Phase retrieval of large-scale time-varying aberrations using a non-linear Kalman filtering framework

Piscaer, Pieter; Soloviev, Oleg; Verhaegen, Michel

DOI
10.1364/JOSAA.405712

Publication date
2021

Document Version
Final published version

Published in
Journal of the Optical Society of America A: Optics and Image Science, and Vision

Research Article Vol. 38, No. 1 / January 2021 / Journal of the Optical Society of America A 25

Journal of the Optical Society of America A

OPTICS, IMAGE SCIENCE, AND VISION

Phase retrieval of large-scale time-varying aberrations using a non-linear Kalman filtering framework

PIETER PISCAER,* OLEG SOLOVIEV, AND MICHEL VERHAEGEN

Delft Center for Systems and Control, Delft University of Technology, Mekelweg 2, 2628 CD, Delft, The Netherlands
**Corresponding author: P.J.Piscaer@tudelft.nl*

Received 19 August 2020; revised 22 October 2020; accepted 16 November 2020; posted 17 November 2020 (Doc. ID 405712); published 8 December 2020

This paper presents a computationally efficient framework in which a single focal-plane image is used to obtain a high-resolution reconstruction of dynamic aberrations. Assuming small-phase aberrations, a non-linear Kalman filter implementation is developed whose computational complexity scales close to linearly with the number of pixels of the focal-plane camera. The performance of the method is tested in a simulation of an adaptive optics system, where the small-phase assumption is enforced by considering a closed-loop system that uses a low-resolution wavefront sensor to control a deformable mirror. The results confirm the computational efficiency of the algorithm and show a large robustness against noise and model uncertainties.

https://doi.org/10.1364/JOSAA.405712

1. INTRODUCTION

Adaptive optics (AO) systems correct in real time for phase aberrations in optical systems. A deformable mirror (DM) is used to correct for the phase aberrations and counter their effects on image quality. Often, the light is split between a wavefront sensor (WFS), which measures the phase aberration, and a focal-plane camera, which captures the image of interest. The reason for including a dedicated WFS is that retrieving the aberrations from focal-plane images alone is very challenging, as most methods are based on iterative algorithms, which are difficult to apply in real time for temporally dynamic wavefronts. Despite the challenges, focal-plane wavefront sensing is still an active field, as it simplifies the optical setup and eliminates errors due to non-common-path aberrations (NCPAs). Assuming a point source, the focal-plane camera captures the point spread function (PSF). The problem of retrieving the static phase aberrations from the PSF is known as the phase retrieval (PR) problem.

Many PR algorithms have been developed, and these can be split into three main classes. The first and still the most commonly used class is based on alternating projections (AP). The most well-known AP methods include [1,2], but many more have been developed. See [3] for an overview and [4] for a recent state of the art method. A second class of methods reformulates the PR problem as a convex optimization problem [5,6]. Despite promising convergence results, these methods are very computationally demanding for larger AO systems. More recently, a third class, which solves the PR problem in its original non-convex framework, has gained attention [7–10]. A problem of this class of methods is the need of an accurate initial guess of the wavefront, usually computed via techniques that require a restrictive measurement setup and many more measurements than unknowns [7,8].

Another approach to focal-plane wavefront sensing is to assume that the phase to be estimated is small [11] and to apply a linear or quadratic approximation of the PSF. This approach includes methods such as iterative linear phase diversity (ILPD) [12], linearized focal-plane technique (LIFT) [13], and Fast & Furious [14]. These methods are restricted to small wavefronts (RMS value below 0.5 rad [12]). In addition, in [15], the approach was used to estimate dynamic aberrations using a non-linear Kalman filtering (KF) framework.

Where the KF implementation proposed in [15] is unpractical for large-scale applications due to its large computational complexity, this paper presents a computationally efficient framework in which a single focal-plane image is used to obtain a high-resolution reconstruction of dynamic aberrations. By exploiting special matrix structures and efficient linear algebra operations specific to AO systems, the number of elementary computations scale just more than linearly ($n \log(n)$) with the total number of pixels n of the reconstruction. This computational complexity is of the same order as computationally efficient small-phase methods [14] and much faster than the existing KF-based method [15] (which scales with n^3). The low computational complexity opens up the possibility for real-time large-scale applications.

An additional goal of this paper is to address the advantages of reformulating the PR problem for dynamic aberrations into a KF framework. Due to the similarities between the KF measurement update and the PR problem, the issue of uniqueness of the PR problem becomes an issue of observability within the KF framework. Consequently, even when there is no unique solution to the PR problem when using a single image, the full dynamic system considered by KF can become observable. In addition, it will be discussed how the KF measurement update can be interpreted as an optimally weighted regularization to the static PR problem.

The KF framework makes the method applicable to cases where the phase aberrations change over time, i.e., dynamic phase aberrations. Moreover, by including the temporal and statistical models of the turbulence and sensors, KF is known to be robust against modeling errors and measurement noise, making it particularly suitable for dealing with dynamic aberrations and noisy measurements. Small-phase aberrations have also been assumed in other algorithms that aim to estimate temporal dynamic aberrations [15,16]. Possible applications include estimating NCPAs (which occur in various fields such as astronomy [17,18] and ophthalmology [19,20]), wind induced dynamic non-common path vibrations [21], or the low-wind effect [18]. This method could also be used to estimate spatial (rather than temporal) dynamics due to anisoplanatic conditions in astronomy [22], ophthalmology [23], and microscopy [24]. Small-phase aberrations are the typical situation considering a closed-loop AO system [12], making the algorithm well-suited to estimate AO-corrected atmospheric residuals.

The proposed algorithm establishes a generic framework that can be applied provided its requirements are satisfied (small-phase aberrations with continuous dynamic evolution). As an illustration, a simulation of a closed-loop AO system for astronomy using the open-source Matlab software Object-Oriented Matlab AO (OOMAO) [25] is performed, giving a detailed analysis of the robustness of the method against modeling inaccuracies and noise. A simple WFS-based controller compensates for the dominant low-order modes in the wavefront in order to keep the wavefront sufficiently small. It should be stressed that this case study is merely an example and should be considered as an illustration of the performance of the method under varying temporal dynamics, noise conditions, and RMS values of the aberration.

Table 1. List of Frequently Used Notations and Symbols

c_X	Average number of non-zero elements per row/column of any sparse banded matrix X
m	Number of actuators inside the aperture
$\bar{n}$, n	Size of grid in pupil plane ($\bar{n} \times \bar{n}$) and number of pixels inside aperture, respectively
$\bar{p}$, p	Size of grid in the image plane ($\bar{p} \times \bar{p}$) and $p = \bar{p}^2$
$\bar{q}$	Size of lenslet grid of the WFS ($\bar{q} \times \bar{q}$)
$\boldsymbol{\phi}(k)$	Aberrated residual wavefront
$\boldsymbol{a}$	Pupil plane amplitude
$\boldsymbol{u}(k)$	DM control signal
$\boldsymbol{v}(k)$, $R(k)$	Measurement noise and its covariance matrix
$\boldsymbol{y}(k)$	Point spread function (PSF)
A, Q	Turbulence dynamics model [see Eq. (10)]
J_ϕ, $\boldsymbol{c}_\phi$	Taylor expansion of $\boldsymbol{y}(k)$ [see Eq. (20)]
$\mathcal{F}$, $\mathcal{F}_{\text{vec}}$	Zero-padded 2D DFT and $\mathcal{F}_{\text{vec}}\{\cdot\} := \text{vec}(\mathcal{F}\{\cdot\})$
$\mathcal{J}_\phi$, $\mathcal{M}$, $\mathcal{B}$	Efficient operators [see Eqs. (21), (35), (36)]
$\mathcal{N}(\mu, P)$	Gaussian random process (mean μ, covariance P)
S_X	Set of all matrices that satisfy a chosen sparsity pattern for matrix X

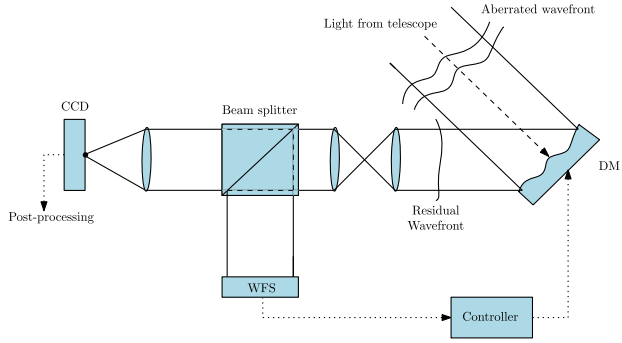

Fig. 1. Schematic of the classic AO system for astronomy.

This paper is structured as follows. Section 2 gives an introduction to modeling AO systems and dynamic aberrations. A computationally efficient representation of the model and its identification procedure is described in Section 3. The advantage of using KFs in WFS-less AO systems is discussed in Section 4. The main contribution of this paper, an efficient implementation of the KF, is introduced in Section 5. Section 6 discusses the simulation environment, and the results are presented in Section 7. Finally, the main conclusions are summarized in Section 8. Frequently used symbols are listed in Table 1.

2. MODELING THE AO SYSTEM

To provide an illustration of the algorithm's performance, a typical closed-loop AO system for astronomy depicted in Fig. 1 will be used. This section introduces the components, aberration dynamics, and measurement noise models.

