

Exploiting the spatiotemporal correlation in adaptive optics using data-driven \mathcal{H}_2 -optimal control

Karel Hinnen and Michel Verhaegen

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

Niek Doelman

TNO Science and Industry, P.O. Box 155, 2600 AD Delft, The Netherlands

Received June 13, 2006; revised December 29, 2006; accepted January 12, 2007;
posted January 29, 2007 (Doc. ID 71874); published May 9, 2007

A recently proposed data-driven \mathcal{H}_2 -optimal control approach is demonstrated on a laboratory setup. Most adaptive optics (AO) systems are based on a control law that neglects the temporal evolution of the wavefront. The proposed control approach is able to exploit the spatiotemporal correlation in the wavefront without assuming any form of decoupling. By analyzing the dynamic behavior of the wavefront sensor (WFS), it is shown that if the wavefront correction device can be considered static, the transfer function from control input to WFS output reduces to a two-tap impulse response and an integer number of samples delay. Considering this model structure, a data-driven identification procedure is developed to estimate the relevant parameters from measurement data. The specific structure allows for an analytical expression of the optimal controller in terms of the system matrices of the minimum-phase spectral factor of the atmospheric disturbance model. The performance of the optimal controller is compared with that of the standard AO control law. An analysis of the dominant error sources shows that optimal control may reduce the temporal error. © 2007 Optical Society of America

OCIS codes: 010.1080, 010.1330.

1. INTRODUCTION

Adaptive optics^{1,2} (AO) is a well-established technique for real-time compensation of the optical wavefront distortions introduced by a turbulent medium. It has found widespread application in ground-based astronomical imaging, where it is used to counteract the devastating effect of atmospheric turbulence on the angular resolution. In this paper we concentrate on the control aspects of AO. Most AO systems are based on a control strategy that is not able to exploit the spatiotemporal correlation in the wavefront. Usually, the control law (see, e.g., Ref. 3) consists of a cascade of a static part, concerned with the problem of finding the actuator inputs that provide the best fit to the wavefront, and a series of parallel feedback loops responsible for stability and closed-loop performance. In the simplest case, the static wavefront reconstruction and fitting step is formulated as a matrix inversion problem. To improve the accuracy, both maximum-likelihood and maximum *a posteriori* techniques have been used to include prior knowledge of the spatial correlation of the wavefront.^{3,4} Prior knowledge of the temporal evolution of the wavefront is usually not included in the control design. Each of the parallel feedback loops typically consists of a first-order lag filter or proportional-integral controller of which the parameters are tuned to make a trade-off among disturbance rejection, noise propagation, and closed-loop stability. To relax the trade-off, modal control optimization has been proposed.^{5,6} In this approach the wavefront is decomposed into a set of modes of which the

corresponding servo gains are optimized independently.

The separation of the control law into static wavefront reconstruction and temporal compensation is based on the assumption that the spatial and temporal dynamics can be decoupled. The Taylor hypothesis,^{7,8} which states that atmospheric turbulence evolves at a time scale that is long compared with the time it takes for the wind-blown inhomogeneities to cross the line of sight, shows that the decoupling assumption will generally fail. When the turbulence can be considered as a frozen layer, there exists a strong correlation between the spatial and temporal dynamics that may be used to the benefit of the controller. By including *a priori* knowledge of the spatiotemporal correlation, wavefront sensor (WFS) measurements from the past and from neighboring channels may be used to predict future wavefront distortions. In this way, it should be possible to reduce the temporal error due to system delays. Also, the sensitivity to measurement noise may be reduced. Since the temporal error is often one of the major error sources in AO,³ advanced control is expected to improve the overall system performance. This may be an improvement in terms of the ability to suppress wavefront distortions or in terms of the limiting magnitude of the guide star.

To exploit the spatiotemporal correlation in the wavefront, we have recently proposed a data-driven \mathcal{H}_2 -optimal control strategy (see Refs. 9 and 10). This consists of a dedicated subspace-identification algorithm that is used to identify a multivariable atmospheric distur-

bance model on the basis of open-loop WFS data. Given a model of the AO system, the identified model is used to compute the controller by formulating the control problem in an \mathcal{H}_2 -optimal control framework. This approach is closely related to the minimum-variance controller design or linear quadratic Gaussian (LQG) formalism used in Refs. 11–15. The main difference is hence not in the framework of analysis, but in the data-driven modeling of the atmosphere and in the way the minimum-variance controller is computed. The subspace-identification algorithm provides an efficient way of identifying an atmospheric disturbance model without assuming any form of decoupling or imposing a restriction on the length of the impulse response. Since the identified disturbance model is sufficiently general to capture the spatiotemporal correlation in the wavefront, the proposed \mathcal{H}_2 -optimal control strategy should be able to benefit from the correlation imposed by, for instance the Taylor hypothesis.

The existing LQG approaches, on the other hand, are based on an atmospheric disturbance model that either assumes modal decoupling¹² or consists only of a first-order autoregressive model.^{11,13–15} A consequence of this restrictive structure is that the disturbance model is able to describe only part of the spatiotemporal correlation in the wavefront. This implies in turn that the LQG controller computed on the basis of such a model is not able to fully exploit the spatiotemporal correlation in the wavefront.

Another advantage of the proposed control approach is that the subspace-identification algorithm directly estimates an atmospheric disturbance model that is minimum-phase with respect to the WFS output. This property can be used to derive an analytical expression for the optimal controller in the case that the transfer function from control input to WFS output consists of a gain matrix and a sample delay. This results in a noniterative way to go from open-loop WFS data to closed-loop controller design. The relation between the minimum-variance controller and the commonly applied control law consisting of a static wavefront reconstructor followed by a series of first-order parallel lag filters has been investigated before. It has been shown that under nearly ideal conditions of no loop and no deformable-mirror dynamics, this type of control law is optimal for isotropic, first-order atmospheric wavefront distortions (see, e.g., Refs. 14 and 15).

The goal of this paper is twofold. First, by analyzing the dynamic behavior of an AO system it will be shown that if the wavefront correction device can be considered static, the transfer function from control input to WFS output can be modeled as a two-tap impulse response and an integer number of samples delay. Considering this model structure is useful, as it allows the analytical expression for the optimal controller to be extended. This will be demonstrated in a forthcoming paper.¹⁶ After deriving the model structure, a data-driven identification procedure is developed to estimate the relevant model parameters from measurement data. Together with the \mathcal{H}_2 -optimal control strategy, this leads to a control design procedure in which all control relevant models, except static optical transformation from phase to slopes, are estimated from data. The second goal of this paper is to demonstrate that

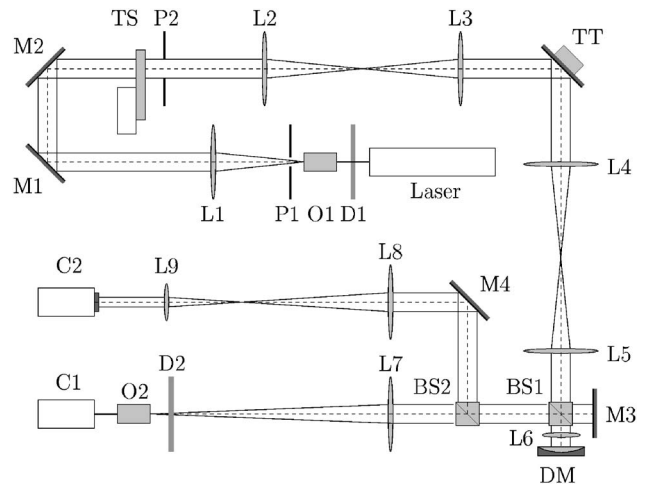


Fig. 1. Schematic representation of AO test bench.

the data-driven \mathcal{H}_2 -optimal control approach is able to reduce the temporal error caused by latencies and bandwidth limitations. The validation is performed on an AO laboratory setup. The remainder of this paper is organized as follows. Section 2 provides a brief description of the AO setup used to validate the proposed optimal control approach, while Section 3 is focused on modeling of the AO system. Section 4 summarizes the main steps of the proposed control approach. The performance of the optimal controller will be compared with that of a conventional AO control law, which will be briefly reviewed in Section 5. This section also provides an overview of the dominant error sources and the criteria used for performance evaluation. The outcome of the experiments is described in Section 6. Section 7 concludes the paper.

