1) & \phi(2) & \dots & \phi(N-q) \end{bmatrix}.$$

Figure 2: Visualization of (a) the circular ‘close neighbourhood’ with radius $r_{A_1} = 2$ and (b) the corresponding sparsity pattern of A_1 for a circular aperture with $\bar{n} = 9$. $\times$ symbols represent the sampling locations (pixels). The red $\times$ is the pixel to be estimated, the blue $\times$ are the pixels considered to be in the ‘close neighbourhood’.

To simplify notations, the algorithm to efficiently solve (8) is presented for a VAR-1 model. Denote the non-zero entries in the i -th row of the matrix A_1 by ℓ_i . The optimization problem can then be separated for each row of A_1 into a set small-scale least-squares problems of the form:

$$\hat{A}_1^{(i, \ell_i)} = \arg \min_{A^{(i, \ell_i)} \in \mathbb{R}^{\kappa_i}} \|\Phi_f^{(i, \star)} - A_1^{(i, \ell_i)} \Phi_p^{(\ell_i, \star)}\|_2^2 \quad (9)$$

where the notation $\star$ denotes the selection of all rows or columns. Assuming the case of $r_A = 2$ as in Fig. 2b, the maximum number of non-zero values per row is 13, such that the identification of an n -by- n matrix as in (8) can be broken down as the estimation of a set of n independent vectors of length ≤ 13 . Since $n \gg m_A$, this new series of smaller least-squares problems is significantly faster than considering a dense A_1 and due to the lower number of unknowns, the optimization problems will require much less data (i.e. smaller N). The code for this algorithm is implemented in the function `Sparse_ARid.m`. To prevent a severe overestimate of the number of non-zero elements each vector $A_1^{(i, \ell_i)}$, sparsity inducing ℓ_1 -regularization could be added. As a circle is generally not the most optimal shape to describe the region of influence, this regularization will be able to find a more ‘optimal’ shape of this region. An ADMM implementation to solve this regularized least-squares problem is included in the function `Sparse_ARid.m`.

3.3 Modelling the atmospheric turbulence stochasticity

With the matrix A identified, the next task is to obtain an accurate estimation of the process noise covariance Q . A straightforward approximation of the matrix Q would be to use the residual of the optimization problem 8, $E = \Phi_f - A\Phi_p$, and construct a *sample covariance matrix* of Q , $Q_s := (EE^T)/(N - 1 - q)$, which requires $N \gg n^2$ in order to obtain an accurate representation. However, the storage required to perform this matrix-matrix multiplication itself can already lead to problems when n and N are very large.

Following the same graphical modelling approach that explains the sparsity of the AR coefficients, it is expected that the inverse of the covariance matrix, Q^{-1} , is highly sparse.¹⁵ This same phenomenon is also visible when looking at the covariance matrix of the wavefront $C_\phi = E[\phi(k)\phi^T(k)]$. To demonstrate this Fig. 3 displays the values of a straightforward inversion of the sample covariance matrices $C_{\phi, s}$ and Q_s . Unfortunately, this inversion of the sample covariance matrix only leads to an approximately sparse matrix (i.e. many are small, but not exactly zero) and truncation of the small values might lead to errors. Therefore, it is better to enforce sparsity within the identification procedure.

The problem of computing Q_s^{-1} could be redefined as follows: find a square invertible matrix $W \in \mathbb{R}^{n \times n}$ such that $\frac{1}{N-1} EE^T W^T W = I$, which is equivalent to:

$$WEE^T W^T = (N - 1)I. \quad (10)$$

Figure 3: Sparsity patterns arising in the inverse matrices (a) C_ϕ^{-1} and (b) Q^{-1} . Matrices shown are inverted sample covariance matrices using $N = 20,000$ datapoints.

In other words, find W such that $\frac{1}{\sqrt{N-1}}WE$ is an orthonormal matrix. Finding such an orthonormal matrix can be found using the Gram-Schmidt method. Since the Gram-Schmidt method will find a lower-triangular matrix W , the desired sparsity pattern of Q^{-1} can easily be obtained by enforcing the lower-triangular part of the sparsity pattern of Q^{-1} on W . The Matlab code of this method is given by the function `ComputeW.m` and the total identification of the matrices A_i and Q^{-1} is included in `ModelIdentification.m`.