A. Overview of the AO System

The effect of the phase aberration on the optical system will be represented in terms of the generalized pupil function (GPF):

$$\mathcal{X}(\rho, \theta) = \mathcal{A}(\rho, \theta) \exp\left(i\Phi(\rho, \theta)\right), \tag{1}$$

where ρ, θ represent the normalized polar coordinates in the pupil plane, and $\mathcal{A}(\rho, \theta)$ and $\Phi(\rho, \theta)$ represent the amplitude apodization function and phase aberration, respectively. Assuming uniform illumination, $\mathcal{A}(\rho, \theta)$ will correspond to the characteristic function with a constant value inside the pupil and zero outside. The image along the optical axis around the focal plane can be expressed in terms of the GPF via the following integral:

$$\mathcal{I}(\zeta, \omega, \delta) = \frac{1}{\pi} \int_0^1 \int_0^{2\pi} \exp\left(i\delta\zeta^2\right) \mathcal{X}(\rho, \theta) \ldots \times \exp(2\pi i \zeta \rho \cos(\theta - \omega)) \rho \mathrm{d}\rho \mathrm{d}\theta, \tag{2}$$

where ζ, ω are the normalized polar coordinates in the focal plane, and δ represents the position along the optical axis with respect to the focal plane. The PSF corresponding to the GPF is defined as the intensity of $\mathcal{I}(\zeta, \omega, \delta)$:

$$\mathcal{Y}(\zeta, \omega, \delta) = |\mathcal{I}(\zeta, \omega, \delta)|^2. \tag{3}$$

The pupil plane coordinates will be sampled on a regular $\bar{n} \times \bar{n}$ grid of pixels. The discretized GPF, its amplitude and its phase (the wavefront) will be denoted by $X \in \mathbb{C}^{\bar{n}\times\bar{n}}$, $\mathcal{A} \in \mathbb{R}^{\bar{n}\times\bar{n}}$, and $\Phi \in \mathbb{R}^{\bar{n}\times\bar{n}}$, respectively, related via

$$X = \mathcal{A} \odot \exp i\Phi, \tag{4}$$

where $\odot$ represents an element-wise product of two vectors or matrices. Often, a vectorized representation will be used in which only the pixels inside the aperture, i.e., the non-zero elements in $\mathcal{A}$, will be considered:

$$\boldsymbol{x} = \boldsymbol{a} \odot \exp i\boldsymbol{\phi}, \tag{5}$$

so that $\boldsymbol{a} \in \mathbb{R}^n$, $\boldsymbol{x} \in \mathbb{C}^n$, and $\boldsymbol{\phi} \in \mathbb{R}^n$. Moreover, due to the circular aperture, it can be concluded that $n < \bar{n}^2$.

Similarly, the image plane coordinates are sampled on a $\bar{p} \times \bar{p}$ grid, and the PSF will be denoted by $Y \in \mathbb{R}^{\bar{p}\times\bar{p}}$. A discrete adaptation of Eq. (2) becomes the two-dimensional (2D) discrete Fourier transform (DFT) denoted by $\mathcal{F}(\cdot)$:

$$Y = |\mathcal{F}\{X\}|^2. \tag{6}$$

To obtain a Nyquist sampled measurement, it is necessary for simulation purposes that the matrix $X \in \mathbb{C}^{\bar{n}\times\bar{n}}$ is padded with zeros to the dimension of $\bar{p} \times \bar{p}$, where $\bar{p} = 2\bar{n}$. A vectorized representation of Eq. (6) is defined by

$$\boldsymbol{y} = |\mathcal{F}_{\text{vec}}\{\boldsymbol{x}\}|^2, \tag{7}$$

where $\mathcal{F}_{\text{vec}}\{\cdot\} := \text{vec}(\mathcal{F}\{\cdot\})$ and $\boldsymbol{y} \in \mathbb{R}^p$: $\boldsymbol{y} := \text{vec}(Y)$ such that $p = \bar{p}^2$, with $\text{vec}(\cdot)$ representing a vectorization of a matrix.

The WFS is assumed to have a square grid of $\bar{q} \times \bar{q}$ lenslets. Measurements from lenslets that do not collect enough light due to their position outside the aperture will be discarded. The effect of the DM on the wavefront is modeled by its influence functions. The actuators are placed such that their centers lie on a $(\bar{q}+1) \times (\bar{q}+1)$ grid located at the corners of the WFS lenslet subapertures (known as Fried geometry). Actuators outside of the aperture will be discarded, hence the number of active actuators $m < (\bar{q}+1)^2$. The wavefront correction by the DM will be denoted by $\Phi_{\text{dm}} \in \mathbb{R}^{\bar{n}\times\bar{n}}$, and it is defined as

$$\Phi_{\text{dm}} = \sum_{\ell=1}^{m} B_\ell u_\ell, \tag{8}$$

where $B_\ell \in \mathbb{R}^{\bar{n}\times\bar{n}}$ represents the DM influence functions. The control commands $u_\ell \in \mathbb{R}$, $\ell = 1, \ldots, m$ are stored in a vector $\boldsymbol{u} \in \mathbb{R}^m$. A vectorized formulation of Φ_{dm} is represented by $\boldsymbol{\phi}_{\text{dm}} \in \mathbb{R}^n$: $\boldsymbol{\phi}_{\text{dm}} := \text{vec}(\Phi_{\text{dm}})$, and

$$\boldsymbol{\phi}_{\text{dm}} = B\boldsymbol{u}, \tag{9}$$

such that each column ℓ of $B \in \mathbb{R}^{n\times m}$ is equal to $\text{vec}(B_\ell)$.

B. Dynamic Aberration Model

In astronomy, the phase aberrations are caused by turbulence in the earth's atmosphere. Since atmospheric turbulence is a well-studied source of dynamic aberrations, it is one of the main reasons for choosing the astronomy example to illustrate the algorithm's performance. The temporal dynamics are usually modeled by assuming Taylor's frozen turbulence assumption, where the atmosphere is modeled by multiple layers, each moving in a constant direction and speed. The layer's combined influence on the wavefront is denoted by $\boldsymbol{\phi}_t \in \mathbb{R}^n$. The temporal dynamics of $\boldsymbol{\phi}_t$ are estimated by a vector auto-regressive (VAR) model of order one, i.e.,

$$\boldsymbol{\phi}_t(k+1) = A\boldsymbol{\phi}_t(k) + \boldsymbol{w}(k), \tag{10}$$

with k specifying the time index, $A \in \mathbb{R}^{n\times n}$, and where $\boldsymbol{w}(k) \in \mathbb{R}^n$ is a Gaussian random process with zero mean and covariance matrix $Q \in \mathbb{R}^{n\times n}$, which will be denoted in this paper as $\boldsymbol{w}(k) \sim \mathcal{N}(0, Q)$, where the symbol $\sim$ declares the statistical distribution of the variable on its left. VAR models have been a popular choice to model aberration dynamics and have been particularly used to model turbulence dynamics [26]. During the simulations, the turbulence will be simulated using OOMAO, so the model of Eq. (10) is not the true model driving the turbulence dynamics. More detailed information in how to obtain the VAR model from data is discussed in Section 3.

Closing the loop using the DM model of Eq. (9), the residual wavefront $\boldsymbol{\phi}(k) \in \mathbb{R}^n$, defined as $\boldsymbol{\phi}(k) = \boldsymbol{\phi}_t(k) - \boldsymbol{\phi}_{\text{dm}}(k)$, which is combined with Eq. (10), becomes

$$\boldsymbol{\phi}(k+1) = A\boldsymbol{\phi}(k) - B\boldsymbol{u}(k) + AB\boldsymbol{u}(k-1) + \boldsymbol{w}(k). \tag{11}$$

Finally, using $\boldsymbol{\phi}$ as the state and Eq. (11) as the state update equation, a state-space model [27] can be formulated by including the measured PSF as an output equation:

$$\boldsymbol{y}(k) = |\mathcal{F}_{\text{vec}}\{\boldsymbol{a} \odot \exp(i\boldsymbol{\phi}(k))\}|^2 + \boldsymbol{v}(k). \tag{12}$$

The measurement noise $\boldsymbol{v}(k)$ is in reality a combination of Guassian read-out noise and Poissonian shot noise, but will be modeled as a Gaussian process: $\boldsymbol{v}(k) \sim \mathcal{N}(0, R(k))$, this will be motivated in the next subsection.

C. Measurement Noise Model

The true measurement noise can be seen as combination of Gaussian read-out noise and Poissonian shot noise, i.e.,

$$\boldsymbol{y} = \boldsymbol{y}_{\text{true}} + \boldsymbol{v}_{\text{shot}} + \boldsymbol{v}_{\text{read}}, \tag{13}$$

where $\boldsymbol{y}_{\text{true}}$ is the number of photons that would have arrived at each pixel of the camera in the noiseless case. The read-out noise is a zero-mean white Gaussian noise $\boldsymbol{v}_{\text{read}} \sim \mathcal{N}(0, \sigma_r^2 I)$. The shot noise is known to correspond to a Poisson distribution: $(\boldsymbol{y}_{\text{true}} + \boldsymbol{v}_{\text{shot}}) \sim \text{Pois}(\boldsymbol{y}_{\text{true}})$.

By approximating the Poisson distribution by a Gaussian distribution, the use of the Kalman filter theory is allowed. An important reason this approximation is adopted is that the Poisson distribution is known to converge to a Gaussian distribution when y_{true} is large, i.e.,

$$\lim_{y_{\text{true}}\to\infty} \text{Pois}(y_{\text{true}}) = \mathcal{N}(y_{\text{true}}, y_{\text{true}}), \tag{14}$$

where in this case y_{true} represents a single pixel. Consequently, if y_{true} is large,

$$y \approx y_{\text{true}} + \tilde{v}_{\text{shot}} + v_{\text{read}}, \tag{15}$$

with $\tilde{v}_{\text{shot}} \sim \mathcal{N}(0, y_{\text{true}})$ and $v_{\text{read}} \sim \mathcal{N}(0, \sigma_r^2)$. Since the read-out and shot noise can be considered to be uncorrelated, the total noise is $v := \tilde{v}_{\text{shot}} + v_{\text{read}}$, $v \sim \mathcal{N}(0, y_{\text{true}} + \sigma_r^2)$. Of course, this remains an approximation, so the robustness of the method against this modeling inaccuracy will be discussed in Section 7.