2. EXPERIMENTAL SETUP

Figure 1 shows the layout of the AO system used to validate the proposed control approach. In the setup, light from a HeNe laser ($\lambda = 633$ nm) is focused on a $20 \mu\text{m}$ pinhole P1 and is then collimated by the lens L1 to mimic a distant point source. The atmospheric turbulence is simulated by a turbulence simulator TS consisting of a circular plane-parallel glass plate that is rotated through the collimated beam. One side of the glass plate has been machined in such a way that the resulting wavefront distortions approximate a Kolmogorov distribution. The distortions are characterized by a Fried parameter¹ of $r_0 = 2$ mm. With an entrance pupil of $D = 10$ mm, this gives rise to a D/r_0 of 5. By adjusting the rotational speed of the glass plate it is possible to simulate different windspeeds of a single layer of frozen turbulence.

The distorted light is directed to a tip-tilt mirror (TT) that is conjugated to the entrance pupil P2. Separate tip-tilt compensation is important since compensation by the deformable mirror (DM) would demand too much of its dynamic range. Via the beam splitter BS1, the entrance pupil is re-imaged on both the DM and the calibration mirror M3. During normal operation the mirror M3 is shielded—it is used only to calibrate the wavefront sensor (WFS). The DM is a 37-channel electrostatic membrane mirror provided by OKO Technologies.¹⁷ The mirror has a

clear aperture of 15 mm diameter and electrostatic actuators arranged in a hexagonal grid with an interactor spacing of 1.8 mm.

A disadvantage of electrostatic actuation is that the actuators are able to apply only a pulling force on the membrane. To allow bidirectional actuation, a bias is added to the control input. This bias introduces additional focus which is compensated by the negative lens L6. In the control design it is assumed that the transfer function from control input to WFS output can be described by a linear time invariant (LTI) system. The DM movement, however, is proportional to the applied voltage squared and therefore needs to be linearized. This is achieved by taking the DM actuator voltage equal to the square root of the biased control input. This neutralizes the static non-linearity.

The beam splitter BS2 divides the light into a science path and a WFS path. The camera C1 in the science path provides an enlarged image of the point source after wavefront correction. In the WFS path a Shack–Hartmann sensor is used to probe the residual phase errors. This consists of a hexagonal array of 127 microlenses with a focal distance of 15 mm and a pitch of 300 μm mounted on the digital camera C2. The beam size in the WFS path is reduced to 3.3 mm and the WFS is conjugated to the entrance pupil. Figure 2 provides an impression of the sensor–actuator layout of the system. The figure shows a CCD image of the WFS spots in the case that there are no atmospheric wavefront distortions. The crosses on top of the image show the approximate positions of the DM actuators in the WFS image plane. The control computer is a general purpose PC with a 3 GHz

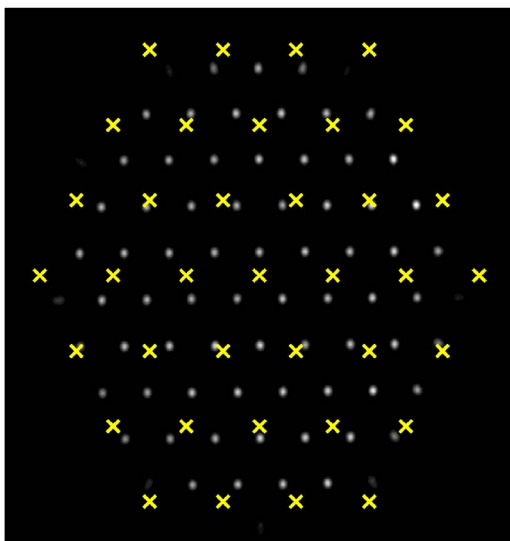


Fig. 2. (Color online) Impression of the sensor–actuator layout.

Intel Pentium IV processor running real-time LINUX. While implementing the control algorithm, the control computer is also responsible for processing the frames from the WFS camera. Using a standard center-of-mass type of algorithm with background compensation and an adjustable threshold level,^{1,18} the control computer has to estimate the deviation of the spots from their nominal position.

A second general-purpose PC is used for simultaneous recording of the images from the cameras C1 and C2. These images are used for performance evaluation. To have a well-established time reference, the frame grabbers of both cameras are triggered by an external pulse generator. Precise sampling and timing of the control action is of utmost importance to demonstrate the optimal control approach. Jitter on the sample frequency causes a mismatch between the expected and actual time of compensation, implying a loss in performance.

3. MODELING THE AO SYSTEM

Computing the \mathcal{H}_2 -optimal controller requires a control-relevant model of both the AO system and the atmospheric distortions. In this section we will consider the problem of modeling the discrete-time transfer function from control input to WFS output. It will be assumed that the wavefront distortions can be represented by a finite-dimensional vector signal $\phi(\cdot) \in \mathbb{R}^{m_\phi}$. Whether the signal $\phi(\cdot)$ provides a zonal or modal representation of the wavefront is irrelevant as long as its mean square error provides a good approximation of the mean square wavefront over the aperture. A similar representation will be used for the phase correction applied by the DM and TT mirror $\phi_m(\cdot)$. Furthermore, we will use the argument of a signal [e.g., $\phi(\cdot)$] to distinguish between its continuous-time [$\phi(t), t \in \mathbb{R}$] and discrete-time counterpart [$\phi(k), k \in \mathbb{N}$]. The sampling time is denoted by T .

A. Active Mirror and WFS Model Structure

In order to derive the model structure of the discrete-time transfer function from control input to WFS output, consider the block scheme of the Shack–Hartmann WFS in Fig. 3. The WFS is not able to directly measure the wavefront phase $\phi(t)$, but provides a discrete-time signal $s(k)$ that is a filtered version of its slope. It will be assumed that the optical transformation from phase $\phi(t)$ to slope $c(t)$ can be modeled by the static mapping $c(t) = G\phi(t)$, with $G \in \mathbb{R}^{m_c \times m_\phi}$ the so-called geometry matrix. The matrix G is determined by the WFS geometry and the basis used to parametrize the phase (see, e.g., Refs. 1 and 2). Since the phase cannot be measured, this relation is the only part of an AO system model that cannot be estimated from data. The shaded block in Fig. 3 models the




Fig. 3. (Color online) Schematic representation of Shack–Hartmann WFS.




Fig. 4. (Color online) Schematic representation of AO system as seen from the controller.

dynamics introduced by the CCD camera. The spot positions on the CCD camera cannot be observed instantaneously but the camera integrates the image over an exposure time $t_e \in \mathbb{R}$. At the end of the integration, the image is read from the CCD camera, which is then again reset to zero. Furthermore, the time required to read and process the frames introduces a time delay $t_d \in \mathbb{R}$. The entire process is repeated with a sampling interval T . With the measurement noise represented by the additive zero-mean white noise term $\eta(k)$, the WFS output at the discrete-time instants $t = kT$, $k \in \mathbb{N}$ can hence be expressed as

$$s(k) = G\phi(k) + \eta(k), \quad (1)$$

where $\phi(k) \doteq \frac{1}{t_e} \int_{kT-t_e}^{kT} \phi(\tau - t_d) d\tau$. Since the delay and the integrating action affect all channels in a similar way, the WFS dynamics are fully decoupled. Furthermore, it is clear that only the part of $\phi(k)$ that is in the row space of G is able to contribute to $s(k)$. This observation can be used to introduce a reduced basis that parametrizes only the informative part of $s(k)$. The main advantage of such a basis is that it reduces the effective number of WFS channels that have to be modeled. Moreover, it improves the numerical conditioning of the identification and control problem by removing the linear dependence between the components of the signal $G\phi(k)$. The reduced basis is obtained by considering the singular value decomposition (SVD),

$$G = U\Sigma V^T = [U_1 \ U_2] \begin{bmatrix} \Sigma_1 & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} V_1^T \\ V_2^T \end{bmatrix}, \quad (2)$$

where U and V are partitioned such that the matrix Σ_1 contains all nonzero singular values. By substituting the SVD into Eq. (1) and exploiting the orthogonality of U , i.e., $U_1 U_1^T + U_2 U_2^T = I$, the signal $s(k)$ can be decomposed as

$$s(k) = U_1 \Sigma_1 \varphi(k) + (U_1 U_1^T + U_2 U_2^T) \eta(k) = U_1 y(k) + U_2 U_2^T \eta(k),$$

where $\varphi(k) \doteq V_1^T \phi(k)$, $y(k) \doteq \Sigma_1 \varphi(k) + n(k)$, and $n(k) \doteq U_2^T \eta(k)$. Since the spaces spanned by U_1 and U_2 are orthogonal and $\eta(k)$ is a zero-mean white noise process uncorrelated to $\phi(k)$, the second term is not related to the turbulence process. For this reason, it is possible to replace the WFS signal $s(k)$ with the lower-dimensional signal $y(k)$ without losing any relevant information. Furthermore, by multiplying Eq. (3) from the left by U_1^T it is clear that the signal $y(k)$ can be simply obtained as $y(k) \doteq U_1^T s(k)$. Since the wavefront cannot be measured directly, the only way to relate $\phi(k)$ to $y(k)$ is via the reduced WFS model

$$y(k) = \Sigma_1 \varphi(k) + n(k). \quad (3)$$

Because of the orthogonality of V , $\phi(k)$ can be decomposed as $\phi(k) = V_1 V_1^T \phi(k) + V_2 V_2^T \phi(k)$. By substituting this in the definition of $\varphi(k)$ it is clear that only the first term, i.e., $V_1 V_1^T \phi(k)$, can be reconstructed from the measurements. Furthermore, since the signals $V_1 V_1^T \phi(k)$ and $\varphi(k)$ have the same 2-norm, the control problem can be reformulated as finding the controller that minimizes the variance of $\varphi(k)$.