3.4 Acquiring identification data via wavefront reconstruction

A topic that is often overlooked is the process used to obtain accurate identification data. Although any focal-plane or pupil-plane wavefront sensing technique can be used to retrieve the identification data, this section will provide an example using a SH sensor. With an available SH sensor, the wavefront reconstruction problem can be formulated as a simple matrix-vector product of the form:

$$\hat{\phi}(k) = Ms(k), \quad (11)$$

where M is a *reconstruction matrix*. Using the description of the WFS as in (2), it can be found that $M = G^\dagger$, i.e. the pseudo-inverse of the matrix G . Since G is rank deficient, the computation of this pseudo inverse requires either the use of the truncated SVD or additional regularization. For a minimum variance wavefront reconstruction, $M = (G^T G + \sigma^2 C_\phi^{-1})^{-1} G^T$. Computing M can be complicated in practice for very large-scale systems. Section 4 will discuss a different approach to wavefront reconstruction, which is able to exploit the sparsity in the system matrices. Of course, the accuracy of the wavefront reconstruction is depending on the measurement noise conditions. Since the measurement conditions are in turn affected by the source brightness and the resolution of the WFS, one should be extra careful when modelling the dynamics for high-resolution wavefronts or from fainter sources.

3.5 Accuracy of sparse turbulence models

The accuracy of a simple VAR model is first tested for a single phase screen. Fig. 4 shows the performance of different VAR models versus the wind speed of the phase aberration generated by a single phase screen. Although there is a general upwards trend in MSE with respect to the wind speed, there is dip in MSE around a wind speed moving at $\nu = 1$. Since this corresponds to a shift that coincides with exactly one sampling distance δ per sampling time, this can be more accurately captured than a shift of approximately $\nu = 0.5$. Increasing the AR order can increase the performance for windspeeds around $\nu = 0.5$.

Next, similar tests are performed under slightly more realistic circumstances. The turbulence is modelled by 3 separate phase screens moving at different speeds and directions. The wavefront data used for identification (i.e. Φ_p and Φ_f) are constructed using reconstructed wavefront data using the simple wavefront reconstruction

Figure 4: Accuracy of identified VAR models in simulation. The wind is moving exactly parallel to the grid of pixels. The phase screen is simulated on a 30×30 grid of pixels using OOMAO with parameters: $D = 0.75m$, $r_0 = 0.15m$, $L_0 = 25m$. (a) Results for varying r_{A_1} for a VAR-1 model. (b) Results for varying VAR orders, $r_{A_1} = 3$ and $r_{A_i} = i r_{A_1}$.

Figure 5: Accuracy of identified VAR models in simulation for three phase screens. The wind speed ν corresponds to the ground layer wind speed. With respect to layer 1, layers 2 and 3 move at 1.5 times the speed at an angle of 90 and 45 degrees respectively. The phase screen is simulated with parameters: $D = 1m$, $r_0 = 0.15m$, $L_0 = 25m$. The identification data is reconstructed using measurements from a WFS with 15×15 lenslets using $M = G^\dagger$ in (11) observing a point source with (Vega) magnitude 8 with a sampling frequency of $f_s = 500Hz$. (a) Results for varying r_{A_1} for a VAR-2 model. (b) Results for varying VAR orders, $r_{A_1} = 3$ and $r_{A_i} = i r_{A_1}$.

with $M = G^\dagger$ in (11). Figure 5 shows the performance of the VAR models under these conditions. What is interesting to see is that in this multi-layered turbulence case, the VAR-2 and VAR-3 models show similar oscillating behaviour as the single layer VAR-1 case and that increasing the VAR order does not always lead to better results in practice, making the selection of the VAR order not as simple as trade-off between accuracy and computational complexity.

4. COMPUTATIONAL EFFICIENT MODELLING OF THE SENSORS

Besides the turbulence models, the AO components have to be modelled in an efficient manner for high-resolution AO applications. This section will discuss computationally efficient models for the SH wavefront sensor and focal-plane camera.