When the approximation of Eq. (15) is extended to the multivariate case of Eq. (13), the approximation of the shot noise becomes $\tilde{\boldsymbol{v}}_{\text{shot}}(k) \sim \mathcal{N}(0, \text{diag}(\boldsymbol{y}(k)))$, with $\text{diag}(\boldsymbol{y})$ a diagonal matrix with the entries of $\boldsymbol{y}$ on its diagonal. The total measurement noise, given by $\boldsymbol{v}(k)$ in Eq. (12), will be modeled by

$$\boldsymbol{v}(k) \sim \mathcal{N}(0,\ R(k)), \quad \text{with} \quad R(k) = \sigma_r^2 I + \text{diag}(\boldsymbol{y}(k)). \tag{16}$$

In the simulations, the shot noise and read-out noise will be simulated according to their true distributions, and the approximation of Eq. (16) is not used to simulate the noise.

Furthermore, since the number of arriving photons (y_{true}) influences the severity of the noise, the brightness of the point source becomes an important parameter. A commonly used unit for the flux of photons arriving at the telescope is the (Vega) magnitude, β_m. The relation β_m and the photon flux β_f in photons/m^2/s depends on the photometric system:

$$\beta_f = c_\beta 10^{-\frac{2}{5}\beta_m}, \tag{17}$$

with constant $c_\beta \approx 8.97 \cdot 10^9$ for the simulations in Section 7.

3. STRUCTURES AND EFFICIENT COMPUTATIONS IN AO SYSTEMS

This section discusses a number of structures of matrices that appear in the state-space model defined by Eqs. (11) and (12). The model structures and efficient operators presented in this section will be used to develop an efficient non-linear Kalman filter implementation discussed in Section 5.

A. Identification of Structured Matrices in Dynamic Aberration Models

The key realization for efficiently modeling the aberration dynamics is that the VAR with exogenious variables (VARX) model of Eq. (11) can be accurately represented by highly sparse matrices A, Q^{-1}, and B. When a type of DM is considered, for which its actuators only have a local influence on the total wavefront, B will be a sparse matrix. The remainder of this subsection will focus on the sparsity of the matrices A and Q^{-1}.

To support this assumption, a graphical model representation of the VARX model can be considered [28]. An intuitive interpretation of the theory in [28] to atmospheric turbulence implies that the matrices A and Q^{-1} are sparse banded matrices. This banded sparsity corresponds to the ability of predicting a single pixel only using those in its close neighbourhood. Such an assumption holds, for example, in situations where the turbulence can indeed be approximated by Taylor's frozen flow hypothesis, as is done in this work. The sparsity structure of the matrices A and Q^{-1} is defined by the sets $\mathcal{S}_A$ and $\mathcal{S}_{Q^{-1}}$, respectively, each denoting the location of the non-zero entries in their corresponding matrix. A more detailed analysis of sparsity structures in dynamic AO systems can be found in [29,30]. Sparsity structures in AO systems are not unique to the dynamic aberration model. For example, [31] shows that the sparsity structure can also be exploited for static wavefront reconstruction.

To compute A and Q^{-1}, different approaches are possible, depending on the available information, such as computing them from first principles by manual tuning or via system identification techniques [27]. The remainder of this subsection will show how A and Q^{-1} can be computed via system identification; however, if the covariance matrix $C_\phi = E[\boldsymbol{\phi}(k)\boldsymbol{\phi}^T(k)]$ is known, A and Q are related via $Q = C_\phi - AC_\phi A^T$.

To identify A and Q^{-1} from data, a time series with N_{id} time samples of open-loop wavefront data is retrieved using either a dedicated WFS or any WFS-less wavefront reconstruction method. This batch will be denoted by $\{\boldsymbol{\phi}_{id}(i) | i = 1, 2, \ldots, N_{id}\}$. The matrix A can be identified by solving the following constrained least squares problem [27]:

$$\begin{aligned} \min_{A} \quad & \|A\Phi_1 - \Phi_2\|_F^2 \\ \text{s.t.} \quad & A \in \mathcal{S}_A, \end{aligned} \tag{18}$$

where $\Phi_1 = [\boldsymbol{\phi}_{id}(1) \cdots \boldsymbol{\phi}_{id}(N_{id} - 1)]$ and $\Phi_2 = [\boldsymbol{\phi}_{id}(2) \cdots \boldsymbol{\phi}_{id}(N_{id})]$. The set $\mathcal{S}_A$ describes the set of all matrices A that have a desired sparsity pattern. Since the exact optimal sparsity pattern is usually unknown, but rather an over-approximation is used, one-norm regularization on A can be added to further increase the sparsity [30].

Having computed an estimate of A, an estimate of $\boldsymbol{w}(k)$ can be found by computing $\hat{\boldsymbol{w}}(k) = \boldsymbol{\phi}_{id}(k+1) - A\boldsymbol{\phi}_{id}(k)$ for $k = 1, \ldots, N_{id} - 1$. Afterwards, a sample covariance matrix,

$$Q_s = \frac{1}{N_{id} - 1} \sum_{i=1}^{N_{id}-1} \hat{\boldsymbol{w}}(k)\hat{\boldsymbol{w}}^T(k), \tag{19}$$

can be computed such that a sparse approximation of Q_s^{-1} could serve as an estimate of Q^{-1}. Alternatively, Q_s^{-1} can be estimated directly from $\hat{\boldsymbol{w}}(k)$ by creating a matrix W, where each kth column of W corresponds to $\frac{1}{\sqrt{N_{id}-1}}\hat{\boldsymbol{w}}(k)$. Next, a matrix Ξ is computed such that $\Xi W = O$, $O \in \mathbb{R}^{n \times n}$ being an orthonormal matrix. This matrix Ξ can be found using the Gramm–Schmidt method, and Q_s^{-1} can then be computed via $Q_s^{-1} = \Xi^T\Xi$, since $Q_s = WW^T = \Xi^{-1}OO^T\Xi^{-T} = \Xi^{-1}\Xi^{-T}$. Hence, Ξ can be seen as the inverse square root of Q_s. Since Q^{-1} is expected to be banded and sparse, its square root Ξ is expected to have an accurate sparse approximation. During the Gramm–Schmidt procedure, the desired sparsity pattern of Ξ can easily be enforced, both speeding up the process and finding a sparse approximation.

Since the model of Eq. (10) is not capable of exactly representing the true turbulence dynamics, manually fine-tuning $\hat{Q}$ could improve the performance. Also, when n is very large, and the dataset N_{id} is limited, problems could occur regarding ill-conditioning of the matrix Q_s. One simple way of tuning $\hat{Q}$ is to make it equal to $Q_s + \gamma I$. As a

rule of thumb throughout this paper, γ is tuned such that $\text{mean}(\text{diag}(\Xi^T\Xi)) = 1/\text{mean}(\text{diag}(\hat{Q}))$.

Furthermore, for applications where the sparsity of Q^{-1} would be unknown, finding a sparse inverse of an sample covariance Q_s has been studied well in graphical modeling literature. Several algorithms have been presented that solve this problem, which can be used to increase the sparsity of the estimate Q^{-1} further when desired, see for example [32].

B. Efficient Computations of the Output Equation

Since the output equation is non-linear, a first-order Taylor approximation with respect to $\boldsymbol{\phi}$ can be used to linearlize Eq. (12). The first-order Taylor approximation around the initial estimate $\hat{\boldsymbol{\phi}}$ will be denoted by

$$\boldsymbol{y}(k) \approx J_{\hat{\phi}}(\boldsymbol{\phi}(k) - \hat{\boldsymbol{\phi}}) + \boldsymbol{c}_{\hat{\phi}} + \boldsymbol{v}(k), \tag{20}$$

where, if we define $f(\boldsymbol{\phi}) = |\mathcal{F}_{\text{vec}}\{\boldsymbol{a} \odot \exp(i\boldsymbol{\phi})\}|^2$, $\boldsymbol{c}_{\hat{\phi}} = f(\hat{\boldsymbol{\phi}})$ and $J_{\hat{\phi}} = f'(\hat{\boldsymbol{\phi}})$, i.e., the Jacobian matrix corresponding $f(\boldsymbol{\phi})$ evaluated at $\hat{\boldsymbol{\phi}}$.