In accordance with the above definitions, the reduced representation of the applied phase correction is defined as $\varphi_m(k) \doteq V_1^T \int_{kT-t_e}^{kT} \phi_m(\tau - t_d) d\tau$, while the corresponding residual phase error is $\varepsilon(k) \doteq \varphi(k) - \varphi_m(k)$. Since the WFS is linear in its input, the output corresponding to $\varepsilon(k)$ can be expressed as $r(k) \doteq y(k) - y_m(k)$, where $y(k)$ and $y_m(k)$ denote the contributions due to $\varphi(k)$ and $\varphi_m(k)$, respectively.

Figure 4 provides a schematic representation of the relation between control input $u(k) \in \mathbb{R}^{m_u}$ and WFS output $y_m(k)$. The digital-to-analog (D/A) converter makes the discrete-time signal $u(k)$ into a continuous-time signal $u(t)$, which is the input to the DM and the TT mirror. As a result, $y_m(k)$ can be modeled as the output of the discrete-time system formed by the cascade of the D/A converter, active mirrors, and WFS.

In the proposed \mathcal{H}_2 -optimal control approach there is no separate loop for controlling the TT mirror. Controlling the TT mirror is seen as an integral part of the control design. The DM and TT mirror are described by a single integrated model. The discrete-time transfer function from $u(k)$ to $y_m(k)$ is denoted by $\Sigma_1 \mathcal{H}(z)$, where $\varphi_m(z) = \mathcal{H}(z)u(z)$ describes the mirror dynamics. It has been recently shown¹⁹ that given a continuous-time mirror model, the discrete-time transfer function from control input to WFS output can be computed using the step-invariant transformation.^{20,21} Even though this could be used to derive $\mathcal{H}(z)$, the analysis will be performed in the time-domain as this provides more insight.

Since the active mirrors used in the experimental setup have a time constant that is short compared with the CCD exposure time, they can be considered static, and the only dynamics results from the WFS sampling process. The projected wavefront $V_1^T \phi_m(t)$ can hence be expressed as $V_1^T \phi_m(t) = \bar{H}u(t)$, with $\bar{H} \in \mathbb{R}^{m_y \times m_u}$ a static influence matrix. By substituting this static mirror model in the definition of $\varphi_m(k)$, we can write

$$\varphi_m(k) = \frac{1}{t_e} \int_{kT-t_e}^{kT-t_d} \bar{H}u(\tau) d\tau. \quad (4)$$

The continuous-time actuator input $u(t)$ is obtained from a zero-order hold type of D/A converter operated at the same sample frequency as the camera, i.e.,

$$u(t) \doteq u(k) \quad \text{for } kT \leq t < (k+1)T. \quad (5)$$

Because of the physical limitations of the CCD camera, the exposure time should always be in the range $0 < t_e \leq T$. According to Eq. (4), this implies that $\varphi_m(k)$ depends at most on two past samples of $u(k)$. To elaborate the integral, divide the time delay t_d into an integer number of samples delay $d \in \mathbb{N}$ with a remainder $\tau_d \in \mathbb{R}$ as $t_d = dT - \tau_d$, where $d \geq 1$ and $0 < \tau_d \leq T$. Furthermore, let us assume for the moment that $t_e > \tau_d$. Then by substituting Eq. (5) in Eq. (4) $\varphi_m(k)$ can be expressed as

$$\varphi_m(k) = \bar{H} \left(\frac{1}{t_e} \int_0^{\tau_d} u(k-d) d\tau + \frac{1}{t_e} \int_{T+\tau_d-t_e}^T u(k-d-1) d\tau \right) \quad (6a)$$

$$= \bar{H}(\alpha_1 u(k-d) + \alpha_2 u(k-d-1)), \quad (6b)$$

with $\alpha_1 = \tau_d/t_e$ and $\alpha_2 = (t_e - \tau_d)/t_e$.

A similar analysis can be performed for $t_e \leq \tau_d$; the WFS output $y_m(k)$ can still be expressed as in Eq. (6b), but with $\alpha_1 = 1$ and $\alpha_2 = 0$. By introducing the definitions $H \doteq \alpha_1 \bar{H}$ and $\alpha \doteq \alpha_2/\alpha_1$ it is clear that the discrete-time-transfer function from $u(k)$ to $\varphi_m(k)$ can be expressed as

$$\mathcal{H}(z) = z^{-d}(H + \alpha z^{-1}H). \quad (7)$$

The above model structure with $\alpha \in \mathbb{R}$ will also hold when the DM is not purely static but has a time constant that is short compared with t_e . This is nicely illustrated by the simulation example in Ref. 21. If the time constant is too large, the mirror model substituted into Eq. (4) should be replaced by a dynamic one. Because of this, the WFS output $y_m(k)$ will depend on more than two samples of $u(k)$, and so the required number of finite impulse response (FIR) taps for modeling $g(z)$ will increase.

B. Data-Driven Modeling of the AO System

In this section we will develop a data-driven identification procedure for estimating the model parameters α and H in Eq. (7). For the time being, it will be assumed that both parameters are unknown. This is for instance the case when either τ_d or t_e is not known, or when the DM differs from being purely static. The strategy of estimating the model parameters from measurement data fits perfectly with the philosophy of using data-driven identification. If the parameter α is known in advance, the procedure for estimating H reduces to a standard linear least-squares problem.

The unknown model parameters are estimated on the basis of the WFS response $y_m(k)$ measured by exciting the DM and TT mirror with a zero-mean white noise sequence $u(k)$. Given the measured response $y_m(k)$, the goal is to minimize in a prediction error sense²² the difference between $y_m(k)$ and $\hat{y}_m(k) = \sum_1 \mathcal{H}(z)u(k)$, i.e.,

$$\min_{\mathcal{H}(z)} \sum_{k=d+1}^{N-1} \|y_m(k) - \hat{y}_m(k)\|_2^2. \quad (8)$$

The main difficulty in solving the above optimization problem is its nonconvexity due to the product of α and H in Eq. (7). Hence, for numerical optimization it would be

desirable to have an efficient way of initialization. To this end, we consider the following more general FIR model for $\mathcal{H}(z)$ that results in a convex optimization problem:

$$\mathcal{H}(z) = \sum_{i=p_1}^{p_2} z^{-i} H_{i-p_1+1}, \quad (9)$$

where $p \doteq p_2 - p_1 + 1$ denotes the number of nonzero taps. Indeed, substituting this expression into Eq. (8), via $\hat{y}_m(k)$, and adjusting the lower limit in the summation to $k = p_1 + 1$ results in a linear least-squares problem in the parameters H_i , $i \in \{1, \dots, p\}$. The solution to this problem can be computed analytically as

$$[\hat{H}_p \dots \hat{H}_2 \hat{H}_1] = \Sigma_1^{-1} Y_{p_2, N-1} (U_{0,p, N-p_2})^\dagger, \quad (10)$$

where

$$U_{0,p, N-p_2} \doteq \begin{bmatrix} u_m(0) & \dots & u_m(N-p_2-1) \\ \vdots & \ddots & \vdots \\ u_m(p-1) & \dots & u_m(N-p_1) \end{bmatrix},$$

and $Y_{p_2, N-1} \doteq [y_m(p_2) \dots y_m(N-1)]$. By making a proper choice for p_1 and p_2 , exploratory identification experiments can be used both to obtain an initial estimate of α and to determine delay d . Indeed, if p_1 and p_2 are such that $p_1 \leq d \leq p_2 - 1$, and if the model structure Eq. (7) is correct, the first j coefficient matrices H_1, \dots, H_j have a negligible norm, implying $d = j + p_1$. Furthermore, it follows from the derivation of Eq. (7) that there are two cases to be considered. First, if $t_e \leq \tau_d$, α is zero, which implies that the exploratory experiments should give rise to only one FIR coefficient significantly different from zero. If this is the case, the problem of identifying H reduces to a standard linear least-squares problem, and H can be estimated by setting $p_1 = p_2 = d$ in Eq. (10).