4.1 Pupil-plane wavefront sensing

Recall that the Shack-Hartmann WFS was modelled using (2). Assuming Fried geometry, the matrix G only has 4 non-zero values per row, making it highly sparse for high-resolution WFSs. Since the total measurement noise is approximated by a zero-mean Gaussian white noise signal with covariance matrix $R = \sigma^2 I$, R is per definition very structured. Besides the sparsity in G , there is another interesting structure present in SH sensors assuming Fried geometry. When only considering the slopes in x - and y -direction separately and assuming a square grid of lenslets, the following Kronecker structure appears:

$$\mathbf{s}(k) = \begin{bmatrix} \mathbf{s}_x(k) \\ \mathbf{s}_y(k) \end{bmatrix} = \begin{bmatrix} G_1 \otimes G_2 \\ G_2 \otimes G_1 \end{bmatrix} \boldsymbol{\phi}(k), \quad (12)$$

with G_1 and G_2 sparse banded matrices. Unfortunately, it is not straightforward to use this structure for efficient wavefront reconstruction. Although the pseudo-inverse of a matrix $(G_1 \otimes G_2)$ can be efficiently computed via $G_1^\dagger \otimes G_2^\dagger$, pseudo-inverting the total matrix G would involve computing the inverse of the following sum of Kronecker products $(G_1^T G_1 \otimes G_2^T G_2) + (G_2^T G_2 \otimes G_1^T G_1)$, for which no such efficient closed-form solutions exist.

Since the slopes $\mathbf{s}$ can be seen as the local gradient of the wavefront $\boldsymbol{\phi}$, wavefront reconstruction essentially becomes the integration of the slopes $\mathbf{s}$. Based on this intuitive reasoning, the undesired dense nature of the reconstruction matrix M in (11) becomes apparent. Where for reconstructing modes with higher spatial-frequency a local integration will suffice to capture their shape, this is not the case for lower spatial-frequency modes. In order to capture these low-spatial frequency modes accurately, the matrix M cannot be simply enforced to be sparse and different matrix structures have to be explored. For example, by reformulating the minimum variance wavefront reconstruction problem as the problem of solving the following normal equations:

$$\left(G_1^T G_1 \otimes G_2^T G_2 + G_2^T G_2 \otimes G_1^T G_1 + \sigma^2 C_\phi^{-1} \right) \hat{\boldsymbol{\phi}}(k) = G^T \mathbf{s}(k). \quad (13)$$

The highly structured and sparse matrix on the left hand side can be exploited by computationally efficient solvers such as the conjugate gradient method.¹⁶

4.2 Focal-plane sensing

For the WFSless case, the output equation (1) does not contain such clear sparsity structures compared to the matrix G . However, the 2D-DFT defining the relation between the PSF and wavefront has itself another special structure that can be used in efficient computations. As a matrix vector product, the 2D-DFT can be written as:

$$\mathcal{F}_{vec}(\mathbf{x}) = (D \otimes D) \mathbf{x} = \text{vec}(DXD^T), \quad (14)$$

where D is the DFT matrix and $\mathbf{x} \in \mathbb{R}^n$ the vectorized generalized pupil function (GPF) X . For the last equation, the relation $\text{vec}(AXB) = (B^H \otimes A)\text{vec}(X) \in \mathbb{R}^{\tilde{n} \times \tilde{n}}$ is used. The DFT itself has more properties, which are exploited by the efficient *fast Fourier transform* (FFT) algorithms. Due to this structured representation, focal-plane wavefront sensing can be performed very efficiently using iterative methods. For example, alternating projection methods exploit the speed of the DFT and inverse DFT to solve the phase retrieval problem. Also in the case of dynamic aberrations, the structure in the DFT can be exploited in the development of optimal wavefront prediction algorithms, such an example will be given in Section 5.

4.3 Modelling the camera noise

Although the measurement noise of the SH sensor slope signal $\mathbf{s}$ can be accurately modelled as a zero-mean white Gaussian noise, this is not true for camera noise in general. There are multiple sources of measurement noise in CCD cameras. The most notable ones being the *shot noise*, which has a Poissonian distribution, and *read-out noise*, which is usually assumed to be Gaussian. Considering a single pixel value of a CCD camera measurement y , this combination of measurement noises can be modelled as follows:

$$y = y_{true} + v_{shot} + v_{read}, \quad (15)$$

Figure 6: KL divergence for different noise approximations. The x-axis indicates the value of y_{true} . The Gaussian and Poissonian approximations are $\mathcal{N}(y, y)$ and $Pois(y)$ respectively, with y drawn from a random distribution $Pois(y_{true})$.

where y_{true} represent the number of photons in a hypothetical noiseless case. The read-out noise is a zero-mean white Gaussian noise $v_{read} \sim \mathcal{N}(0, \sigma_r^2)$ and the contribution of the Poissonian shot noise, v_{shot} , is such that $(y_{true} + v_{shot}) \sim Pois(y_{true})$.

