

High-order coronagraphic wavefront control with algorithmic differentiation: first experimental demonstration

Scott D. Will^{a,b,c,*} Marshall D. Perrin^{b,c} Emiel H. Por^c James Noss^c
Ananya Sahoo^c Peter Petrone^{c,d} Iva Laginja^{b,e} Raphaël Pourcelot^c
Susan F. Redmond^{b,f} Laurent Pueyo^c Tyler D. Groff^{b,a}
James R. Fienup^b and Remi Soummer^c

^aNASA Goddard Space Flight Center, Greenbelt, Maryland, United States

^bThe Institute of Optics, University of Rochester, Rochester, New York, United States

^cSpace Telescope Science Institute, Baltimore, Maryland, United States

^dSigma Space, Lanham, Maryland, United States

^eLESIA, Université PSL, Sorbonne Université, Université Paris Cité, Observatoire de Paris, CNRS, Meudon, France

^fPrinceton University, Department of Mechanical and Aerospace Engineering, Princeton, New Jersey, United States

ABSTRACT. Future space-based coronagraphs will rely critically on focal-plane wavefront sensing and control with deformable mirrors (DMs) to reach deep contrast by mitigating optical aberrations in the primary beam path. Until now, most focal-plane wavefront control algorithms have been formulated in terms of Jacobian matrices, which encode the predicted effect of each DM actuator on the focal-plane electric field. A disadvantage of these methods is that Jacobian matrices can be cumbersome to compute and manipulate, particularly when the number of DM actuators is large. Recently, we proposed a new class of focal-plane wavefront control algorithms that utilize gradient-based optimization with algorithmic differentiation to compute wavefront control solutions while avoiding the explicit computation and manipulation of Jacobian matrices entirely. In simulations using a coronagraph design for the proposed Large UV/Optical/Infrared Surveyor, we showed that our approach reduces overall CPU time and memory consumption compared to a Jacobian-based algorithm. Here, we expand on these results by implementing the proposed algorithm on the High-contrast Imager for Complex Aperture Telescopes tested at the Space Telescope Science Institute and present initial experimental results, demonstrating contrast suppression capabilities equivalent to Jacobian-based methods.

© The Authors. Published by SPIE under a Creative Commons Attribution 4.0 International License. Distribution or reproduction of this work in whole or in part requires full attribution of the original publication, including its DOI. [DOI: [10.1117/1.JATIS.9.4.045004](https://doi.org/10.1117/1.JATIS.9.4.045004)]

Keywords: coronagraphy; high-contrast imaging; wavefront control; optimization

Paper 23098G received Aug. 18, 2023; revised Oct. 26, 2023; accepted Nov. 27, 2023; published Dec. 19, 2023.

1 Introduction

Future space coronagraphs attempting to image and characterize Earth-like planets around nearby solar-type stars will rely critically on closed-loop wavefront sensing and control (WFS&C) using deformable mirrors (DMs) to mitigate optical aberrations in the primary beam path. These aberrations, primarily mid-spatial frequency wavefront errors and optical misalignments in the telescope and coronagraph optics, give rise to a speckle floor that is coherent with

*Address all correspondence to Scott D. Will, scott.d.will@nasa.gov

the star and evolves slowly over time in response to miniscule drifts in the thermal and mechanical state of the observatory. If uncorrected, the speckle floor overwhelms the faint image of the orbiting planet, which is expected to be 10^{10} times fainter than the host star at 0.1 arcseconds of separation or less at visible wavelengths.¹

The current state-of-the-art wavefront control algorithms, stroke minimization (SM)² and electric field conjugation (EFC),³ compute DM command solutions using a first-order approximation of the focal-plane electric field in an optimal control framework. In both cases, the optimal DM update is written down in closed form as the solution of a linear system of equations constructed from a Jacobian matrix that describes the influence of each DM actuator on the focal-plane electric field.

Until recently, little attention has been paid to the computational demands of SM and EFC. The complexity of SM and EFC is dominated by the cost of computing and manipulating the Jacobian matrix, whose size is proportional to the product of the number of DM actuators and number of dark-zone pixels. In broadband imaging, this is compounded by the requirement to compute a separate Jacobian for each controlled wavelength. The Jacobian is most often model-based, in which case an optical diffraction model of the coronagraph is evaluated repeatedly to predict the focal-plane influence of each actuator. Moreover, the Jacobian is a linearization of the true, nonlinear behavior of the DMs and must be recalculated periodically as the state of the DMs evolves over time.

As direct imaging missions demand DMs with ever-higher actuator density to enable wider and wider search areas, computational aspects will inevitably become a point of concern from a systems engineering standpoint. In on-orbit WFS&C, all sensing and control computations are processed by the flight computer; conversely, in ground-in-the-loop scenarios, raw data is communicated to a ground-based computing node that calculates the DM correction and relays the correction back to the observatory. Though each approach has tradeoffs, a major advantage of on-orbit WFS&C is the ability to update DM commands more frequently without relying on the continuous availability of a communication link with the ground station. On-orbit WFS&C can help to relax observatory-level wavefront stability requirements by enabling the high-contrast dark zone to be maintained over shorter time intervals. Successful deployment of an on-orbit architecture is predicated on the availability of sufficient computational resources; however, radiation-hard, space-qualified computing hardware lags behind conventional hardware by decades and is extremely resource-limited. This poses a substantial computation capability gap if the current algorithms are expected to be deployed on-orbit on a future direct imaging mission, such as NASA's planned Habitable Worlds Observatory (HWO).