On the other hand, if $t_e > \tau_d$, there should be precisely two FIR coefficient matrices H_{j+1} and H_{j+2} with a norm significantly different from zero. The ratio of these norms, i.e. $\|H_{j+2}\|_2 / \|H_{j+1}\|_2$, provides an estimate of α . The so-computed α is then used to initialize a numerical optimization procedure for solving Eq. (8). Even though this optimization problem is nonconvex, it can be efficiently solved by using separable least-squares.²³ To this end, note that for a fixed value of α , Eq. (8) reduces to a standard least-squares problem in H that has the solution

$$H(\alpha) = \Sigma_1^{-1} Y_{p_2, N-1} [(I \alpha I) U_{0.2, N-d-1}]^\dagger. \quad (11)$$

By substituting this expression in Eq. (8) it is possible to eliminate H , which gives rise to a scalar nonlinear optimization problem over the parameter α only. Having a good initial estimate, this optimization problem can be efficiently solved using a numerical algorithm based on, e.g., Levenberg–Marquardt or Gauss–Newton iterations. Given the optimal value for α , the corresponding H is given by Eq. (11).

4. DATA-DRIVEN OPTIMAL CONTROL

This section provides a brief outline of the data-driven \mathcal{H}_2 -optimal control design strategy validated in this paper. The control approach is based on the work presented

in Refs. 9 and 10 and departs from the generalized plan depicted in Fig. 5. The main components, indicated by the shaded boxes, are the AO system model $\mathcal{H}(z)$ and an atmospheric disturbance model $\mathcal{S}(z)$ describing the uncorrected wavefront distortions $\varphi(k)$ and corresponding WFS output $y(k)$. In this section we will restrict our attention to the special case that the AO system model reduces to a gain matrix and an integer number of samples delay, i.e., $\mathcal{H}(z) = z^{-d}H$, since this is the type of AO system considered in the validation experiments. However, as will be shown in a forthcoming paper,¹⁶ all results can be extended to an AO system of the form of Eq. (7).

After identifying $\mathcal{H}(z)$, the first step of the control design strategy is to determine the disturbance model $\mathcal{S}(z)$. Here, it is assumed that the second-order statistics of the signal $y(k)$ can be described as white noise filtered by a LTI system. Without loss of generality, this shaping filter is in innovation form with respect to $y(k)$. This in combination with Eq. (3) gives rise to the following model structure:

$$\mathcal{S}: \begin{cases} x(k+1) = A_d x(k) + K_d v(k) \\ y(k) = \Sigma_1 C_d x(k) + v(k) \\ \varphi(k) = C_d x(k) + \zeta(k) \end{cases}, \quad (12)$$

where $(A_d - K_d \Sigma_1 C_d) \in \mathbb{R}^{n_d \times n_d}$ and $A_d \in \mathbb{R}^{n_d \times n_d}$ are stable, $v(k)$ is a zero-mean white innovation sequence, and the $\zeta(k) \doteq \Sigma_1^{-1}[v(k) - n(k)]$ is defined in such a way that $\varphi(k)$ is independent from the measurement noise $n(k)$. The system matrices A_d , K_d , and C_d are full and no additional structure is imposed. Considering use of such a model is reasonable if the statistical properties of the wavefront change on a time scale that is long with respect to the time scale of the fluctuations themselves. Indeed, validation experiments on open-loop WFS data from the William Herschel Telescope have shown that a model of this form can be used to predict future wavefront distortions.⁹ The system matrices of the atmospheric disturbance model $\mathcal{S}(z)$ are identified on the basis of open-loop WFS

data $y(k)$. Data-driven modeling has the advantage that it provides a good match with the prevalent turbulence conditions. Moreover, since the disturbance model of Eq. (12) does not assume any form of decoupling, it is sufficiently general to capture the spatiotemporal correlation imposed by a frozen flow. A consequence of this extensive description is that already relatively small AO systems give rise to a sizeable identification problem. For this reason, a dedicated subspace-identification algorithm has been proposed in Ref. 9. Using an efficient implementation this step can be performed on a general purpose PC for AO systems with up to a few hundred WFS channels. For significantly larger systems more efficient algorithms need to be developed.

Given the identified atmospheric disturbance model $\mathcal{S}(z)$ and the transfer function $\mathcal{H}(z)$, the final step of the control design strategy is to compute the optimal controller. The control objective is to find the controller that minimizes the 2-norm of the performance output $e(k)$. As can be easily verified, this is equivalent to finding the controller $C(z)$ that minimizes the cost function

$$J = \mathcal{E}\{\varepsilon^T(k)\varepsilon(k)\} + \mathcal{E}\{u^T(k)Qu(k)\}, \quad (13)$$

where \mathcal{E} denotes the conditional expectation, and the regularization matrix $Q = Q^T \geq 0$ is typically chosen diagonal. The regularization makes a trade-off between the objectives of minimizing the expected residual wavefront error and minimizing the control effort.

By expressing the control problem in the generalized plant framework, the problem of finding the optimal controller $C(z)$ was reduced to a standard \mathcal{H}_2 -optimal control problem (see, e.g., Ref. 20). Computing the \mathcal{H}_2 -optimal controller generally involves the numerical solution of two algebraic Riccati equations. However, because of the minimum-phase property of the atmospheric disturbance model of Eq. (12) and the special structure of the AO system model (i.e., $\mathcal{H} = z^{-d}H$), these Riccati equations can

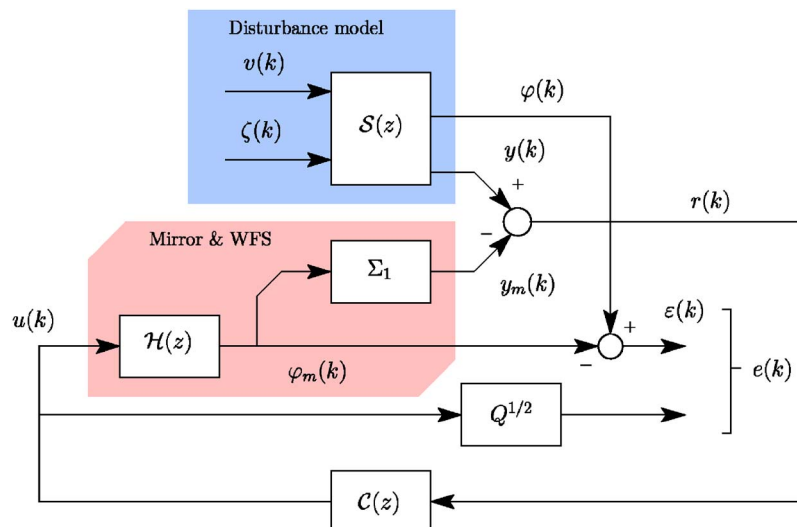


Fig. 5. (Color online) Schematic representation of the closed-loop AO system.

both be avoided. This gives rise to the following state-space expression for the optimal feedback controller $C(z)$:

$$\begin{bmatrix} \xi(k+1) \\ u(k) \end{bmatrix} = \begin{bmatrix} \tilde{A} + z^{-d+1}K_d\Sigma_1HF \\ F(\tilde{A} + z^{-d+1}K_d\Sigma_1HF) \end{bmatrix} \begin{bmatrix} K_d \\ FK_d \end{bmatrix} \begin{bmatrix} \xi(k) \\ r(k) \end{bmatrix}, \quad (14)$$

where $\tilde{A} \doteq A - K_d\Sigma_1C_d$, $F \doteq H_Q^\dagger C_d A^{d-1}$, and $H_Q^\dagger \doteq (H^T H + Q)^{-1} H^T$. That the controller in Eq. (14) is indeed the optimal controller that is minimizing cost function (13) has been proved⁹ for the case that $d=1$ and $Q = \rho I$, with $\rho \in \mathbb{R}$. The proof for different d and Q can be performed in a completely analogous manner.