Although there are no sparse matrices in the output equation, Eq. (12), the efficient FFT algorithm can be used to compute the 2D DFT. For the first-order Taylor approximation of Eq. (20), an efficient computation of the matrix-vector multiplication $J_{\hat{\phi}}\boldsymbol{\phi}$ can be formulated using the FFT algorithm, thereby avoiding the explicit computation of the complete Jacobian. When introducing $\hat{\boldsymbol{x}} = \boldsymbol{a} \odot \exp(i\hat{\boldsymbol{\phi}})$ and $\hat{\boldsymbol{y}} = \mathcal{F}_{\text{vec}}\{\hat{\boldsymbol{x}}\}$, the following operators can be introduced:

$$J_{\hat{\phi}}\boldsymbol{\phi} \rightarrow \mathcal{J}_{\hat{\phi}}(\boldsymbol{\phi}) \triangleq \Re\left(2i\hat{\boldsymbol{y}} \odot \mathcal{F}_{\text{vec}}\left\{\hat{\boldsymbol{x}} \odot \boldsymbol{\phi}\right\}\right), \tag{21}$$

$$J_{\hat{\phi}}^T\boldsymbol{\psi} \rightarrow \mathcal{J}_{\hat{\phi}}^T(\boldsymbol{\psi}) \triangleq \Re\left(-2i\hat{\boldsymbol{y}}^* \odot \mathcal{F}_{\text{vec}}^{-1}\left\{\hat{\boldsymbol{x}}^* \odot \boldsymbol{\psi}\right\}\right), \tag{22}$$

where $\mathcal{F}_{\text{vec}}^{-1}\{\cdot\}$ represents the vectorized 2D inverse DFT, and $\hat{\boldsymbol{y}}^*$ and $\hat{\boldsymbol{x}}^*$ represent the element-wise complex conjugates of $\hat{\boldsymbol{y}}$ and $\hat{\boldsymbol{x}}$, respectively. In Section 5, both the sparsity structures of A and Q^{-1} and the operators in Eqs. (21) and (22) will be used to develop an efficient non-linear Kalman filter implementation.

4. KALMAN FILTERS FOR WFS-LESS AO

Kalman filters are widely used as a state observer for many applications, and they are treated in many textbooks [27]. At each time step k, the Kalman filter involves two steps: a *measurement update* and a *time update*. The goal of the measurement update is to optimally combine a new measurement $\boldsymbol{y}(k)$ with a previous state estimate $\hat{\boldsymbol{\phi}}(k|k-1)$ into an improved update $\hat{\boldsymbol{\phi}}(k|k)$. The time update uses the dynamical model to predict the state for the next time step, which will be denoted by $\hat{\boldsymbol{\phi}}(k+1|k)$.

The accuracy of the estimates is expressed in terms of the error-covariance matrices $P(k+1|k)$ and $P(k|k)$ for the time update and measurement update, respectively:

$$P(k+1|k) := E[(\boldsymbol{\phi}(k) - \hat{\boldsymbol{\phi}}(k+1|k))(\boldsymbol{\phi}(k) - \hat{\boldsymbol{\phi}}(k+1|k))^T], \tag{23}$$

$$P(k|k) := E[(\boldsymbol{\phi}(k) - \hat{\boldsymbol{\phi}}(k|k))(\boldsymbol{\phi}(k) - \hat{\boldsymbol{\phi}}(k|k))^T]. \tag{24}$$

The Kalman filter computes the estimates $\hat{\boldsymbol{\phi}}(k+1|k)$ and $\hat{\boldsymbol{\phi}}(k|k)$, such that their error-covariance matrices are minimal [27]. Since the state equation, Eq. (11), is linear, the time update will be equivalent to the classical Kalman filter. The measurement update can be formulated as a regularized non-linear least squares problem:

$$\hat{\boldsymbol{\phi}}(k|k) = \arg\min_{\boldsymbol{\phi}\in\mathbb{R}^n} \left\|\hat{\boldsymbol{\phi}}(k|k-1) - \boldsymbol{\phi}\right\|^2_{P^{-1}(k|k-1)} + \cdots$$

$$\left\|\boldsymbol{y}(k) - \left|\mathcal{F}_{\text{vec}}\left\{\boldsymbol{a} \odot \exp(i\boldsymbol{\phi})\right\}\right|^2\right\|^2_{R^{-1}(k)}, \tag{25}$$

where the weight matrix $P(k|k-1)$ is the error-covariance matrix corresponding to $\boldsymbol{\phi}(k) - \hat{\boldsymbol{\phi}}(k|k-1)$. Because of the non-linearity of the output equation of the state-space model, a non-linear adaptation of the measurement update has to be used.

A. Kalman Filtering and the Phase Retrieval Problem

If the Kalman filter measurement update is compared to the PR problem, many similarities appear, see, for example, [33] for an overview of different phase estimation methods. Many algorithms have been proposed that aim to solve the PR problem in its original non-convex optimization framework [7,10]. In general, the optimization framework could be formulated as

$$\min_{\boldsymbol{x}\in\mathbb{C}^n} \| \boldsymbol{y} - |\mathcal{F}_{\text{vec}}(\boldsymbol{x})|^2 \|^2_{W_y} + \| \hat{\boldsymbol{x}} - \boldsymbol{x} \|^2_{W_x}, \tag{26}$$

where W_x, W_y are weight matrices and for any vector $\boldsymbol{\alpha}$ and matrix M: $\| \boldsymbol{\alpha} \|^2_M := \boldsymbol{\alpha}^H M \boldsymbol{\alpha}$. Most methods do not use the second term, as there generally is no prior information $(\hat{\boldsymbol{x}}, W_x)$ considered to be available. Due to the non-convexity of the first term, a very accurate initial condition is computed first before attempting to solve the problem using either a gradient descent or second-order optimization algorithm. Obtaining this initial estimate when considering Fourier measurements is not trivial. This is where a Kalman filter and knowledge of the aberration dynamics can be used to its advantage.

When the amplitude in the pupil plane $\boldsymbol{a}$ is known, the problem can be rewritten to solve for $\boldsymbol{\phi}$ directly:

$$\min_{\boldsymbol{\phi}\in\mathbb{R}^n} \| \boldsymbol{y} - |\mathcal{F}_{\text{vec}}\{\boldsymbol{a} \odot \exp(i\boldsymbol{\phi})\}|^2 \|^2_{W_y} + \| \hat{\boldsymbol{\phi}} - \boldsymbol{\phi} \|^2_{W_\phi}, \tag{27}$$

with W_ϕ as another weighting matrix. Although this formulation of the PR problem is not often used in existing literature due to the required extra information $\boldsymbol{a}$ and the extra non-linearity caused by the exponential $\exp(i\boldsymbol{\phi})$, it is more convenient for the dynamic case due to the convenient linear state update equation, Eq. (11), in terms of $\boldsymbol{\phi}$. Moreover, when aiming to reconstruct the wavefront from a single image, knowledge of $\boldsymbol{a}$ is a significant advantage for obtaining a good performance. Considering the formulation of the PR problem as in Eq. (27) and comparing it to the Kalman filter measurement update of Eq. (25), the similarities are obvious. The Kalman filter provides a framework to define optimal weight matrices W_ϕ and W_y and uses the model dynamics to find the initial estimate $\hat{\boldsymbol{\phi}}$.

B. Solving the Phase Retrieval Problem Using a Kalman Filter

The next step is to solve the non-linear measurement update of Eq. (25). The most well-known method is to use a linear Taylor approximation to solve the non-linear measurement update, which is known in the Kalman filter theory as the extended Kalman filter (EKF). In this paper, the Gauss–Newton algorithm is applied to the non-linear optimization problem. Solving the measurement update of Eq. (25) using the Gauss–Newton algorithm leads to a filter that is known as the iterated EKF (IKF) [34]. The IKF, as the name suggests, is an iterative version of the EKF. It is interesting to draw similarities here with other existing PR algorithms in literature that aim to solve the related formulation of Eq. (26) via a gradient descend or Gauss–Newton optimization schemes [7,10].

When a straightforward implementation of the IKF is used, the iterations of the measurement update become as follows: initializing $\boldsymbol{\phi}_0 = \hat{\boldsymbol{\phi}}(k|k-1)$ and $P_0 = P(k|k-1)$ for $\ell = 0, 1, \ldots, L$:

$$\boldsymbol{\phi}_{\ell+1} = (I - K_\ell J_{\phi_\ell})\hat{\boldsymbol{\phi}}(k|k-1) + K_\ell \left(\boldsymbol{y}(k) - \boldsymbol{c}_{\phi_\ell} + J_{\phi_\ell}\boldsymbol{\phi}_\ell\right), \tag{28}$$

$$P_{\ell+1} = \left(I - K_\ell J_{\phi_\ell}\right) P(k|k-1), \tag{29}$$

with J_{ϕ_ℓ} and $\boldsymbol{c}_{\phi_\ell}$ as defined in Eq. (20), where

$$K_\ell = P(k|k-1) J_{\phi_\ell}^T \left(J_{\phi_\ell} P(k|k-1) J_{\phi_\ell}^T + R(k)\right)^{-1}, \tag{30}$$

and, after the last iteration, $\hat{\boldsymbol{\phi}}(k|k) = \boldsymbol{\phi}_L$ and $P(k|k) = P_L$. The time update becomes equal to the standard KF update:

$$\hat{\boldsymbol{\phi}}(k+1|k) = A\hat{\boldsymbol{\phi}}(k|k) - B\boldsymbol{u}(k) + AB\boldsymbol{u}(k-1), \tag{31}$$

$$P(k+1|k) = AP(k|k)A^T + Q. \tag{32}$$

The most important problem with this method is its large computational complexity; therefore, this paper presents a computationally efficient formulation of the IKF in Section 5.

C. Advantage Dynamic versus Static Approach

This subsection discusses the advantages of considering the dynamic case with respect to the static case. For both cases, we restrict this discussion to the case in which the pupil plane amplitude is considered to be known, i.e., the optimization problem of Eq. (27).

In the static case, without any knowledge of the model dynamics, the PR aims to solve Eq. (27) for $W_y = I$ and $W_\phi = 0$ (i.e., disregarding the second term). This often creates problems regarding the non-uniqueness of the solution when only a single image is used, and, without an accurate initial guess, algorithms usually do not converge to an accurate solution. Usually, the uniqueness problem is overcome by considering multiple images along the optical axis [also known as phase diversity (PD) [35]], but this involves splitting the light further and re-introduces the problem with NCPAs.