As a case in point, the current state-of-the-art testbed for coronagraph laboratory demonstrations, the decadal survey testbed (DST) at NASA's Jet Propulsion Laboratory, has successfully demonstrated 3.82×10^{-10} instrumental contrast over a 10% fractional bandpass within an annular dark zone extending from $3 \lambda_0/D$ to $8 \lambda_0/D$, where λ_0 is the central imaging wavelength and D is the diameter of the telescope aperture, using two DMs each with 48×48 actuators.⁴ The baseline requirements for an HWO-like mission will be considerably steeper, using the Large UV/Optical/Infrared Surveyor (LUVOIR) and Habitable Exoplanet Observatory (HabEx) flagship mission concepts formulated for the Astro2020 Decadal Survey^{5,6} as representative reference designs. The HabEx design features two 64×64 DMs, and a search area with a maximal outer radius of $32 \lambda_0/D$ over a 20% bandpass—nearly double the total actuator count, a factor of 4 increase in search radius, and a factor of two increase in control bandwidth. Meanwhile, the LUVOIR Architecture “A” reference design includes a pair of 128×128 DMs and a dark zone with a $64 \lambda_0/D$ maximal outer radius over a 10% fractional bandpass or wider. Because the number of detector pixels in the dark zone scales with the area of the dark zone rather than its radius, these parameters correspond to an increase by a factor of 32 (HabEx) and 512 (LUVOIR-A) of the worst-case Jacobian dimensionality, for each control wavelength, compared to DST.

In recent work, we formulated an alternative wavefront control framework that iteratively compute DM updates using gradient-based optimization techniques, which eliminates the need to calculate and manipulate the Jacobian.⁷ Our approach is based in part on a numerical technique known as algorithmic differentiation (AD),⁸ to calculate the gradients of the cost function for optimal control accurately and efficiently. We described an AD-based counterpart to SM,

which we termed AD penalty stroke minimization (AD-PSM), and compared it to SM in simulations with a small-angle LUVUOIR design. Our results indicated superior computational efficiency with AD-PSM and comparable starlight suppression performance. While the CPU time and memory consumption of SM grew superlinearly with actuator count, the increase in both for AD-PSM was negligible (e.g., with 128×128 actuators, AD-PSM utilized 95% less memory and CPU time), suggesting that iterative methods are a promising alternative to Jacobian-based techniques for on-orbit WFS&C with high actuator counts and large dark zones.

In this paper, we report on the first experimental demonstration of AD-PSM as well as an AD-based counterpart to the EFC algorithm, termed AD-EFC, using the High-contrast Imager for Complex Aperture Telescopes (HiCAT) at the Space Telescope Science Institute in Baltimore, Maryland. We benchmark the contrast performance of AD-PSM and AD-EFC as a function of several key parameters including regularization and the termination tolerance of the nonlinear optimizer, and compare it to SM and EFC.

This paper is structured as follows. In Sec. 2, we review concepts from our earlier work, including AD and the mathematical formulation of AD-PSM and AD-EFC. In Sec. 3, we provide an overview of HiCAT and discuss our experimental setup, including algorithm implementation details that are pertinent to our demonstration. In Sec. 4, we present and discuss our experimental results. Finally, in Sec. 5, we provide our conclusions and briefly describe our planned future work.

1.1 Notation

In this paper, our principal concern is with algorithms that operate on discrete vector-valued quantities, which are represented in boldface. Many of these quantities vary as a function of control iteration, and are denoted with the subscript k . These may be truly discrete, such as the vector \mathbf{a}_k of DM actuator command updates, or may represent discretizations of functions of continuous spatial variables, such as electric fields. We denote \mathbf{x} as a column vector, its transpose \mathbf{x}^T as a row vector, and $\|\mathbf{x}\|$ as its Euclidean length. For complex-valued quantities, \dagger denotes the Hermitian transpose. Scalar quantities are denoted in ordinary (i.e., non-boldface) typographic weight.