Having an analytical expression for the optimal controller is useful as this leads to an efficient implementation. Together with the subspace-identification algorithm, it gives rise to a noniterative way to go from open-loop measurement data $y(k)$ to closed-loop controller design. By comparing the state-space equations of the optimal and the atmospheric disturbance model, it can be shown that $u(k)$ can be expressed as

$$u(k) = H_Q^\dagger \hat{\varphi}(k+d|k), \quad (15)$$

where H_Q^\dagger can be interpreted as a regularized inverse of the DM influence matrix H , and $\hat{\varphi}(k+d|k)$ denotes the conditional estimate of $\varphi(k+d)$ given the past closed-loop WFS data $r(j) = y(j) - y_m(j)$, $j \leq k$. From this interpretation it is clear that the optimal controller consists of a part that is concerned with predicting future wavefront distortions and a static mapping that projects the estimated wavefront on the actuator space. Also the common AO control approach, which will be briefly reviewed in Section 5, decomposes into a wavefront reconstruction step followed by a projection on the actuator space. The main difference between both approaches, however, is that in the optimal control approach the static wavefront reconstruction is replaced by dynamic prediction.

5. VALIDATION PROCEDURE AND PERFORMANCE MEASURES

This section focuses on the validation procedure used to demonstrate the benefit of optimal control. Before going into detail on this, it is important to highlight some of the peculiarities of the setup. The setup has an unusual geometry with significantly more microlenses than actuators (see Fig. 2). There are approximately 1.78 microlenses per r_0 , whereas the number of actuators per r_0 is only 1.11. The relatively low number of actuators gives rise to a considerable fitting error. In this paper we attribute to the fitting error any error that is caused by the difference between the estimated required phase correction [i.e., $\hat{\varphi}(k+d|k)$ in the optimal control approach] and the actual phase correction by the mirror. This includes both errors due to the physical inability of DM to take an arbitrary shape and errors caused by imperfections in the projection on the actuator space. Since there are approximately 1.78 microlenses per r_0 the wavefront is still reasonably sampled, which is attractive for performance evaluation. A consequence of this is also that a considerable part of the fitting error can be actually observed.

Another complication of the setup is that the DM has a rather limited actuator range. The DM has a maximum deflection of $9.0 \mu\text{m}$ at the center of the mirror. The deflection at the edges is however much smaller since the mirror membrane is clamped. The available actuator range is insufficient for suppressing the peaks of the distortions. To avoid actuator saturation, we employ the possibility of penalizing the control input. Avoiding actuator saturation is important as it leads to violation of the linearity assumption and may destabilize the control loop. In most AO systems, actuator saturation is not an issue. When it is a problem, increasing the input regularization is usually not the best option as it is overly conservative. For this reason, actuator saturation should ideally be accounted for in the control design. This is, however, beyond the scope of this paper. Applying an input regularization does not need to be a problem in demonstrating the proposed control approach. From Eq. (15) it is clear that increasing the input regularization will lead to a further increase of the wavefront fitting error by perturbing the projection on the actuator space. By a proper design of the validation procedure it is possible to single out the contribution of the fitting error.

The proposed data-driven \mathcal{H}_2 -optimal control approach has been compared with a regularized version of a commonly applied AO control law.^{2,3} For conformity with the rest of the paper, this control law will be reviewed in terms of the reduced signal representation. The applied AO control law decomposes into a static wavefront fitting and reconstruction step and a temporal compensator. Given a WFS measurement $y(k)$, the static part is concerned with finding the actuator input $u(k)$ that would provide the best fit to the wavefront. Let the static relation between $u(k)$ and $y(k)$ be given by $u(k) = R y(k)$ and let the mirrors be modeled as $\varphi_m(k) = H u(k)$. Then, with the WFS model of Eq. (3), the problem of finding the control matrix is formulated as

$$R = \arg \min_R \{ \mathcal{E} \{ [\varphi(k) - H R y(k)]^T [\varphi(k) - H R y(k)] + \mathcal{E} \{ u(k) Q u(k) \} \} \}. \quad (16)$$

In comparison with the usual minimum-variance formulation of the reconstruction problem, Eq. (16) includes a penalty on the control effort. The additional regularization is necessary both to avoid actuator saturation and to enable a fair comparison with the optimal control approach. Under the assumption that wavefront $\varphi(k)$ and the measurement noise $n(k)$ are uncorrelated, R is given by

$$R = \underbrace{(H^T H + Q)^{-1} H^T}_{F} C_\varphi \underbrace{\Sigma_1 (\Sigma_1 C_\varphi \Sigma_1 + C_n)^{-1}}_E, \quad (17)$$

where $C_\varphi \doteq \mathcal{E} \{ \varphi(k) \varphi^T(k) \}$ and $C_n \doteq \mathcal{E} \{ n(k) n^T(k) \}$. In Eq. (17) the control matrix R has been split into two parts. The estimation matrix E provides a minimum-variance estimate of the wavefront $[\hat{\varphi}(k) = E y(k)]$, while F can be interpreted as a projection onto the actuator space. With $F = H_Q^\dagger$, this decomposition shows that the input regularization enables a fair comparison with the optimal control approach if the only dynamics in $\mathcal{H}(z)$ is an integer number of samples delay. Since the AO system operates in closed

loop, the static reconstruction $Rr(k)$ can be interpreted as an estimate of the increment needed to the current actuator commands. In order to ensure closed-loop performance, the temporal compensator has therefore to possess integrating action. The control law used for performance comparison is given by $u(k)=c_1/(1-c_2z^{-1})Rr(k)$, with $c_1 \in \mathbb{R}$ and $c_2 \in \mathbb{R}$ as user-defined control parameters.

In evaluating the reconstruction matrix R , the covariance matrix C_φ is computed theoretically assuming a perfect Kolmogorov spatial distribution. Furthermore, it is assumed that the measurement noise has a covariance matrix of the form $C_n=\sigma_n^2I$, where the variance of the noise σ_n^2 is estimated from open-loop WFS data $y(k)$ for a static distortion. Since the AO system is operated in closed loop, the reconstruction matrix should actually be computed using the closed-loop covariance matrix $C_\varepsilon \doteq \mathcal{E}\{\varepsilon(k)\varepsilon^T(k)\}$ rather than C_φ (see, e.g., Refs. 24 and 15). This modification in wavefront statistics, however, is usually neglected.

Also in this paper we will simply use the open-loop covariance matrix C_φ . Note that since the measurement noise in the setup is rather small, the effect of replacing the closed-loop covariance with the open-loop covariance is almost negligible. The choice of the control parameters c_1 and c_2 will be considered in Section 6.

In order to show that optimal control is indeed able to achieve a performance improvement with respect to the common approach, it is important to have a better insight into the dominant error sources. In the experimental setup, the wavefront error is dominated by the fitting error and the temporal error. From the above discussion it is clear that if the only dynamics in $\mathcal{H}(z)$ is an integer number of samples delay, the common and optimal control approach use the same projection onto the actuator space. This implies that both control approaches give rise to precisely the same fitting error. Any difference in performance is hence caused by a difference in the ability to accurately predict the wavefront at the time of correction. The wavefront prediction accuracy in the validation experiments is determined mainly by the temporal error. Here, the temporal error refers to the error that is caused by the dynamic mismatch between the moment of estimating the wavefront and the actual correction. Both bandwidth limitations and pure time delays may contribute to this error.

To better distinguish between the contributions of the fitting and temporal error it is useful to study their dependence on the sample frequency and the turbulence conditions. Considering the static projection operator H_Q^\dagger , it is clear that for a fixed DM the fitting error depends only on the spatial distribution of the turbulence. Since the Fried parameter r_0 in the setup is fixed, the fitting error can also be considered constant. The temporal error on the other hand depends on the temporal dynamics of the controller and the atmosphere. Under the assumption that the bandwidth of the temporal compensator in the common control approach is proportional to the sample frequency f , the temporal error for Kolmogorov turbulence may be expressed as $\sigma_t^2 = \kappa(f_G/f)^{5/3}$, where $\kappa \in \mathbb{R}$ is a scaling constant and f_G is known as the Greenwood frequency.^{1,25} The Greenwood frequency is a characteris-

tic frequency of the turbulence and for a single frozen layer with wind velocity v it is given by $f_G=0.427(v/r_0)$. Since the fitting error and temporal error are uncorrelated the total wavefront error is obtained by summing the variances, which forms the motivation to consider the error model

$$\sigma_\varepsilon^2 \approx a_0 + a_1(f_G/f)^{a_2}, \quad (18)$$

where a_0 , a_1 , and $a_2 \geq 0$. For the common control law, the constant a_2 should be close to 5/3. It will be assumed that the residual wavefront error obtained with the optimal control approach satisfies the same expression, possibly with a different constant a_2 .