When considering the dynamic case, the PR problem is reformulated as the measurement update of the Kalman filter, i.e., Eq. (25). Compared to the static case, the problem of uniqueness becomes the problem of *observability* [27]. Where the uniqueness in the static case is purely defined by the output equation, observability depends on both the output equation and the state dynamics. Even when the output equation alone does not lead to a unique solution in the static case, adding knowledge of the dynamics can make the system observable. As a consequence, the Kalman filter measurement update in Eq. (25) is much more likely to result in an accurate estimate compared to the static case. The prior information $\hat{\boldsymbol{\phi}}(k|k-1)$, $P(k|k-1)$, $R(k)$ obtained from the dynamic model and noise model, acts as an optimally weighted *regularization* to the static PR problem. Since the regularization term $\|\hat{\boldsymbol{\phi}}(k|k-1) - \boldsymbol{\phi}\|^2_{P(k|k-1)^{-1}}$ is convex, it significantly helps solvers aiming to solve the non-linear optimization problem in Eq. (25). The addition of this extra knowledge opens the possibility for considering situations impossible to accurately solve in the static case, such as only considering a single out of focus image, or even considering a single in focus image.

Another advantage of this optimally weighted regularization term is that due to the extra knowledge of the measurement noise, the algorithm also becomes much less sensitive to noise. A common method to deal with measurement noise in the static case is to hard-threshold the image and discard all measured pixels in $\boldsymbol{y}(k)$ smaller than $y_{\min}$ from the measurements. Since the optimal choice of $y_{\min}$ depends on the noise level σ_r, the value of $y_{\min}$ is difficult to optimally tune. The Kalman filter does not need this truncation parameter, as the noise is already taken into account by the application of the weighting matrices $R(k)$ and $P(k|k-1)$, resulting in a more consistent algorithm under noisy circumstances.

5. EFFICIENT NON-LINEAR KALMAN FILTER IMPLEMENTATION

In Section 3, it was discussed how, due to the underlying nature of the dynamics, the inverse covariance matrices corresponding to $\boldsymbol{\phi}(k)$ and $\boldsymbol{w}(k)$ (i.e., C_ϕ^{-1} and Q^{-1}) are expected to be sparse. As a consequence, the inverse error-covariance matrix of the KF, P^{-1}, is expected to have an accurate sparse approximation as well. Therefore, the *information filter* representation of the KF is used, which only uses the inverse error-covariance matrices $P^{-1}(k|k)$ and $P^{-1}(k+1|k)$. To summarize, the new filter consists of four main steps: (1) compute the measurement update $\hat{\boldsymbol{\phi}}(k|k)$ via Eq. (25), (2) compute inverse error-covariance matrix $P^{-1}(k|k)$, (3) compute time update $\hat{\boldsymbol{\phi}}(k+1|k)$ via Eq. (31), and (4) compute inverse error-covariance matrix $P^{-1}(k+1|k)$. The computation of the time update $\hat{\boldsymbol{\phi}}(k+1|k)$ (step 3) is straightforward. Since the matrix A and B are sparse, this update can be performed very efficiently. Efficient implementations of the other steps will be discussed in the remainder of this section.

The notation $\mathcal{O}(\cdot)$ will be used to describe how the number of elementary computations ($+$, $-$, $\times$, $\div$) scale with respect to certain parameters. For example, the notation $\mathcal{O}(n^3)$ represents that if the system dimension n doubles, the number of required

Table 2. Computational Complexities Per Step of the Algorithm[a]

Step in Algorithm	Computational Complexity
Compute $\phi(k\|k)$ by solving Eq. (34) via CG method	$\mathcal{O}(n(c_{P^{-1}}+\log n)\mathrm{LL}_{\mathrm{CG}})$
Explicitly computing p_c rows of the Jacobian	$\mathcal{O}(np_c)$
Sparse low-rank approximation of $P^{-1}(k\|k)$ in Eq. (37)	$\mathcal{O}(np_cc_{P^{-1}})$
Compute $\hat{\phi}(k+1\|k)$ via Eq. (31)	$\mathcal{O}(nc_A+mc_B)$
Computing sparse inverse approximation via solving Eq. (39)	$\mathcal{O}(n(c_\Gamma^3+c_{M_2}c_\Gamma^2))$
Sparse approximation of $P^{-1}(k+1\|k)$ in Eq. (40)	$\mathcal{O}(nc_{M_1}^2c_{P^{-1}})$

[a] p_c represents the rank of the approximated term in Eq. (37), L is the number of the IKF iterations, and L_{CG} is the average number of CG iterations. The other symbols are included in Table 1.

elementary computations increase by a factor $2^3=8$. A method is considered to scale better with respect to the system dimensions when the exponent is as low as possible. All computational complexities will be summarized in Table 2.

An important advantage of the implementation presented is that the state measurement update is computed efficiently using the full PSF as input, i.e., no truncation or cropping of the image is performed. The only necessary approximations are in the updates of the matrix P^{-1}. Since P^{-1} acts as just a weight matrix in the measurement update, it is expected that small inaccuracies in P^{-1} affect the performance of the algorithm much less significantly than manually truncating the data $\boldsymbol{y}(k)$.

A. State Measurement Update

Recall the Kalman filter measurement update formulation given by Eq. (25). Introducing the increment $\delta\boldsymbol{\phi}_\ell := \boldsymbol{\phi}_{\ell+1}-\boldsymbol{\phi}_\ell$, the ℓth iteration of the IKF measurement update can be reformulated into the following optimization problem:

$$\delta\boldsymbol{\phi}_\ell = \arg\min_{\delta\boldsymbol{\phi}\in\mathbb{R}^n} \left\|\hat{\boldsymbol{\phi}}(k|k-1)-\boldsymbol{\phi}_\ell-\delta\boldsymbol{\phi}\right\|^2_{P^{-1}(k|k-1)} + \left\|\boldsymbol{y}(k)-J_{\phi_\ell}\delta\boldsymbol{\phi}-\boldsymbol{c}_{\phi_\ell}\right\|^2_{R^{-1}(k)}, \tag{33}$$

where $\boldsymbol{\phi}_{\ell+1}=\delta\boldsymbol{\phi}_\ell+\boldsymbol{\phi}_\ell$ is equal to the results presented in Eq. (28). Since computing Eq. (28) by using Eq. (30) has a computational complexity of $\mathcal{O}(n^3)$, it is too computationally demanding for large values of n. Therefore, the above least squares problem is solved using an iterative solver that exploits the fast operators of Eqs. (21) and (22). One possible iterative algorithm is the conjugate gradient (CG) algorithm, which can be found in many textbooks, e.g., [36].

First, the normal equations corresponding to the least squares problem are formulated:

$$\left(J_{\phi_\ell}^TR^{-1}(k)J_{\phi_\ell}+P^{-1}(k|k-1)\right)\delta\boldsymbol{\phi}=\ldots$$

$$J_{\phi_\ell}^TR^{-1}(k)\left(\boldsymbol{y}(k)-\boldsymbol{c}_{\phi_\ell}\right)+P^{-1}(k|k-1)\left(\hat{\boldsymbol{\phi}}(k|k-1)-\boldsymbol{\phi}_\ell\right), \tag{34}$$

Algorithm 1. Efficient IKF Meausurement Update

procedure IKF-MU($\boldsymbol{y}(k)$, $\delta\boldsymbol{\phi}_0$, tol)
  $\boldsymbol{\phi}_0 \leftarrow \hat{\boldsymbol{\phi}}(k|k-1)$
  for $\ell \leftarrow 1:L$ **do** ▷ Start IKF iterations
    $\boldsymbol{r}_0 \leftarrow \mathcal{B}(\boldsymbol{y}(k);\ell,k)-\mathcal{M}(\delta\boldsymbol{\phi}_0;\ell,k)$ ▷ Eq. (35), (36)
    $\boldsymbol{p}_0 \leftarrow \boldsymbol{r}_0$
    $i \leftarrow 0$
    while $\boldsymbol{r}^T\boldsymbol{r} > tol$ **do** ▷ Start CG iterations
      $\boldsymbol{z}_i \leftarrow \mathcal{M}(\delta\boldsymbol{\phi}_i;\ell,k)$ ▷ Eq. (35)
      $\alpha_i \leftarrow \frac{\boldsymbol{r}_i^T\boldsymbol{r}_i}{\boldsymbol{p}_i^T\boldsymbol{z}_i}$
      $\delta\boldsymbol{\phi}_{i+1} \leftarrow \delta\boldsymbol{\phi}_i+\alpha_i\boldsymbol{p}_i$
      $\boldsymbol{r}_{i+1} \leftarrow \boldsymbol{r}_i-\alpha_i\boldsymbol{z}_i$
      $\beta_i \leftarrow \frac{\boldsymbol{r}_{i+1}^T\boldsymbol{r}_{i+1}}{\boldsymbol{r}_i^T\boldsymbol{r}_i}$
      $\boldsymbol{p}_{i+1} = \boldsymbol{r}_{i+1}-\beta_i\boldsymbol{p}_k$
      $i \leftarrow i+1$
    $\boldsymbol{\phi}_{\ell+1} \leftarrow \delta\boldsymbol{\phi}_{i+1}+\boldsymbol{\phi}_\ell$ ▷ End CG iterations
  $\hat{\boldsymbol{\phi}}(k|k) \leftarrow \boldsymbol{\phi}_\ell$ ▷ End IKF iterations
  return