By convention, the derivative of a scalar with respect to a column vector is a row vector, i.e.

$$\frac{\partial J}{\partial \mathbf{x}} = \left[\frac{\partial J}{\partial \mathbf{x}[1]} \cdots \frac{\partial J}{\partial \mathbf{x}[N]} \right], \quad (1)$$

where $\mathbf{x}[n]$ is the n 'th element of \mathbf{x} . Consequently, the derivative of a column vector with respect to another column vector is a matrix of row vectors:

$$\frac{\partial \mathbf{y}}{\partial \mathbf{x}} = \begin{bmatrix} \frac{\partial \mathbf{y}[1]}{\partial \mathbf{x}} \\ \dots \\ \frac{\partial \mathbf{y}[M]}{\partial \mathbf{x}} \end{bmatrix}. \quad (2)$$

2 Wavefront Control Using Algorithmic Differentiation

The goal of the WFS&C loop in coronagraphy is to drive starlight within the dark zone toward zero over time so that a faint orbiting companion becomes detectable against the reduced starlight background. Each iteration of the WFS&C loop, indexed by the integer k , consists of two steps: (1) an estimation step, in which an estimate $\hat{\mathbf{E}}_{\text{ab},k}$ of the true aberrated electric field $\mathbf{E}_{\text{ab},k}$ is formed within the dark zone from focal-plane intensity measurements, and (2) a control step, in which the DM correction is updated to reduce the energy in $\mathbf{E}_{\text{ab},k}$, as shown in Fig. 1. In this paper, we focus principally on the control step.

Modern model-based wavefront control algorithms find the DM correction update \mathbf{a}_k by minimizing some cost function $J_k(\mathbf{a}_k; \hat{\mathbf{E}}_{\text{ab},k})$ with respect to \mathbf{a}_k . Usually, J_k is constructed to trade off between minimizing starlight and minimizing the size of the correction, which helps to regularize the problem and stabilize the solution. In general, the true relationship between the starlight in the dark zone and the DM correction is highly nonlinear and nonconvex, owing to the fact that the DMs impart phase-only corrections of the form $\exp\{i\phi_{\text{DM}}\}$ at or near the

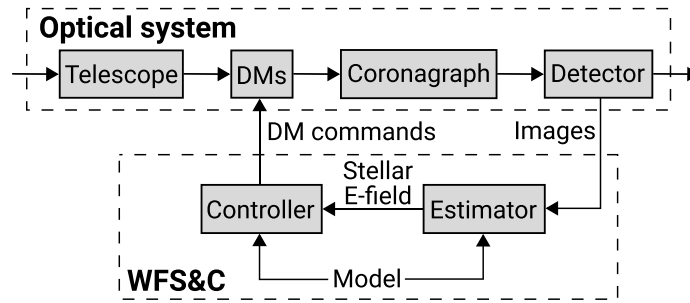


Fig. 1 Closed-loop coronagraphic WFS&C uses images from the science camera, rather than an external wavefront sensing instrument, to estimate the electric field from the host star within the dark zone and drive it toward zero. Model-based WFS&C algorithms use a numerical model of the coronagraph to solve an inverse problem for the unknown electric field and corresponding DM correction, respectively.

coronagraph entrance pupil. However, when the optical aberrations are small, we can approximate the true electric field in the coronagraph entrance pupil with a first-order Taylor series expansion about a small update to the DM commands. In this case, the corrected electric field in the dark zone has the form

$$\mathbf{E}_{\text{DZ},k}(\mathbf{a}_k) \approx \hat{\mathbf{E}}_{\text{ab},k} + \mathbf{E}_{\text{DM},k}(\mathbf{a}_k), \quad (3)$$

where $\hat{\mathbf{E}}_{\text{ab},k}$ is the estimate of the aberrated dark-zone electric field produced by the estimation step, and $\mathbf{E}_{\text{DM},k}$ is the change in electric field resulting from the update to the DM correction. We can also write $\mathbf{E}_{\text{DZ},k}$ in the form

$$\mathbf{E}_{\text{DZ},k}(\mathbf{a}_k) = \hat{\mathbf{E}}_{\text{ab},k} + \mathbf{G}_k \mathbf{a}_k, \quad (4)$$

where $\mathbf{G}_k \triangleq \partial \mathbf{E}_{\text{DZ},k} / \partial \mathbf{a}_k$ is the Jacobian matrix with dimensions $N_{\text{pix}} \times N_{\text{act}}$, N_{pix} is the number of dark-zone pixels, and N_{act} is the total number of controllable DM actuators. The intensity from the corrected electric field, integrated over the dark zone, can be written in terms of the Jacobian as

$$\|\mathbf{E}_{\text{DZ},k}(\mathbf{a}_k)\|^2 = \|\mathbf{G}_k \mathbf{a}_k\|^2 + \|\hat{\mathbf{E}}_{\text{ab},k}\|^2 + 2 \text{Re}\{\hat{\mathbf{E}}_{\text{ab},k}^\dagger \mathbf{G}_k\} \mathbf{a}_k. \quad (5)$$

This is a quadratic function of \mathbf{a}_k , meaning that under this approximation, there exists a unique, optimal DM correction that minimizes the dark-zone starlight.

SM and EFC utilize the relationship in Eq. (5) to derive closed-form expressions for this optimal correction in terms of \mathbf{G}_k that can be evaluated by solving a linear system of dimension $N_{\text{act}} \times N_{\text{act}}$, as illustrated in Fig. 2. As we discuss in Appendix B, this is equivalent to minimizing the cost function using Newton's method with an exact Hessian matrix. Alternatively, it is possible to find the DM correction by minimizing the cost function J_k with respect to \mathbf{a}_k iteratively, rather than analytically, using gradient-based nonlinear optimization as shown in Fig. 3.