The performance of both controllers has been compared at different Greenwood-to-sample frequency ratios. At each Greenwood-to-sample frequency ratio, the mean square residual wavefront error is estimated on the basis of $N_s=5500$ samples of WFS signal $r(k)$. Given these data, the sample estimate of the mean square residual wavefront error is computed as

$$\hat{\sigma}_\varepsilon^2 = \frac{1}{(N_s - 1)m_\phi} \sum_{k=1}^{N_s} \hat{\varepsilon}(k)^T \hat{\varepsilon}(k), \quad (19)$$

with $\hat{\varepsilon}(k) = \Sigma_1^{-1}r(k)$. To verify if the error model (18) indeed provides a good description of the residual wavefront error, it has been fitted to the observed values of $\hat{\sigma}_\varepsilon^2$. For a fixed a_0 , the problem of estimating the coefficients a_1 and a_2 boils down to fitting an exponential relation $a_1(f/f_G)^{a_2}$ to the measurements $(\hat{\sigma}_\varepsilon^2 - a_0)$. Such a fitting problem is conveniently solved on a logarithmic scale, as this renders the error model linear in the unknowns $\log_{10}(a_1)$ and a_2 . This forms the motivation for defining the following least-squares problem to estimate a_0 , a_1 , and a_2 :

$$\min_{a_1, a_2} \|\log_{10}(\bar{\sigma}_\varepsilon^2 - a_0) - \mathbf{1} \log_{10}(a_1) - a_2 \log_{10}(\bar{f})\|_2^2, \quad (20)$$

where $\bar{\sigma}_\varepsilon^2$ and \bar{f} are the vectors obtained by stacking the different observations of $\hat{\sigma}_\varepsilon^2$ and the corresponding f_G/f ratios, respectively, and $\mathbf{1}$ is a vector of the same dimension as $\bar{\sigma}_\varepsilon^2$ and \bar{f} with all elements equal to 1. Since the above optimization problem is a linear least-squares problem for fixed a_0 , separable least-squares can be used to transform into a single parameter optimization problem.

Besides $\hat{\sigma}_\varepsilon^2$, a number of other criteria have been used for characterizing the performance. Since the \mathcal{H}_2 -optimal controller is designed to minimize the cost function (13), an obvious choice is to look at the relative improvement of this function. Let \hat{J}_c and \hat{J}_o denote the sample estimates of the cost function defined in accordance with Eq. (19). Then the relative improvement is computed as \hat{J}_c/\hat{J}_o . Furthermore, it is interesting to have performance measures that directly relate to the quality of the corrected image. For this reason, the Strehl ratio and the normalized encircled energy have also been computed. Both of these performance measures are derived from the long-exposure image mimicked by averaging $N_f=250$ frames of the science camera. The averaged image is background compensated to account for the CCD dark pattern. Since the science camera has a fixed exposure time of 5 ms, the total recording time in each experiment is constant.

The Strehl ratio is the most commonly used performance metric in AO.^{1,2} Let $\bar{I}(p)$ denote the background-compensated, long-exposure image, with $p \in \mathbb{N}^2$ the pixel coordinate in the CCD frame, and let $\mathcal{B}_r(p_0) = \{p \in \mathbb{N}^2 \mid \|p - p_0\|_2 < r\}$ denote all pixels $p \in \mathbb{N}^2$ in a circle with radius r around the point $p_0 \in \mathbb{R}^2$. Further, let q_0 denote the radius of the theoretical diffraction-limited spot computed on the basis of the aperture size. Then the first step in estimating the Strehl ratio is to extract a neighborhood with a radius $0.2q_0$ around the pixel $p_m \in \mathbb{N}$ with maximum intensity. This neighborhood is used to obtain refined estimates of the peak intensity I_q and position p_q by fitting a quadratic form to the measured intensity $\bar{I}(p)$, $p \in \mathcal{B}_{0.2q_0}(p_m)$. With the refined estimate of the peak position a larger neighborhood $\mathcal{B}_{2p_0}(p_q)$ is extracted over which the total flux is computed as $I_t = \sum_{p \in \mathcal{B}_{2q_0}(p_q)} \bar{I}(p)$. Also the diffraction limited total flux I_d on this neighborhood is computed from the theoretical diffraction pattern with unit peak intensity by oversampling the pixels by a factor of 8.

In the final step, the Strehl ratio is computed by weighting the estimated peak intensity by the ratio of the measured and the theoretically computed total flux, i.e., $S = (I_q/I_t)I_d$. Using the same notation, the normalized encircled energy is defined as $E(r) = \max_{p_0} (\sum_{p \in \mathcal{B}_r(p_0)} \bar{I}(p) / \sum_p \bar{I}(p))$. This provides a measure of the fraction of the total incident energy that is contained in the central core of the spot. The better the wavefront correction, the more energy is concentrated in the central core, and the faster $E(r)$ increases with r .

6. EXPERIMENTAL RESULTS

Both the \mathcal{H}_2 -optimal control approach and the common AO control law have been implemented on the experimental setup. After aligning the WFS, only 69 of the 127 microlenses are illuminated sufficiently to be used for wavefront sensing. This implies that the unreduced WFS signal consists of $m_s = 138$ channels. The geometry matrix G specifying the relation between slope measurements and phase is defined by adapting the well-known Fried configuration for a hexagonal grid. Just as for a rectangular grid, the position of the phase points is determined by shifting the hexagonal grid over half the pitch size. The reduced WFS signal $y(k)$, obtained by projecting out the modes that cannot be related to the wavefront, consists of $m_y = 88$ channels. During the experiments, the WFS exposure time was adjusted to 5 ms. The WFS gain has been calibrated using the TT mirror.

The exploratory identification experiments in the data-driven identification procedure for estimating the model parameters α and H in $\mathcal{H}(z)$ show that there is only one FIR coefficient that differs significantly from zero. This is in perfect agreement with the observation that the real-time software is implemented in such a way that, for the given exposure time and sample frequencies, the time delay between control input and WFS output is always 2 samples, i.e., $d=2$ and $\tau_d=T$. From Subsection 3.A it is clear that, for $t_e \leq \tau_d$ and $\alpha=0$, $\mathcal{H}(z)$ reduces to $\mathcal{H}(z) = z^{-2}H$. An even more important argument for considering this reduced model structure is the close agreement between measured and predicted WFS output. Validation

experiments on a data set of $N=1000$ samples show a mean variance accounted for²⁶ of more than 99.2% over the channels. This implies that more than 99.2% of the variance $y(k)$ can be explained by the identified model. The identified transfer function therefore provides an accurate description of the true system.

Both control algorithms have been tested for Greenwood-to-sample frequency ratios in the range $f_G/f \in [0.015, 0.43]$. Realistic values for the Greenwood frequency can range tens to hundreds of Hertz (see, e.g., Ref. 1). For an AO system with a sample frequency of $f = 500$ Hz, this gives rise to Greenwood-to-sample frequency ratios that are roughly in the range $f_G/f \in [0.04, 0.4]$. Most AO systems, however, will probably be designed to operate at the lower end of this range. The upper bound of $f_G/f=0.4$ especially is rather large as it implies a characteristic turbulence frequency close to the Nyquist frequency. The reason for considering these high Greenwood-to-sample frequency ratios is that the fitting error in the experimental setup is relatively large. As a result of this, it takes higher f_G/f ratios for the temporal error to become the dominant factor. Furthermore, comparing both control algorithms over a wide range of Greenwood-to-sample frequency ratios is interesting as it provides some insight into the possible increase in Greenwood-to-sample frequency ratio while maintaining a certain performance level.

To verify if the residual wavefront error indeed depends only on the ratio of f_G and f , and not on their respective values, the f_G/f ratios have been obtained by considering different Greenwood and sample frequency combinations. Some of these combinations give rise to the same or approximately the same f/f_G ratio. The sample rates used to obtain the different f/f_G ratios include $f \in \{4.44, 6.44, 8.33, 10.41, 12.5, 14.29, 16.67, 20.0\}$ Hz.