Since many existing filtering and prediction techniques are based on the assumption that the measurement noise is Gaussian, a Gaussian approximation can be interesting. The Poisson distribution is known to converge to a Gaussian distribution when the number of true arriving photons, y_{true} , is large:

$$Pois(y_{true}) \approx \mathcal{N}(y_{true}, y_{true}) \quad (16)$$

and since the read-out noise and shot noise can be considered to be uncorrelated, $v_{shot} + v_{read} \sim \mathcal{N}(0, y_{true} + \sigma_r^2)$.

Besides the errors caused by the Gaussian approximation, another fundamental source of modelling errors is that y_{true} is unknown in general. Therefore, an estimate of y_{true} , denoted by y_{est} , has to be used. Since only a single measurement y of the pixel is observed, the best available estimate is $y_{est} = y$. Using this approximation, the Gaussian approximation for the full camera measurement vector becomes:

$$\mathbf{v}_y(k) \sim \mathcal{N}(0, R_y(k)), \quad \text{with } R_y(k) = \sigma_r^2 I + \text{diag}(\mathbf{y}(k)) \quad (17)$$

and the SNR of a pixel with value y can be approximated by:

$$SNR(y) \approx \frac{y}{\sqrt{y + \sigma_r^2}}. \quad (18)$$

As the SNR is very low for dark pixels, many algorithms that directly use CCD camera measurements, such as phase retrieval algorithms, employ a value truncation of the data.

Since y_{true} is unknown no matter which statistical distribution is chosen, even Poissonian noise models will not be exact. A simple simulation is performed to compare the accuracy of the Gaussian and Poissonian approximations of the shot noise. The Kullback-Leibler (KL) divergence can be used as a measure of the accuracy of certain estimated distribution with respect to the true distribution. Fig. 6 shows the KL divergence over a Monte Carlo simulation with 30000 independent draws. It can be concluded that both the Poissonian and Gaussian model have significant modelling errors, and that the Gaussian approximation only causes a very slight increase in KL divergence when $y_{true} \lesssim 25$ photons.

5. IMPLEMENTATION AND SIMULATION OF ADVANCED CONTROL METHODS FOR AO

This section will address challenges regarding the implementation of predictive wavefront reconstruction methods using the efficient turbulence and sensor models discussed in the previous sections.

5.1 Potential of structured modelling in AO

To illustrate both the potential of structured modelling as well as their implementation challenges, this subsection will provide two examples of wavefront estimation methods for high-resolution AO systems.

5.1.1 Example 1: non-linear Kalman filtering for WFSless AO

This subsection discusses an example of an efficient non-linear Kalman filter (KF) implementation for focal-plane wavefront sensing.⁷ In this method, it is assumed that the phase is small so that a first-order Taylor approximation can be used to linearize the output equation. For the relation in (1), the first order Taylor approximation is:

$$\mathbf{y}(k) = |\mathcal{F}_{vec}(\mathbf{a} \odot \exp(j\epsilon(k)))|^2 + \mathbf{v}_y(k) \approx J_{\hat{\phi}}(\phi(k) - \hat{\phi}) + \mathbf{c}_{\hat{\phi}} + \mathbf{v}_y(k), \quad (19)$$

where $f(\phi) = |\mathcal{F}_{vec}(\mathbf{a} \odot \exp(i\phi))|^2$, $\mathbf{c}_{\hat{\phi}} = f(\hat{\phi})$ and $J_{\hat{\phi}} = f'(\hat{\phi})$, i.e. the Jacobian matrix corresponding $f(\phi)$ evaluated at $\hat{\phi}$. The explicit computation of the Jacobian matrix has a large computational complexity, but the matrix-vector multiplication $J_{\hat{\phi}}\phi$ can be solved efficiently using the FFT. When introducing $\hat{\mathbf{x}} = \mathbf{a} \odot \exp(i\hat{\phi})$ and $\hat{\mathbf{y}} = \mathcal{F}_{vec}(\hat{\mathbf{x}})$, the following operator can be defined:

$$J_{\hat{\phi}}\phi \rightarrow \mathcal{J}_{\hat{\phi}}(\phi) \triangleq \Re(2i\hat{\mathbf{y}} \odot \mathcal{F}_{vec}\{\hat{\mathbf{x}} \odot \phi\}), \quad (20)$$