such that the matrix on the left-hand side is square, symmetric, and positive definite. With the insights of Section 3, evaluating the vectors on both the left- and right-hand side of this system of equations requires only matrix-vector multiplications with sparse matrices P^{-1} and R^{-1} and operators of Eqs. (21) and (22). Hence, two new efficient operators are introduced:

$$\mathcal{M}(\delta\boldsymbol{\phi};\ell,k) := \mathcal{J}_{\boldsymbol{\phi}_\ell}^T\left(R^{-1}(k)\mathcal{J}_{\phi_\ell}(\delta\boldsymbol{\phi})\right)+P^{-1}(k|k-1)\delta\boldsymbol{\phi}, \tag{35}$$

$$\mathcal{B}(\boldsymbol{y};\ell,k) := \mathcal{J}_{\boldsymbol{\phi}_\ell}^T\left(R^{-1}(k)\left(\boldsymbol{y}(k)-\boldsymbol{c}_{\phi_\ell}\right)\right)+\cdots P^{-1}(k|k-1)\left(\hat{\boldsymbol{\phi}}(k|k-1)-\boldsymbol{\phi}_\ell\right), \tag{36}$$

such that $\mathcal{M}(\delta\boldsymbol{\phi};\ell,k)$ is equal to the vector on the left-hand side of Eq. (34) and $\mathcal{B}(\boldsymbol{y}(k);\ell,k)$ to the right-hand side. Since R^{-1} is diagonal, both Eqs. (35) and (36) have a computational complexity of $\mathcal{O}(n(c_{P^{-1}}+\log n))$, $c_{P^{-1}}$ being the average number of non-zero elements per row/column of P^{-1}. This allows the CG algorithm to find the solution $\hat{\boldsymbol{\phi}}(k|k)$ in a computationally efficient manner, presented in Algorithm 1.

B. Error-Covariance Measurement Update

The next step is to find an update of the matrix $P^{-1}(k|k)$:

$$P^{-1}(k|k) = P^{-1}(k|k-1)+J_{\phi_\ell}^TR^{-1}(k)J_{\phi_\ell}, \tag{37}$$

which has two difficulties. First of all, computing term $J_{\phi_\ell}^TR^{-1}(k)J_{\phi_\ell}$ using the operators of Eqs. (21) and (22) still has a complexity of $\mathcal{O}(n^2\log(n))$. Second, the same term is not necessarily sparse, so in order to maintain a sparse matrix $P^{-1}(k|k)$, a sparse approximation has to be made. The approximation that is proposed in this paper is to represent the matrix $J_{\phi_\ell}^TR^{-1}(k)J_{\phi_\ell}$ as a low-rank, sparse matrix. By assuming the low-rank property, only a small part of the Jacobian has to be computed. The sparsity property is necessary to obtain a sparse updated matrix that can be efficiently used in the next time step.

The method for computing the low-rank approximation is based on a common procedure in PR algorithms, which is to discard dark pixels in the PSF that are highly corrupted by noise. Since the proposed method is to be used for small-phase aberrations only, it is expected that the PSF will be relatively sharp, and its brightest pixels will be concentrated around its center. Hence, most of the information used for the measurement update is contained in the center pixels, and, consequently, the update $P^{-1}(k|k)$ will mainly depend on the part of the Jacobian corresponding to these center pixels. Explicitly computing the rows of the Jacobian corresponding to the pixels located in the (much smaller) $\bar{p}_c \times \bar{p}_c$ center square has a computational complexity of $\mathcal{O}(n p_c)$, where $p_c := \bar{p}_c^2$.

Finally, it is used so that the desired sparsity pattern of $P^{-1}(k|k)$ is known; therefore, only the elements of $J_{\phi_\ell}^T R^{-1}(k) J_{\phi_\ell}$ that correspond to this sparsity pattern have to be computed. When there are an average of $c_{P^{-1}}$ non-zero elements per row of $P^{-1}(k|k)$, the total computation of this step is $\mathcal{O}(n p_c c_{P^{-1}})$.

C. Error-Covariance Time Update

After computing the time update via Eq. (31), the corresponding inverse error-covariance matrix can be computed via

$$P^{-1}(k+1|k) = Q^{-1} - M_1\big(P^{-1}(k|k) + M_2\big)^{-1} M_1^T, \quad \textbf{(38)}$$

where $M_1 := Q^{-1}A$ and $M_2 := A^T Q^{-1} A$ can be precomputed offline to speed up the computations. Notice that all matrices A, Q^{-1} and $P^{-1}(k|k)$ are sparse banded matrices, but that the inverse $(P^{-1}(k|k) + M_2)^{-1}$ is not sparse in general. However, it is known that since the matrix $P^{-1}(k|k) + M_2$ is a banded and positive definite matrix, its inverse belongs to the class of off-diagonal decaying matrices and can be approximated by a sparse matrix [37].

The computation of this approximation of $P^{-1}(k+1|k)$ is split in two steps. First, the following approximate sparse inverse matrix is computed via

$$\hat{\Gamma} = \arg\min_{\Gamma} \quad \big\|(P^{-1}(k|k) + M_2)\Gamma - I\big\|_F^2$$
$$\text{s.t.} \quad \Gamma \in \mathcal{S}_\Gamma, \quad \textbf{(39)}$$

such that $\hat{\Gamma}$ is a sparse estimate of $(P^{-1}(k|k) + M_2)^{-1}$, and $\mathcal{S}_\Gamma$ is the set of all matrices corresponding to the desired (chosen) sparsity pattern for Γ. The second step is to compute

$$P^{-1}(k+1|k) = Q^{-1} - M_1 \Gamma M_1^T. \quad \textbf{(40)}$$

The sparsity of $P^{-1}(k+1|k)$ (defined by the set $\mathcal{S}_{P^{-1}}$) is known. Hence, when computing Eq. (40), only the entries in $P^{-1}(k+1|k)$ corresponding to this desired sparsity pattern $\mathcal{S}_{P^{-1}}$ have to be computed.

To analyze the computational complexity, the average number of non-zeros per row/column of the matrices Γ, M_1, and M_2 are denoted by c_Γ, c_{M_1}, and c_{M_2}, respectively, and the sets of matrices corresponding to the chosen sparsity patterns of M_1 and M_2 will be defined as $\mathcal{S}_{M_1}$ and $\mathcal{S}_{M_2}$, respectively. For the sake of simplifying the notations, we restrict this analysis to the typical case where all matrices are banded and where $c_\Gamma \geq c_{M_1}$ is chosen such that $\mathcal{S}_{M_1} \subseteq \mathcal{S}_\Gamma$. By exploiting the sparsity structures and only computing the elements of $P^{-1}(k+1|k)$ that are in the sparsity set $\mathcal{S}_{P^{-1}}$, the computational complexity of Eq. (40) will be of the order $\mathcal{O}(n c_{M_1}^2 c_{P^{-1}})$, and computing the matrix $\hat{\Gamma}$ in Eq. (39) will be of the order $\mathcal{O}(n(c_\Gamma^3 + c_{M_2} c_\Gamma^2))$.

6. SIMULATION DESIGN

This section discusses the simulation environment used to provide an illustration of the IKF's performance. For the first-order Taylor approximation in Eq. (20) to hold, the RMS value of the wavefront observed by the focal-plane camera, denoted by RMS($\boldsymbol{\phi}$), should be sufficiently small. Unfortunately, for larger telescope diameters and realistic atmospheric conditions, RMS($\boldsymbol{\phi}$) will be too large; therefore, it is chosen to position the IKF within a closed-loop AO system, as small-phase aberrations are the typical situation within a control loop [12]. In this example, a WFS-based minimum variance estimate (WFS-MVE) of the wavefront is used to drive the controller [38]. The goal of this parallel controller is to compensate for dominant low-order aberrations in the wavefront in order to keep RMS($\boldsymbol{\phi}$) sufficiently low. The open-source Matlab toolbox OOMAO [25] is used to simulate the turbulence and AO system.

The WFS-MVE will also be used as a baseline during the performance evaluations, as it defines the value of RMS($\boldsymbol{\phi}$) to be estimated by the focal-plane wavefront sensing methods. It is important to note that the WFS-MVE should not be interpreted as the optimal performance of WFS-based methods in general, as the Shack–Hartmann (SH) sensor only estimates the low-order aberrations due to its limited number of lenslets.

In order to put the performance of the IKF in further perspective, it is compared to an AP method based on the Gerchberg–Saxton [1] algorithm in a PD framework [35] (referred to as AP-PD). To improve the convergence of the AP-PD algorithm, its estimate at the previous time instance is used as an initial estimate, and all pixels observing less than $y_{\min} = 5$ photons are set to zero to decrease the influence of the measurement noise.

It should be emphasized that the AP-PD method assumes to have two phase diversity images taken simultaneously along the optical axis, while the proposed IKF only uses a single focal-plane image. Although the AP-PD method has this unfair advantage over the IKF method, it will provide an estimate of the expected maximum achievable performance of classical PR methods. Due to its access to two diversity images, the performance of AP-PD should not be significantly affected by the aberration dynamics and therefore serves as a benchmark to see if the IKF can achieve a similar performance under highly dynamic circumstances.