At each Greenwood-to-sample frequency ratio, the performance of the common control approach is determined using the same value for c_1 and c_2 . The control parameters were tuned to minimize $\hat{\sigma}_e^2$ at a f_G/f ratio of 0.043. This resulted in $c_1=0.48$ and $c_2=0.98$. These parameters give rise to a reasonable performance over the entire Greenwood-to-sample frequency range without loss of stability. In fact, the values are close to the ones found by repeating the tuning procedure at a number of test points over the considered range. The reason for this weak dependence on the Greenwood-to-sample frequency ratio is that the setup uses a bright light source so that the measurement noise contribution is rather small. As a result, the problem of finding the optimal control parameters turns into the one of maximizing the control bandwidth without losing stability.

Since the only dynamic in the AO system is a two-sample delay, i.e., $\mathcal{H}(z) = z^{-2}H$, the control parameters for the system is still stable are independent of the Greenwood-to-sample frequency ratio. For a fair comparison, the same input regularization matrix Q has been used in both control approaches. The regularization matrix is chosen diagonal, i.e., $Q = \text{diag}\{q_1, q_2, \dots, q_{m_y}\}$, and the control effort weighting q_i on each of the actuators is tuned to avoid actuator saturation and to ensure that the DM stays within a linear range. This results in an additional weight on the actuators with a small dynamic

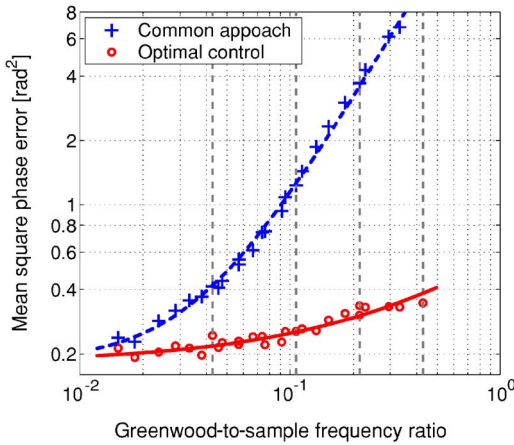


Fig. 6. (Color online) Mean square error as a function of the Greenwood-to-sample frequency ratio.

range, like the actuators near the edge of the DM. The TT mirror, on the other hand, does not require regularization. In the optimal control approach, the atmospheric disturbance model $S(z)$ is identified on the basis of $N_s = 5500$ samples of open-loop WFS data $y(k)$. The model order, i.e., the dimension of the state vector in all validation experiments is chosen equal to $n_d = 256$. The subspace-identification algorithm is able to identify atmospheric disturbance models of higher order (e.g., $n_d = 300$, $n_d = 350$) but this does not lead to any performance improvement.

The estimated mean square residual phase error $\hat{\sigma}_\varepsilon^2$ obtained in the different experiments is depicted in Fig. 6. Each circle and each cross is the result of an experiment at a specified Greenwood-to-sample frequency ratio using the optimal and common control approach, respectively. As expected from the discussion in Section 5, the performance improvement increases with Greenwood-to-sample frequency ratio. Whereas the improvement is rather low at low Greenwood-to-sample frequency ratios, a considerable improvement is observed at high ratios. The dashed curve in Fig. 6 represents the fit of the error model of Eq. (18) to the residual phase error $\hat{\sigma}_\varepsilon^2$ obtained with the common control approach. The corresponding fit for the optimal control approach is depicted by the solid curve. Since the residual phase error for the optimal control approach shows a rather weak dependence on the Greenwood-to-sample frequency ratio, the estimate of the parameters a_1 and a_2 is quite sensitive to a slight variation in a_0 . For this reason, optimization over a_0 is omitted and its value is fixed to the value of a_0 found for the common control approach. This is a reasonable assumption, since the fitting error in both approaches should be the same. The fitted error models obtained in this way are given by

$$\hat{\sigma}_\varepsilon^2 \approx 0.1833 + 44.21(f_G/f)^{1.6616},$$

$$\hat{\sigma}_\varepsilon^2 \approx 0.1833 + 0.3853(f_G/f)^{0.7642},$$

for the common and the optimal approach, respectively.

Figure 6 shows that the above relations provide an accurate fit to the mean square residual phase error observed in the experiments. Furthermore, the fitted exponent a_2 for the common control approach is in close

agreement with the theoretical value $5/3 (\approx 1.667)$. These observations support the error model (18) and show that optimal control is indeed effective in reducing the temporal error. Since at low Greenwood-to-sample frequency ratios the fitting error becomes the limiting factor, only little can be gained by optimal control in this regime.

To prove that the performance at low Greenwood-to-sample frequency ratios is indeed limited by the fitting error, it is useful to estimate this error on the basis of the available open-loop data $y(k)$. This is achieved by first reconstructing the uncorrected wavefront as $\hat{\varphi}(k) = \Sigma_1^{-1} y(k)$. Since the operator H_Q^\dagger can be interpreted as the projection of $\hat{\varphi}(k)$ on the actuator space, the actuator commands can be computed as $\hat{u}(k) = H_Q^\dagger u(k)$. The computed actuator commands are then used to determine the wavefront correction $\hat{\varphi}_m(k) = H \hat{u}(k)$. Neglecting all dynamics, the fitting error can now be estimated as the mean square error of $\hat{\varepsilon}(k) = \hat{\varphi}(k) - \hat{\varphi}_m(k)$. To demonstrate the effect of the regularization, $\hat{u}(k)$ has been computed both with and without input regularization. Furthermore, the effect of actuator saturation has been investigated by chopping off the signals that are out of range.

The averaged estimated mean square fitting error and the standard deviation over the different data sets are shown in Table 1. The table shows that when accounting for both the input regularization and the actuator saturation the estimated fitting error is in close agreement with the constant a_0 obtained from the error model. The fitting error estimated on the basis of the open-loop WFS data is within 2 standard deviations of the estimated value of a_0 . Furthermore, Table 1 shows that actuator saturation almost doubles the observed fitting error. The additional increase of the fitting error by accounting for both regularization and saturation is rather small since both error contributions are strongly correlated, as the regularization parameter Q has been tuned to avoid actuator saturation.

As pointed out in Section 5, the relative improvement in cost function and the Strehl ratio have been used as additional measures of performance. They have been evaluated at the Greenwood-to-sample frequency ratios indicated by the vertical dashed lines in Fig. 6 and are summarized in Table 2. The results in this table are consistent with the previous observations. They show a considerable performance improvement for large Greenwood-to-sample frequency ratios, while the performance improvement at low Greenwood-to-sample frequency ratios is rather modest. The relative improvement in cost function is especially large. This is to be expected as the \mathcal{H}_2 -optimal controller is designed to minimize this criterion. Moreover, this suggests that the performance gain may be much larger if no input regularization is needed.

Table 2 shows that optimal control is also able to achieve a performance improvement with respect to the

Table 1. Estimate of DM Fitting Error

Regularization	Saturation	$\hat{\sigma}_f^2$	Std
no	no	0.087	0.003
no	yes	0.170	0.009
yes	yes	0.197	0.007

Table 2. Performance Comparison of Common and Optimal Control Approach

f_G/f ($\times 10^{-1}$)	Strehl Common	Strehl Optimal	\hat{J}_c/\hat{J}_0 Relative
0.43	0.79	0.85	1.70
1.1	0.57	0.85	4.66
2.1	0.20	0.81	11.7
4.3	0.12	0.78	19.3




Fig. 7. (Color online) Normalized encircled energy at different Greenwood-to-sample frequency ratios.

Strehl ratio. Especially at large Greenwood-to-sample frequency ratios, optimal control gives rise to a considerably higher Strehl ratio. Figure 7 shows the normalized encircled energy achieved with both the optimal and the common AO control approach for different Greenwood-to-sample frequency ratios. Also, this figure shows that optimal control is able to improve the science image and that the gain in performance increases with the Greenwood-to-sample frequency ratio. Whereas the encircled energy curves obtained with the common control approach indicate a rather spread out intensity pattern, optimal control gives rise to a science image in which a larger fraction of the incident energy is concentrated in the central core.

Finally, note that optimal control significantly reduces the sensitivity of the Strehl ratio and the normalized energy to variation in the Greenwood-to-sample frequency ratio. This is in accordance with the observation that the mean square residual phase error shows a weaker dependence on the Greenwood-to-sample frequency ratio for the optimal than for the common control approach. Optimal control is therefore also attractive from the viewpoint of performance robustness.