Fig. 2 In EFC and SM, the Jacobian matrix is precomputed using a computer model of the coronagraph to predict the effect of an update to each of the N_{act} DM actuators individually on the dark-zone electric field, here represented by the Kronecker δ functions δ_n , where $\delta_n[i] = 1$ if $i = n$ and zero otherwise. The Jacobian \mathbf{G}_k and the estimate of the aberrated dark-zone electric field $\hat{\mathbf{E}}_{\text{ab},k}$ together are used to construct a linear system of equations whose solution is the desired DM update \mathbf{a}_k^* . We show in Appendix B that this is equivalent to minimizing the wavefront control cost function using Newton's method.

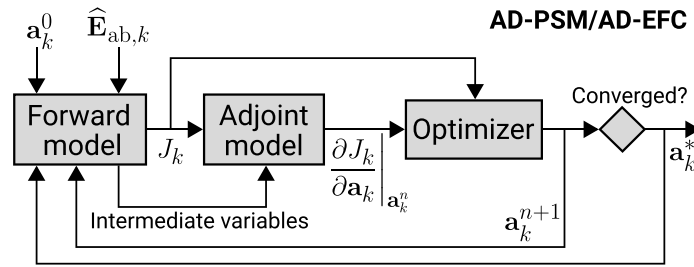


Fig. 3 Our AD-based wavefront controllers AD-PSM and AD-EFC use RMAD to differentiate the wavefront control cost function J_k with respect to the DM correction update \mathbf{a}_k , yielding the gradient vector $\partial J_k / \partial \mathbf{a}_k$ evaluated at the current iterate \mathbf{a}_k^n . A nonlinear optimization algorithm calculates a new iterate \mathbf{a}_k^{n+1} that reduces the value of the cost function, i.e., $J_k(\mathbf{a}_k^{n+1}) \leq J_k(\mathbf{a}_k^n)$. This procedure is repeated until the gradient becomes sufficiently small, indicating that the solution \mathbf{a}_k^* is at or near a local minimum of the cost function. A starting guess for the solution \mathbf{a}_k^0 as well as the aberrated electric field $\hat{\mathbf{E}}_{ab,k}$ are the input parameters.

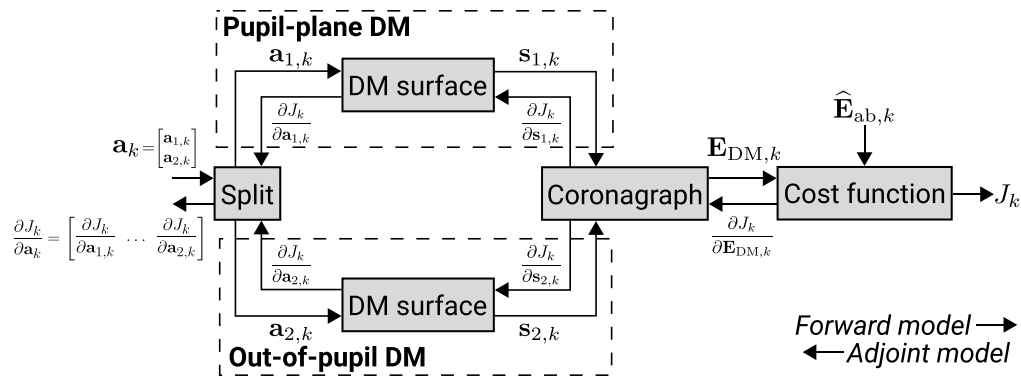


Fig. 4 The forward model for the wavefront control problem maps DM command updates \mathbf{a}_k to a scalar cost function J_k . The DM command vector is split into independent command vectors $\mathbf{a}_{1,k}$ and $\mathbf{a}_{2,k}$ for the pupil-plane and out-of-pupil DM, respectively. These are mapped onto DM surfaces $\mathbf{s}_{1,k}$ and $\mathbf{s}_{2,k}$ using the model in [Appendix C](#), and propagated through an end-to-end coronagraph model to predict the resulting change in dark-zone electric field $\mathbf{E}_{DM,k}$. Reverse-mode AD transforms the forward model into an adjoint model that backpropagates the partial derivatives of J_k with respect to each intermediate variable $\mathbf{a}_{1,k}$, $\mathbf{a}_{2,k}$, $\mathbf{s}_{1,k}$, $\mathbf{s}_{2,k}$, and $\mathbf{E}_{DM,k}$ in reverse order, starting from the output on the right. The derivatives with respect to the individual DM command vectors $\mathbf{a}_{1,k}$ and $\mathbf{a}_{2,k}$ are concatenated to form the full gradient vector for optimization.