As a measure of performance, the distance $d(k)$ in the pupil plane between the true and estimated GPF will be used in order to avoid issues regarding phase unwrapping:

$$d(k) = \min_{\phi_p \in [0, 2\pi]} \| \boldsymbol{x}(k) - e^{i\phi_p} \hat{\boldsymbol{x}}(k|k) \|_2^2, \quad \textbf{(41)}$$

where $\boldsymbol{x}(k) = \boldsymbol{a} \odot e^{i\boldsymbol{\phi}(k)}$ and $\hat{\boldsymbol{x}}(k|k) = \boldsymbol{a} \odot e^{i\hat{\boldsymbol{\phi}}(k|k)}$ for $\hat{\boldsymbol{\phi}}(k|k)$, as given in Eq. (25). The scalar ϕ_p expresses the piston offset between the real and estimated phase. A normalization of this distance is defined as

$$d_n(k) = \frac{d(k)}{\| \boldsymbol{x}(k) \|_2^2}. \tag{42}$$

Once a time sequence of k_{tot} steps has been obtained, a scalar measure of the accuracy over the full time sequence is computed by taking the RMS: RMS($\boldsymbol{d}_n$), where $\boldsymbol{d}_n \in \mathbb{R}^{k_{\text{tot}}}$ is a vector containing $d_n(k), k = 1, \ldots, k_{\text{tot}}$.

7. RESULTS

This section presents the results of the simulations discussed in the previous section. All parameters have their standard value of Table 3 unless mentioned otherwise.

A. Computational Efficiency

The main result of the paper is the scalable IKF implementation; therefore, the computational efficiency of the algorithm in terms of execution time will be tested for an increasing system size. In this experiment, the telescope diameter D, number of WFS lenslets $\bar{q}$, and resolution of the reconstruction $\bar{n}$ are increased in a way such that their relative ratios are still the same as in Table 3. For each parameter setting, a time sequences of 50 steps is simulated, and the execution time is tracked for each time step. The increase in computational burden due to the higher resolution is visualized in Fig. 2 and compared with the straightforward IKF presentation of Section 4. Comparing the results of the efficient IKF with the reference line of $\mathcal{O}(n)$ confirms the almost linear complexity [$\mathcal{O}(n \log(n))$] of the efficient implementation. The improvement in terms of the scalability of the new algorithm opens up the possibility for a real-time implementation.

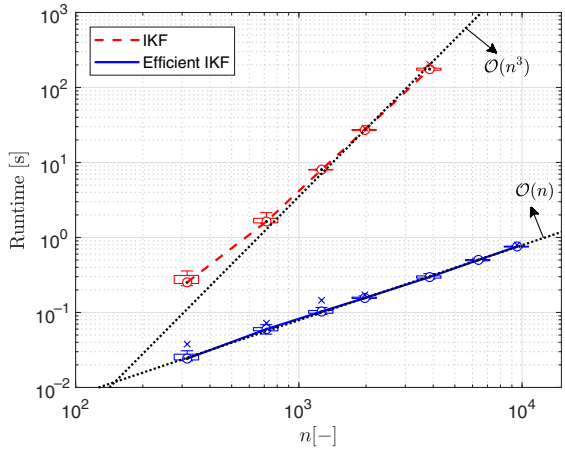

Fig. 2. Computation time per time step for an increasing resolution n. The boxes indicate the 25th and 75th percentiles over the time sequence. Lines are drawn through the medians. The dotted lines present a slope corresponding to a complexity of $\mathcal{O}(n^3)$ and $\mathcal{O}(n)$ for reference purposes. "Efficient IKF" is the new method, "IKF" is the IKF implementation of Section 4.

B. Influence of the Wavefront Dynamics and RMS($\boldsymbol{\phi}$)

Next, the robustness of the method with respect to the wavefront dynamics and value of RMS($\boldsymbol{\phi}$) is investigated by varying v and r_0. For each parameter setting, 10 independent time sequences of 500 steps are simulated in a Monte Carlo simulation. The results in Fig. 3 show that the IKF algorithm performs significantly better for moderate conditions (low v, large r_0). This can be explained by two factors.

Table 3. Simulation Parameters[a]

Parameter	Standard Value
Fried parameter r_0 [m]	0.2
Wind speed v, layer 1 [m/s]	12
Wind speed v, layer 2 [m/s]	16
Source magnitude β_m	8
read-out noise σ_r [photons]	2
Telescope diameter D [m]	1
WFS lenslet grid size $\bar{q}$	6
DM grid size $\bar{q}+1$	7
Wavefront grid size $\bar{n}$	30
Sample frequency f_s [Hz]	500
Outer scale L_0 [m]	15

[a]If not mentioned otherwise, the standard values are used. Turbulence layer 2 is located at an altitude of 5000 m and is moving at an angle of 90° w.r.t. layer 1 located at 0 m. The conversion of magnitude β_m to photon flux is given by Eq. (17). The source is a single natural guide star.

Firstly, since the VAR-1 aberration dynamics model is more accurate at mild conditions, varying v and r_0 will provide information on the limits of the method with respect to the accompanying model. This effect explains the decrease in performance when increasing v, as the wind speed has the largest influence on the accuracy of the model.

Secondly, since the IKF relies on a linear Taylor approximation, it is expected to rely significantly on the magnitude of the residual aberrations, RMS($\boldsymbol{\phi}$). When RMS($\boldsymbol{\phi}$) becomes too large, the Taylor approximation will no longer be valid at this range. Consequently, the performance will rely on the accuracy of the controller, since the accuracy of the WFS-MVE determines the value of RMS($\boldsymbol{\phi}$). The clear dependence on RMS($\boldsymbol{\phi}$) is visualized in Fig. 3(b), which shows corresponding values of RMS($\boldsymbol{\phi}$) for a certain choice of r_0 on the top horizontal axis. This shows that the method is able to track aberrations up to a value of RMS($\boldsymbol{\phi}$) ≈ 1 rad.

Since the wavefront estimates $\hat{\boldsymbol{\phi}}(k|k)$ are compared and not the predictions $\hat{\boldsymbol{\phi}}(k+1|k)$, the AP-PD algorithm is influenced much less by the atmospheric conditions. This is expected, since it uses multiple simultaneous images, and it does not rely on prior information given by a dynamical model or a linear approximation.

The results of Fig. 3(b) are particularly interesting, since it gives useful insights with respect to the estimation of NCPAs. That is, the residual wavefront $\boldsymbol{\phi}$ represents a mismatch between the WFS-based wavefront reconstruction and the corresponding measured PSF image. Hence, this implies that compensating for NCPAs should not be a problem under two assumptions. The first being that the total residual aberration has RMS($\boldsymbol{\phi}$) < 1 rad, and the second that the dynamics of the NCPAs are contained within the dynamical model.

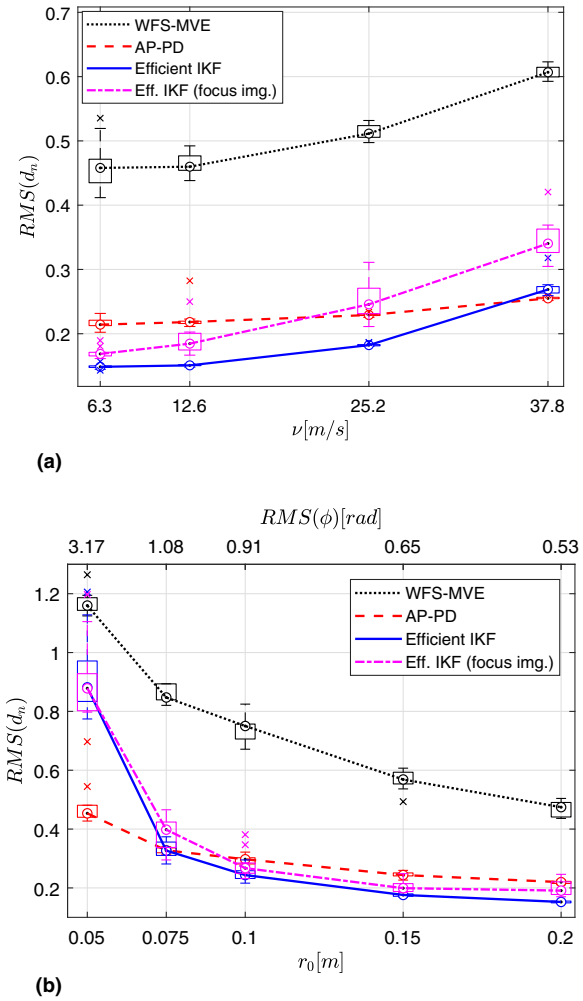

Fig. 3. Results for varying the atmospheric conditions. The boxplots show the RMS of the normalized distance of Eq. (42). The boxes indicate the 25th and 75th percentiles of the results in a Monte Carlo simulation. Lines are drawn through the medians. "Efficient IKF" has a single focal-plane camera out of focus, whereas "Eff. IKF (focus img.)" has a camera placed in focus. The AP-PD method uses two phase diversity images. (a) Results for varying ν. The value shown on the x axis is the speed of the layer 1. The speed of layer 2 is changed accordingly. (b) Results when varying r_0. The top horizontal axis shows the value of RMS(ϕ) for the corresponding r_0.

C. Robustness to Measurement Noise

The robustness to measurement noise was previously addressed in Section 4, where it was argued that the KF algorithm is expected to have an increased performance under noisy conditions without user defined tuning parameters. The influence of the shot noise is now investigated by varying the brightness of the source. Figure 4 shows the performance for an increasing magnitude β_m, i.e., a decreasing brightness in terms of photon flux β_f [see Eq. (17)]. The performance of the IKF deteriorates much less significantly when the magnitude is increased compared to the AP-PD algorithm. This shows that the IKF is more robust for low signal-to-noise ratios, which occur in low brightness conditions, than classical phase diversity methods. Moreover, under low-noise conditions, it shows that the IKF, using only a single focal-plane image, is able to retrieve an estimate that is as accurate as the AP-PD method, which uses two images. The low medians show that even when using a single in focus image, the IKF performs well in most cases, but the performance is much less consistent considering the large spread of the boxplots.