7. CONCLUSIONS

In this paper, we have demonstrated a recently proposed data-driven \mathcal{H}_2 -optimal control approach on an experimental setup. In contrast to existing AO control approaches, this approach does not assume any form of decoupling and has the potential to exploit the spatiotemporal correlation imposed by the Taylor hypothesis. In the first step a dedicated subspace-identification algorithm is used to identify a multivariable atmospheric

disturbance model from open-loop WFS data. Data-driven modeling has the advantage that it provides a good match with the prevalent turbulence conditions.

A second ingredient necessary for computing the optimal controller is the transfer function from control inputs to WFS outputs. By analyzing the dynamic behavior of the WFS, it has been shown that if the wavefront correction device can be considered to be static, the scalar dynamics of the discrete-time transfer function from actuator inputs to WFS output can always be modeled as an integer number of samples delay followed by a two-tap impulse response. This observation is particularly useful as, together with the minimum-phase property of the identified atmospheric disturbance model, it allows the \mathcal{H}_2 -optimal controller to be computed analytically as will be shown in a forthcoming paper. A data-driven identification approach has been developed to identify a transfer function of the desired structure from measurement data.

The data-driven \mathcal{H}_2 -optimal control approach has been validated in an experimental setting by comparing it with a commonly applied AO control law. In this comparison different performance criteria based on measurements from both the science camera and the WFS have been used. The considered performance criteria include an estimate of the mean square residual wavefront error, the reduction in cost function, the Strehl ratio, and the normalized encircled energy. The experiments show that optimal control is able to achieve a performance improvement with respect to each of these criteria, and that the gain in performance increases with the Greenwood-to-sample frequency ratio.

A careful analysis of the dominant error sources has shown that the improved performance can be attributed to a reduction in the temporal error. Since the temporal error is an exponentially increasing function of the Greenwood-to-sample frequency ratio, this also explains that a lot more can be gained at higher ratios than at lower ratios where the fitting error becomes dominant. Optimal control is able to reduce the temporal error by exploiting the spatiotemporal correlation in the wavefront. For a single layer of turbulence it is conceptually clear that the spatiotemporal correlation can be used to improve the predictability of the wavefront. It is however important to note that the proposed control strategy does not depend on this assumption. It is therefore to be expected that comparable or slightly lower performances can be achieved in the multilayer case provided that the WFS has a sufficiently high spatial and temporal resolution to resolve the contributions of the individual layers. The influence of multiple layers on the overall performance remains a point of further investigation.

Another issue that deserves further attention is the effect of measurement noise. Since exploiting the spatiotemporal correlation is expected to reduce the effect of measurement noise, optimal control may also be beneficial in low signal-to-noise situations. This has already been shown indirectly by the fact the optimal control is able to achieve the same performance at much lower sample frequencies, enabling a longer integration time.

Finally it is important to note that a possible disadvantage of the proposed data-driven optimal control approach is that it cannot actively control unseen modes. Unseen

modes are wavefront modes that are in the null space of the geometry matrix G . Since these modes do not influence the WFS output, they can never be modeled using data identification. This does not necessarily imply that these modes are unobservable when starting from an atmospheric disturbance model that directly models the phase. The possibility of suppressing unseen modes has been demonstrated in Ref. 13. From this perspective it may be rewarding to search for methods that can combine the advantages of data-driven identification with the ability to suppress unseen modes.

ACKNOWLEDGMENTS

This research has been conducted in the framework of the Knowledge Center for Aperture Synthesis (KAS). The Knowledge Center is an initiative of TNO Science and Industry to develop fundamental and advanced technologies for optical aperture synthesis. The knowledge center is a long-term cooperation of TNO and (primarily) the Delft University of Technology.

Corresponding author K. Hinnen can be reached by phone, +31(0)6-16328935; fax, +31(0)15-2786679; e-mail, k.j.g.hinnen@planet.nl.

REFERENCES

1. J. W. Hardy, *Adaptive Optics for Astronomical Telescopes* (Oxford U. Press, 1998).
2. F. Roddier, *Adaptive Optics in Astronomy* (Cambridge U. Press, 1999).
3. M. A. van Dam, D. L. Mignant, and B. A. Macintosh, "Performance of Keck Observatory adaptive optics system," *Appl. Opt.* **43**, 5458–5467 (2004).
4. N. F. Law and R. G. Lane, "Wavefront estimation at low light levels," *Opt. Commun.* **126**, 19–24 (1996).
5. B. L. Ellerbroek, "Optimizing closed-loop adaptive-optics performance with use of multiple control bandwidths," *J. Opt. Soc. Am. A* **11**, 2871–2886 (1994).
6. E. Gendron and P. Léna, "Astronomical adaptive optics," *Astron. Astrophys. Arch. Ration. Mech. Anal.* **291**, 337–347 (1994).
7. G. I. Taylor, "The spectrum of turbulence," *Proc. R. Soc. London, Ser. A* **164**, 476–490 (1938).
8. E. Gendron and P. Léna, "Single layer atmospheric turbulence demonstrated by adaptive optics observations," *Astrophys. Space Sci.* **239**, 221–228 (1996).
9. K. J. G. Hinnen, M. Verhaegen, and N. J. Doelman, " \mathcal{H}_2 -optimal control of an adaptive optics system: Part I, Data-driven modeling of the wavefront disturbance," *Proc. SPIE* **5203** 75–85 (2005).
10. K. J. G. Hinnen, N. J. Doelman, and M. Verhaegen, " \mathcal{H}_2 -optimal control of an adaptive optics system: part II, closed-loop controller design," *Proc. SPIE* **5903** 86–99 (2005).
11. R. N. Paschall and D. J. Anderson, "Linear quadratic Gaussian control of a deformable mirror adaptive optics system with time-delayed measurements," *Appl. Opt.* **32**, 6347–6358 (1993).
12. D. P. Looze, M. Kasper, S. Hippler, O. Beker, and R. Weiss, "Optimal compensation and implementation for adaptive optics systems," *Exp. Astron.* **15**, 67–88 (2003).
13. B. Le Roux, J.-M. Conan, C. Kulcsár, H.-F. Raynaud, L. M. Mugnier, and T. Fusco, "Optimal control law for classical and multiconjugate adaptive optics," *J. Opt. Soc. Am. A* **21**, 1261–1276 (2004).
14. C. K. Kulcsár, H. F. Raynaud, C. Petit, J. M. Conan, and P. Viaris de Lesegno, "Optimal control, observers and integrators in adaptive optics," *Opt. Express* **14**, 7464–7476 (2006).
15. D. P. Looze, "Minimum variance control structure for adaptive optics systems," *J. Opt. Soc. Am. A* **23**, 603–612 (2006).
16. K. J. G. Hinnen, M. Verhaegen, and N. J. Doelman, "A Data-Driven \mathcal{H}_2 -Optimal Control Approach for Adaptive Optics," (2006). Submitted to *IEEE Transactions*.
17. G. Vdovin and P. M. Sarro, "Flexible mirror micromachined in silicon," *Appl. Opt.* **34**, 2969–2972 (1995).
18. S. Thomas, "Optimized centroid computing in a Shack–Hartmann sensor," *Proc. SPIE* **5490**, 1238–1246 (2004).
19. K. J. Aström and B. Wittenmark, *Computer-Controlled Systems, Theory and Design* (Prentice-Hall, 1997).
20. T. Chen and B. Francis, *Optimal Sampled-Data Control Systems* (Springer-Verlag, 1995).
21. D. P. Looze, "Realization of systems with CCD-based measurements," *Automatica* **41**, 2005–2009 (2005).
22. L. Ljung, *System Identification: Theory for the User* (Prentice Hall, 1999).
23. G. H. Golub and V. Pereyra, "The differentiation of pseudo-inverses and nonlinear least squares problems whose variables separate," *SIAM (Soc. Ind. Appl. Math.) J. Numer. Anal.* **10**, 413–432 (1973).
24. M. Kasper, "Optimization of an adaptive optics system and its application to high-resolution imaging spectroscopy of T Tauri," Ph.D. thesis (University of Heidelberg, 2000).
25. D. L. Fried, "Time-delay-induced mean-square error in adaptive optics," *J. Opt. Soc. Am. A* **7**, 1224–1225 (1990).
26. V. Verdult and M. Verhaegen, "Identification of multivariable bilinear state space systems based on subspace techniques and separable least squares optimization," *Int. J. Control* **74**, 1824–1836 (2001).