To do so eliminates the need to calculate the Jacobian matrix, but requires a way of calculating the gradient vector $\partial J_k / \partial \mathbf{a}_k$. Reverse-mode algorithmic differentiation (RMAD) provides a way of doing so that is both computationally efficient and accurate, in the sense that the derivatives computed by RMAD are accurate to machine precision and do not utilize finite-difference approximations.⁸ The basic principle of RMAD is that any function that can be written down as a sequence of differentiable operations, called the forward model, can be transformed mechanically to construct a related function, called the adjoint model, that evaluates the derivative of the forward model with respect to any of the intermediate variables encountered during its evaluation, as well as its inputs. Figure 4 illustrates this procedure for the wavefront control cost function J_k . We refer the reader to our earlier work for a more comprehensive discussion of the principles of RMAD and its application to wavefront control.⁷

2.1 Stroke Minimization: From Lagrange Multipliers to Penalty Method

SM finds the smallest DM correction that achieves a desired level of stellar intensity integrated over the dark zone, denoted by $I_{\text{target},k}$.² It is solved by finding the stationary point of the Lagrangian function

$$\mathcal{L}_{SM,k} = \|\mathbf{a}_k\|^2 + \mu_k (\|\mathbf{E}_{DZ,k}(\mathbf{a}_k)\|^2 - I_{\text{target},k}), \quad (6)$$

i.e., a point such that $\partial\mathcal{L}_{\text{SM},k}/\partial\mathbf{a}'_k = 0$, where \mathbf{a}_k is the DM actuator command update, $\mathbf{E}_{\text{DZ},k}$ is the corrected electric field in the dark zone, $\mathbf{a}'_k \triangleq [\mathbf{a}_k^T \mu_k]^T$, and μ_k is the Lagrange multiplier. Because this stationary point is a saddle point, it cannot be reached by minimizing $\mathcal{L}_{\text{SM},k}$ with respect to \mathbf{a}'_k . Instead, one chooses a fixed starting value for the Lagrange multiplier, μ_k^0 , and minimizes $\mathcal{L}_{\text{SM},k}$ with respect to \mathbf{a}_k to find a corresponding DM solution \mathbf{a}_k^0 . If the constraint $\|\mathbf{E}_{\text{DZ},k}(\mathbf{a}_k^0)\|^2 \leq I_{\text{target},k}$ is not satisfied, a larger value $\mu_k^1 > \mu_k^0$ is selected and this procedure is repeated. For any fixed value of μ_k^n , the minimum of Eq. (6) is given in closed form in terms of the Jacobian matrix \mathbf{G}_k :

$$\mathbf{a}_k^n = -\left(\text{Re}\{\mathbf{G}_k^\dagger \mathbf{G}_k\} + \frac{1}{\mu_k^n} \mathbb{I}\right)^{-1} \text{Re}\{\mathbf{G}_k^\dagger \hat{\mathbf{E}}_{\text{ab},k}\}, \quad (7)$$

where \mathbb{I} is the identity matrix and $\hat{\mathbf{E}}_{\text{ab},k}$ is the estimate of the aberrated electric field.

AD-PSM has the same goal, but instead iteratively minimizes the cost function

$$J_{\text{PSM},k} = \|\mathbf{a}_k\|^2 + \eta_k (\|\mathbf{E}_{\text{DZ},k}(\mathbf{a}_k)\|^2 - I_{\text{target},k})^2, \quad (8)$$

where η_k is a parameter that encodes the relative importance of minimizing actuator stroke and driving the integrated intensity toward the target value. The minimum of $J_{\text{PSM},k}$ with respect to \mathbf{a}_k is coincident with the stationary point of $\mathcal{L}_{\text{SM},k}$,⁹ meaning that the DM solutions obtained by SM and AD-PSM are, in principle, identical. In this paper, we choose η_k in each WFS&C iteration as

$$\eta_k = \frac{\eta_{00}}{(\|\hat{\mathbf{E}}_{\text{ab},k}\|^2 - I_{\text{target},k})^2}, \quad (9)$$

where η_{00} , known as the penalty parameter, is a constant set by the experimenter. The denominator $(\|\hat{\mathbf{E}}_{\text{ab},k}\|^2 - I_{\text{target},k})^2$ scales the cost function to be invariant to the energy in the dark-zone electric field, which helps in practice to obtain good solutions in all WFS&C iterations as $\hat{\mathbf{E}}_{\text{ab},k}$ gradually converges toward zero.

2.2 Electric Field Conjugation

EFC attempts to drive the dark zone electric field toward zero, with Tikhonov regularization to mitigate ill-conditioning caused by the presence of actuators that are completely or partially obscured by pupil features. Its cost function for a single correction wavelength is given as

$$J_{\text{EFC},k} = \|\mathbf{E}_{\text{DZ},k}(\mathbf{a}_k)\|^2 + \|\mathbf{\Gamma}_k \mathbf{a}_k\|^2, \quad (10)$$

where $\mathbf{\Gamma}_k$ is the Tikhonov regularization matrix. In the most common case, one chooses $\mathbf{\Gamma}_k = \alpha_k \mathbb{I}$, making EFC identical to SM with a fixed Lagrange multiplier $\mu_k = 1/\alpha_k^2$. For this case, its solution can be obtained using Eq. (7). The general solution for the Jacobian-based formulation of EFC is given as

$$\mathbf{a}_k^* = -(\text{Re}\{\mathbf{G}_k^\dagger \mathbf{G}_k\} + \mathbf{\Gamma}_k^T \mathbf{\Gamma}_k)^{-1} \text{Re}\{\mathbf{G}_k^\dagger \hat{\mathbf{E}}_{\text{ab},k}\}. \quad (11)$$