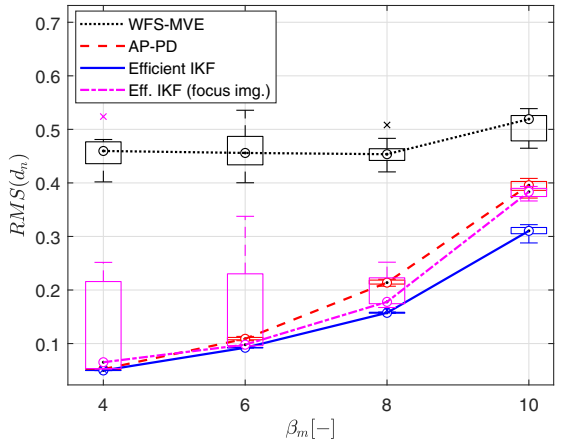

Fig. 4. Results for varying the brightness of the source. The presentation of the boxplots is the same as Fig. 3.

8. CONCLUSIONS

A computationally efficient framework has been proposed in which a single focal-plane image is used to obtain a high-resolution reconstruction of dynamic aberrations. The framework is based on a reformulation of the PR problem for dynamic aberrations into a KF framework using a simple identified model of the dynamics. The computationally efficient implementation scales almost linearly with the number of pixels of the focal-plane camera, making the method suitable for high-resolution AO systems and opening up a real-time implementation as a topic for future research. In a simulation study, the low computational complexity was confirmed, and the accuracy of the method was analyzed under varying conditions. It was discovered that the Kalman filter (using a single focal-plane image) is able to obtain an estimate that is as accurate as phase diversity methods (using two focal-plane images), even when considering highly dynamic aberrations. Moreover, it was shown that the Kalman filter is able to maintain a much better performance than classical phase diversity methods when considering lower signal-to-noise ratios. Finally, although in general measuring the PSF out of focus usually increases the performance, a single in-focus PSF image was able to achieve satisfying results in many cases.

Funding. Seventh Framework Programme (2007-2013); H2020 European Research Council (339681).

Disclosures. The authors declare no conflicts of interest.

REFERENCES

1. R. W. Gerchberg, "A practical algorithm for the determination of phase from image and diffraction plane pictures," Optik **35**, 237–246 (1972).
2. J. R. Fienup, "Reconstruction of an object from the modulus of its Fourier transform," Opt. Lett. **3**, 27–29 (1978).
3. Y. Shechtman, Y. C. Eldar, O. Cohen, H. N. Chapman, J. Miao, and M. Segev, "Phase retrieval with application to optical imaging: a contemporary overview," IEEE Signal Process. Mag. **32**(3), 87–109 (2015).
4. N. Hieu Thao, O. Soloviev, and M. Verhaegen, "Convex combination of alternating projection and Douglas–Rachford operators for phase retrieval," arXiv:2002.01678v1 (2020).
5. E. J. Candes, T. Strohmer, and V. Voroninski, "Phaselift: exact and stable signal recovery from magnitude measurements via convex programming," Commun. Pure Appl. Math. **66**, 1241–1274 (2013).
6. I. Waldspurger, A. d'Aspremont, and S. Mallat, "Phase recovery, maxcut and complex semidefinite programming," Math. Program. **149**, 47–81 (2015).
7. E. J. Candes, X. Li, and M. Soltanolkotabi, "Phase retrieval via Wirtinger flow: theory and algorithms," IEEE Trans. Inf. Theory **61**, 1985–2007 (2015).
8. Y. Chen and E. J. Candès, "Solving random quadratic systems of equations is nearly as easy as solving linear systems," Commun. Pure Appl. Math. **70**, 822–883 (2017).
9. G. Wang, G. B. Giannakis, and Y. C. Eldar, "Solving systems of random quadratic equations via truncated amplitude flow," IEEE Trans. Inf. Theory **64**, 773–794 (2017).
10. C. Ma, X. Liu, and Z. Wen, "Globally convergent Levenberg–Marquardt method for phase retrieval," IEEE Trans. Inf. Theory **65**, 2343–2359 (2018).
11. R. A. Gonsalves, "Small-phase solution to the phase-retrieval problem," Opt. Lett. **26**, 684–685 (2001).
12. C. Smith, R. Marinică, A. Den Dekker, M. Verhaegen, V. Korkiakoski, C. Keller, and N. Doelman, "Iterative linear focal-plane wavefront correction," J. Opt. Soc. Am. A **30**, 2002–2011 (2013).
13. S. Meimon, T. Fusco, and L. M. Mugnier, "Lift: a focal-plane wavefront sensor for real-time low-order sensing on faint sources," Opt. Lett. **35**, 3036–3038 (2010).
14. V. Korkiakoski, C. U. Keller, N. Doelman, M. Kenworthy, G. Otten, and M. Verhaegen, "Fast & furious focal-plane wavefront sensing," Appl. Opt. **53**, 4565–4579 (2014).
15. C. Smith, R. Marinica, and M. Verhaegen, "Real-time wavefront reconstruction from intensity measurements," *Proceedings of the 3rd AO4ELT Conference: Adaptive Optics for Extremely Large Telescopes*, Florence, Italy, May 26–31 (Arcetri Astrophysical Observatory, 2013).
16. R. Doelman, M. Klingspor, A. Hansson, J. Löfberg, and M. Verhaegen, "Identification of the dynamics of time-varying phase aberrations from time histories of the point-spread function," J. Opt. Soc. Am. A **36**, 809–817 (2019).
17. J.-F. Sauvage, T. Fusco, G. Rousset, and C. Petit, "Calibration and precompensation of noncommon path aberrations for extreme adaptive optics," J. Opt. Soc. Am. A **24**, 2334–2346 (2007).
18. M. Wilby, C. Keller, J.-F. Sauvage, T. Fusco, D. Mouillet, J.-L. Beuzit, and K. Dohlen, "A "fast and furious" solution to the low-wind effect for sphere at the VLT," Proc. SPIE **9909**, 99096C (2016).
19. H. Hofer, N. Sredar, H. Queener, C. Li, and J. Porter, "Wavefront sensorless adaptive optics ophthalmoscopy in the human eye," Opt. Express **19**, 14160–14171 (2011).
20. Y. N. Sulai and A. Dubra, "Non-common path aberration correction in an adaptive optics scanning ophthalmoscope," Biomed. Opt. Express **5**, 3059–3073 (2014).
21. L. A. Poyneer and J.-P. Véran, "Kalman filtering to suppress spurious signals in adaptive optics control," J. Opt. Soc. Am. A **27**, A223–A234 (2010).
22. B. M. Welsh, "Fourier-series-based atmospheric phase screen generator for simulating anisoplanatic geometries and temporal evolution," Proc. SPIE **3125**, 327–338 (1997).
23. S. A. Burns, A. E. Elsner, K. A. Sapoznik, R. L. Warner, and T. J. Gast, "Adaptive optics imaging of the human retina," Prog. Retin. Eye Res. **68**, 1–30 (2019).
24. N. Ji, T. R. Sato, and E. Betzig, "Characterization and adaptive optical correction of aberrations during in vivo imaging in the mouse cortex," Proc. Natl. Acad. Sci. USA **109**, 22–27 (2012).
25. R. Conan and C. Correia, "Object-oriented MATLAB adaptive optics toolbox," Proc. SPIE **9148**, 91486C (2014).
26. F. Assémat, R. W. Wilson, and E. Gendron, "Method for simulating infinitely long and non stationary phase screens with optimized memory storage," Opt. Express **14**, 988–999 (2006).
27. M. Verhaegen and V. Verdult, *Filtering and System Identification: A Least Squares Approach* (Cambridge University, 2007).
28. J. Songsiri, J. Dahl, and L. Vandenberghe, "Graphical models of autoregressive processes," in *Convex Optimization in Signal Processing and Communications* (2010), pp. 89–116.
29. C. Yu and M. Verhaegen, "Structured modeling and control of adaptive optics systems," IEEE Trans. Control Syst. Technol. **26**, 664–674 (2017).
30. P. Piscaer, "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.
31. B. L. Ellerbroek, "Efficient computation of minimum-variance wavefront reconstructors with sparse matrix techniques," J. Opt. Soc. Am. A **19**, 1803–1816 (2002).
32. J. Friedman, T. Hastie, and R. Tibshirani, "Sparse inverse covariance estimation with the graphical lasso," Biostatistics **9**, 432–441 (2008).
33. L. M. Mugnier, A. Blanc, and J. Idier, "Phase diversity: a technique for wave-front sensing and for diffraction-limited imaging," Adv. Imaging Electron Phys. **141**, 1–76 (2006).
34. B. M. Bell, "The iterated Kalman smoother as a Gauss–Newton method," SIAM J. Optim. **4**, 626–636 (1994).
35. R. A. Gonsalves, "Phase retrieval and diversity in adaptive optics," Opt. Eng. **21**, 215829 (1982).
36. Y. Saad, *Iterative Methods for Sparse Linear Systems* (SIAM, 2003), Vol. **82**.
37. A. Haber and M. Verhaegen, "Framework to trade optimality for local processing in large-scale wavefront reconstruction problems," Opt. Lett. **41**, 5162–5165 (2016).
38. B. Ellerbroek and T. Rhoadarmer, "Adaptive wavefront control algorithms for closed loop adaptive optics," Math. Comput. Modell. **33**, 145–158 (2001).