where $\mathcal{F}_{vec}^{-1}\{\cdot\}$ represents the vectorized inverse 2D-DFT and $\hat{\mathbf{y}}^*$ and $\hat{\mathbf{x}}^*$ represent the element-wise complex conjugates of $\hat{\mathbf{y}}$ and $\hat{\mathbf{x}}$ respectively. This efficient formulation of the Jacobian can be used in non-linear KFs such as the EKF. Using the same reasoning as for the structured wavefront reconstruction problem in Sec. 4, the EKF measurement update can be formulated as a solution to a system of normal equations that are highly structured:

$$\left(J_{\hat{\phi}}^T R_y^{-1} J_{\hat{\phi}} + P^{-1}\right) \delta\phi = J_{\hat{\phi}}^T R_y^{-1} \left(\mathbf{y}(k) - \mathbf{c}_{\hat{\phi}}\right), \quad (21)$$

such that $\hat{\phi}(k) = \delta\phi + \hat{\phi}$ is the KF measurement update, and where $\hat{\phi}$ and P^{-1} is prior information based on the identified turbulence model, see the original paper for more details.⁷ The code for this EKF algorithm is included within `IEKF_Prediction.m`.

5.1.2 Example 2: basis functions in high-resolution phase retrieval

For the second example, the use of a modal representation of the wavefront using basis functions will be discussed. When choosing a proper basis function in WFS-less AO, one should aim to preserve as many structures as possible that are already present in the 2-dimensional DFT. The DFT matrix is highly structured, which is exploited by FFT algorithms. On top of the structures within the DFT matrix, the 2D-DFT is a separable function, which was given by (14). A regular square grid of Gaussian radial basis functions is one of such sets of basis functions that maintains the separability of the 2D-DFT. Expressing the GPF in terms of these GRBFs gives the following relation between the basis function coefficients and the corresponding PSF:

$$\mathbf{y} = |(C_1 \otimes C_2) \boldsymbol{\alpha}|^2, \quad (22)$$

where $\boldsymbol{\alpha} \in \mathbb{C}^{n_{\alpha}}$ is a vector containing coefficients. Due to the use of a modal representation, the structures present in the DFT matrix D itself will be lost, making the evaluation of the PSF in terms of basis functions generally slower than when using a pixel basis. Despite the increased theoretical computational complexity, there are still advantages of using a modal basis. For example, by limiting the number of basis functions, the number of unknowns can be significantly decreased. The use of smooth basis functions in phase retrieval furthermore results into a higher noise robustness. So even when C_1 and C_2 do not have structures that are exploited, the decrease in the number of parameter due to the modal representation, and the use of the separability can be used to speed up phase retrieval algorithms. One such possible Modal-based alternating projection algorithms is implemented in `ModalAP_Prediction.m`.

5.2 A simulation test-bench for AO

Although matrix structures such as sparsity and tensor structures tend to address the typical challenges in large-scale system modelling, their enforced approximations might require simplifications that result in unrealistic physical circumstances. It was shown that realistic environmental conditions such as finite and circular aperture shapes, non-Gaussian measurement noise and a large number of turbulent phase screens, might lead to modelling errors. Due to the complicated nature of the control algorithms and their dependence on the modelling accuracy and environmental conditions, many algorithms are presented in a simple simulation environment. Unfortunately, using over-simplified simulation environment might provide an inaccurate picture of the new method's performance.

An important step to bridge the gap between astronomy and the control community would be the existence of a realistic, user-friendly simulation test-bench. Due to the lack of such a test-bench, issues regarding unrealistic assumptions regarding the complex temporal evolution the turbulence, its stochastic nature and measurement noise conditions are often overlooked. Open-source AO toolboxes such as OOMAO form an important element in the development as they are widely used and tested by the astronomy community. The steep learning curve and required background knowledge to operate such toolboxes properly are, however, prohibiting them to be used by researchers without a background in astronomy. This paper proposes a user-friendly application of exiting AO toolboxes such as OOMAO to develop a realistic simulation test-bench while opening up the complicated control problems faced in AO to a wider group of researchers in the control community. Matlab code of the main ideas discussed in this paper, and a number of other functions regarding structured modelling, have been included in the supplementary code, which serves as an idea for what functionalities such a test-bench should include.