In AD-EFC, one iteratively minimizes a scaled version of the EFC cost function

$$J'_{\text{EFC},k} = \frac{1}{\|\hat{\mathbf{E}}_{\text{ab},k}\|^2} J_{\text{EFC},k}. \quad (12)$$

Similarly to AD-PSM, the scaling factor $1/\|\hat{\mathbf{E}}_{\text{ab},k}\|^2$ makes the AD-EFC cost function invariant to the energy in the dark-zone electric field to aid in obtaining numerical solutions in practice. The RMAD adjoint model for $J_{\text{EFC},k}$ is provided in [Appendix D](#).

3 Experimental Setup

In this section, we provide an overview of the HiCAT testbed and provide details about the implementation of AD-PSM and AD-EFC, the reference Jacobian-based implementations of SM and EFC, and the electric field estimation algorithm used in the estimation step of the WFS&C loop.

3.1 HiCAT Testbed

HiCAT is a testbed dedicated to technology demonstrations for coronagraphy on segmented-aperture space observatories, with the intent of being directly traceable to a future HWO-like mission.^{10–17} These technologies include Lyot coronagraphy, high-order WFS&C for generating and stabilizing dark zones, and low-order wavefront sensing (LOWFS).¹⁸ HiCAT operates in a mid-contrast regime (10^{-7} to 10^{-8}) which approaches the limit achievable outside of a vacuum environment. The testbed is equipped with two Boston Micromachines Kilo-DMs for high-order sensing and control, with 952 actuators each, making it suitable for our proof-of-concept demonstrations. One DM is placed in a plane conjugate to the HiCAT entrance pupil, while the second DM is placed ~ 300 mm farther along the optical axis, corresponding to a Fresnel number $N_F \approx 98$ at a wavelength of 638 nm. This configuration enables simultaneous control of amplitude and phase aberrations over a dark zone that extends over both halves of the image plane. We conducted our experiments using a Thorlabs MCLS1 laser diode source, which emits monochromatic light with a central wavelength $\lambda_0 = 638$ nm.

HiCAT additionally has an IrisAO segmented DM with 37 hexagonal segments, each with controllable piston/tip/tilt, to act as a telescope pupil simulator and to inform experimental efforts devoted to segment-level tolerancing and stabilization.^{19,20} Figure 5 shows a simplified system layout of HiCAT, including the primary imaging beam path as well as several additional beam paths used by the LOWFS and metrology subsystems.

Our experiments on HiCAT utilized a classical Lyot coronagraph (CLC) design with a hexagonal entrance pupil mask designed to mask the extreme edges of the IrisAO, a circular focal-plane mask, and a circular Lyot stop. Figure 6 shows an overlay of the CLC pupil masks along with a simulated stellar point-spread functions (PSFs). Figure 7 shows example experimental PSFs obtained before and after closed-loop WFS&C using SM, along with the corresponding DM commands.

3.2 Algorithm Implementation

We developed a differentiable model of HiCAT using Python, including a hand-derived adjoint model. To facilitate testing, our model was comprised of several sub-modules each with its own individual forward and adjoint models.

1. A model to compute the DM surface resulting from a given set of actuator commands using a fast convolutional representation. This is described in further detail in [Appendix C](#).

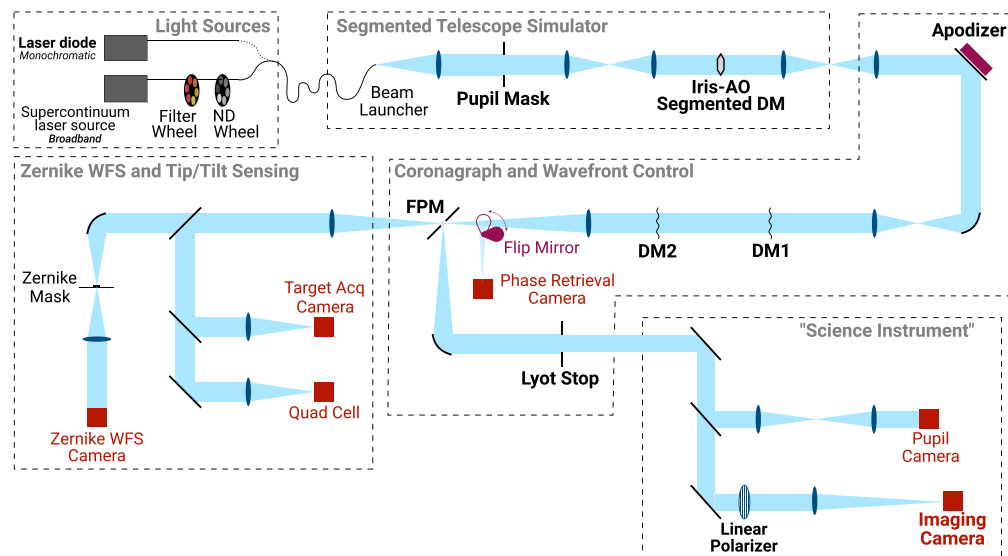


Fig. 5 Simplified, partially-unfolded layout of the HiCAT testbed. The elements encountered by the primary imaging beam path are highlighted in bold. DM1 and DM2 indicate the in-pupil and out-of-pupil DMs, respectively. Our experiments utilized a coronagraph configuration in which the apodizer is replaced with a flat mirror (see Sec. 3.1).




Fig. 6 (a) Overlay of HiCAT pupil masks projected onto the in-pupil DM plane, including the reflective area of the DM, the IrisAO segmented aperture, the entrance pupil mask (with the same geometry, but slightly undersized relative to the IrisAO) and the Lyot stop. (b) Simulated and (c) experimental coronagraphic PSFs with outline of the geometrical edge of the focal-plane mask.

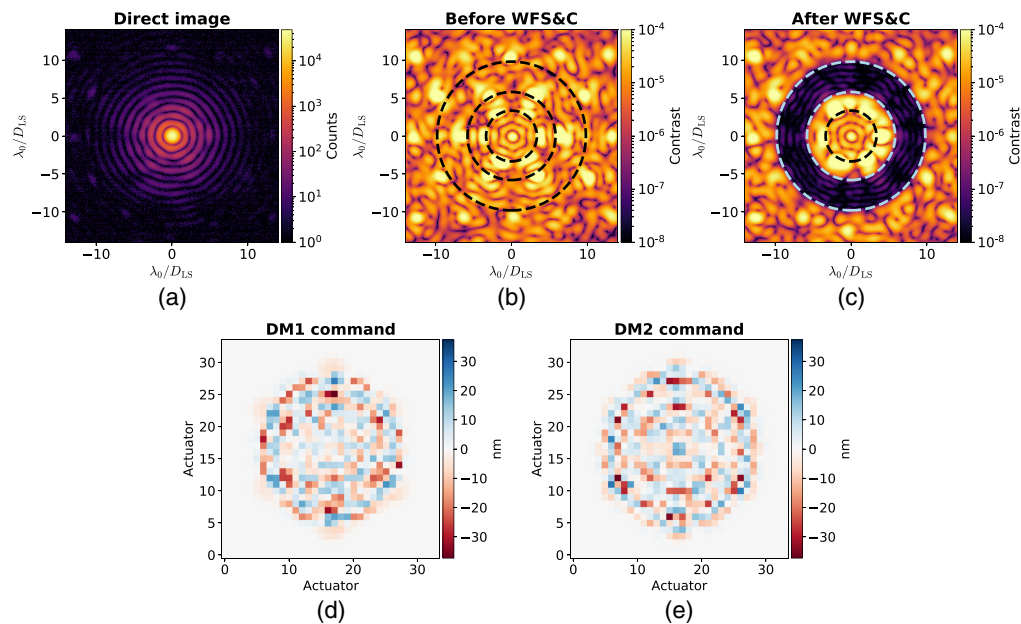


Fig. 7 Experimental on-axis images from HiCAT. (a) Non-coronagraphic image, (b) coronagraphic image prior to the WFS&C loop, (c) coronagraphic image after 80 iterations of SM, and corresponding actuator commands for the in-pupil DM (d), and out-of-pupil DM (e). In panels (b) and (c), the geometrical edge of the focal-plane mask (FPM) with radius $3.34 \lambda_0/D_{LS}$ is shown as well as the inner and outer edges of the dark zone at $5.8 \lambda_0/D_{LS}$ to $9.8 \lambda_0/D_{LS}$, respectively, where D_{LS} is the Lyot stop diameter and $\lambda_0 = 638$ nm.

2. A model to apply the phase corrections imparted by the in-pupil DM (DM1) and out-of-pupil DM (DM2), including the free-space propagation between the two DMs.
3. A model to propagate the electric field after DM correction through the HiCAT CLC. We use the semi-analytical Lyot coronagraph model originally described in Ref. 21.

We refer the reader to our earlier work for detailed descriptions of the operations in the forward and adjoint models for the latter two sub-modules.⁷ A pre-existing numerical model of HiCAT based on the POPPY framework^{22,23} served as a reference for calibrating our differentiable model.

3.2.1 Optimization algorithm

We used the limited-memory Broyden–Fletcher–Goldfarb–Shanno (L-BFGS) algorithm,⁹ as implemented in the SciPy package,²⁴ as the optimization algorithm for AD-PSM and AD-EFC. L-BFGS is a quasi-Newton algorithm, meaning that it uses the gradient vectors collected during

optimization to approximate the inverse Hessian matrix \mathbf{H}_k^{-1} of second derivatives, involving substantially less effort than Newton's method with an exact Hessian (see Appendix B).