6. CONCLUSIONS

Wavefront aberrations are the result of complex physical processes that can be difficult to model accurately in an efficient manner. It was discussed how enforcing matrix structures and exploiting efficient linear algebra operations is a promising approach to identify an accurate model for the turbulence dynamics from a limited dataset. Simple model structures, such as low-order sparse VAR models, were shown to be a good candidate for data-driven modelling of the turbulence using a limited dataset. While more complicated models might be more accurate in theory, their large number of parameters to be estimated will eventually become a problem by itself. Another reason for considering such structured models is that the number of inputs and outputs can become very large when large-scale or high-resolution AO applications are considered, making it a challenge to develop real-time algorithms. Most existing optimal predictive control algorithms are often impractical for such large-scale systems. By exploiting the efficient linear algebra operations, however, these control algorithms can be implemented much more efficiently. Despite their great potential, the underlying assumptions and approximations of these algorithms make them sensitive to the simulation conditions in which they are tested. Therefore, a clear and user-friendly simulation test-bench should be developed in order to properly validate the future generation of advanced AO algorithms.

REFERENCES

- [1] Kulcsár, C., Raynaud, H.-F., Petit, C., Conan, J.-M., and De Lesegno, P. V., "Optimal control, observers and integrators in adaptive optics," *Optics express* **14**(17), 7464–7476 (2006).
- [2] Kulcsár, C., Raynaud, H.-F., Petit, C., and Conan, J.-M., "Minimum variance prediction and control for adaptive optics," *automatica* **48**(9), 1939–1954 (2012).
- [3] Correia, C., Conan, J.-M., Kulcsár, C., Raynaud, H.-F., and Petit, C., "Adapting optimal lqg methods to elt-sized ao systems," in *[1st AO4ELT conference-Adaptive Optics for Extremely Large Telescopes]*, 07003, EDP Sciences (2010).
- [4] Massioni, P., Raynaud, H.-F., Kulcsár, C., and Conan, J.-M., "An approximation of the riccati equation in large-scale systems with application to adaptive optics," *IEEE transactions on control systems technology* **23**(2), 479–487 (2014).
- [5] Piscaer, P., Soloviev, O., and Verhaegen, M., "Predictive wavefront sensorless adaptive optics for time-varying aberrations," *JOSA A* **36**(11), 1810–1819 (2019).

- [6] Smith, C., Marinica, R., and Verhaegen, M., “Real-time wavefront reconstruction from intensity measurements,” in [*Proceedings of the 3rd AO4ELT Conference: Adaptive Optics for Extremely Large Telescopes, Florence, Italy, 26-31 May 2013*], Arcetri Astrophysical Observatory (2013).
- [7] Piscaer, P., Soloviev, O., and Verhaegen, M., “Phase retrieval of large-scale time-varying aberrations using a non-linear kalman,” *JOSA A* (To be published) (2020).
- [8] Hinnen, K., Verhaegen, M., and Doelman, N., “A data-driven $\mathcal{H}_2$ -optimal control approach for adaptive optics,” *IEEE Transactions on Control Systems Technology* **16**(3), 381–395 (2008).
- [9] Piscaer, P., *Sparse VARX Model Identification for Large-Scale Adaptive Optics*, Master’s thesis, TU Delft (2016). Retrieved from <http://resolver.tudelft.nl/uuid:9677ebd9-cf4c-484c-97f8-1f553cd0e873>.
- [10] Siquin, B. and Verhaegen, M., “Tensor-based predictive control for extremely large-scale single conjugate adaptive optics,” *JOSA A* **35**(9), 1612–1626 (2018).
- [11] Yu, C. and Verhaegen, M., “Structured modeling and control of adaptive optics systems,” *IEEE Transactions on Control Systems Technology* **26**(2), 664–674 (2018).
- [12] Fraanje, R., Rice, J., Verhaegen, M., and Doelman, N., “Fast reconstruction and prediction of frozen flow turbulence based on structured kalman filtering,” *JOSA A* **27**(11), A235–A245 (2010).
- [13] Conan, R. and Correia, C., “Object-oriented matlab adaptive optics toolbox,” in [*Adaptive Optics Systems IV*], **9148**, 91486C, International Society for Optics and Photonics (2014).
- [14] Shechtman, Y., Eldar, Y. C., Cohen, O., Chapman, H. N., Miao, J., and Segev, M., “Phase retrieval with application to optical imaging: a contemporary overview,” *IEEE signal processing magazine* **32**(3), 87–109 (2015).
- [15] Songsiri, J., Dahl, J., and Vandenberghe, L., “Graphical models of autoregressive processes,” *Convex Optimization in Signal Processing and Communications* (2009).
- [16] Saad, Y., [*Iterative methods for sparse linear systems*], SIAM (2003).