All numerical optimization algorithms have a termination criterion that determines when the algorithm has reached a location in the parameter space that is sufficiently close to a local optimum. In the L-BFGS implementation used by SciPy, the termination criterion is set by a tolerance parameter ε defined as the magnitude of the largest element of the gradient vector. As $\varepsilon \rightarrow 0$, L-BFGS carries out a greater number of optimization iterations to terminate closer to the minimum. Choosing a larger value of ε reduces total computation with the tradeoff of a less-optimal DM solution. However, as our experiments showed, in certain instances this can help to regularize the solution by terminating before L-BFGS converges to an overly aggressive DM correction that would otherwise make the WFS&C loop unstable.

The termination criterion presents an additional consideration not present in SM and EFC which exactly minimize their respective cost functions in each WFS&C iteration. The value of the tolerance parameter ε should be chosen to optimally trade-off between accuracy and computational effort. We discuss this in further detail below.

3.3 Estimation Algorithm

We used the pairwise probe estimator³ for the estimation step in all experiments. The pairwise estimator forms a least-squares estimate $\hat{\mathbf{E}}_{ab,k}$ of the focal-plane electric field $\mathbf{E}_{ab,k}$ by applying a series of P probing DM commands \mathbf{u}_p to generate probing electric fields $\mathbf{E}_{DM,k}(\mathbf{u}_p)$ that interfere with $\mathbf{E}_{ab,k}$. The data vector for the least-squares estimate is formed by differencing the images resulting from $\mathbf{E}_{DM,k}(\mathbf{u}_p)$ and $\mathbf{E}_{DM,k}(-\mathbf{u}_p)$. The estimate of the m 'th pixel in the dark zone is then found by finding the least-squares solution of the system

$$\begin{bmatrix} \mathbf{I}_{k,1}^+[m] - \mathbf{I}_{k,1}^-[m] \\ \vdots \\ \mathbf{I}_{k,P}^+[m] - \mathbf{I}_{k,P}^-[m] \end{bmatrix} = \begin{bmatrix} \text{Re}\{\mathbf{E}_{DM,k}(\mathbf{u}_1)\}[m] & \text{Im}\{\mathbf{E}_{DM,k}(\mathbf{u}_1)\}[m] \\ \vdots & \vdots \\ \text{Re}\{\mathbf{E}_{DM,k}(\mathbf{u}_P)\}[m] & \text{Im}\{\mathbf{E}_{DM,k}(\mathbf{u}_P)\}[m] \end{bmatrix} \begin{bmatrix} \text{Re}\{\mathbf{E}_{ab,k}\}[m] \\ \text{Im}\{\mathbf{E}_{ab,k}\}[m] \end{bmatrix}, \quad (13)$$

where $\mathbf{I}_{k,p}^\pm \triangleq |\mathbf{E}_{ab,k} + \mathbf{E}_{DM,k}(\pm\mathbf{u}_p)|^2$.

We generated four DM probe functions \mathbf{u}_p that were optimized to produce probing fields of the form

$$\mathbf{E}_{DM,k}^p = \text{sgn}\{\theta_x\} \exp\left\{i\pi \frac{p-1}{4} \text{sgn}\{\theta_x\}\right\}, \quad (14)$$

where $p \in \{1,2,3,4\}$ and $\text{sgn}\{\theta_x\}$ is the sign of the focal-plane x coordinate. These probing fields, which have anti-Hermitian symmetry, correspond to purely real commands applied to the in-pupil DM, which we can calculate via a regularized least-squares approach for each p ,

$$\arg \min_{\mathbf{u}_k^p} \|\mathbf{G}_{1,k} \mathbf{u}_k^p - \mathbf{E}_{DM,k}^p\|^2 + \alpha_{\text{probe}}^2 \|\mathbf{u}_k^p\|^2, \quad (15)$$

where α_{probe}^2 is the Tikhonov regularization term for the probe calculation and $\mathbf{G}_{1,k}$ is the Jacobian for the in-pupil DM. This is identical to the EFC problem in Eq. (10) but with the probing field $\mathbf{E}_{DM,k}^p$ on the right-hand side instead of $-\hat{\mathbf{E}}_{ab,k}$, and has the solution

$$\mathbf{u}_k^p = (\text{Re}\{\mathbf{G}_{1,k}^\dagger \mathbf{G}_{1,k}\} + \alpha_{\text{probe}}^2 \mathbb{I})^{-1} \text{Re}\{\mathbf{G}_{1,k}^\dagger \mathbf{E}_{DM,k}^p\}. \quad (16)$$

For the experiments with SM and EFC, we applied Eq. (16) to compute probe commands for pairwise probe estimation. For experiments with AD-PSM and AD-EFC, we iteratively solved Eq. (15) using the AD-EFC framework. In all cases, we set the Tikhonov regularization parameter for probe generation as $\alpha_{\text{probe}}^2 = 0.7$; for the iterative probes, we set the optimizer tolerance to $\varepsilon = 10^{-5}$. Figure 8 shows the probe commands \mathbf{u}_p along with the corresponding magnitude and phase of $\mathbf{E}_{DM,k}(\mathbf{u}_p)$.

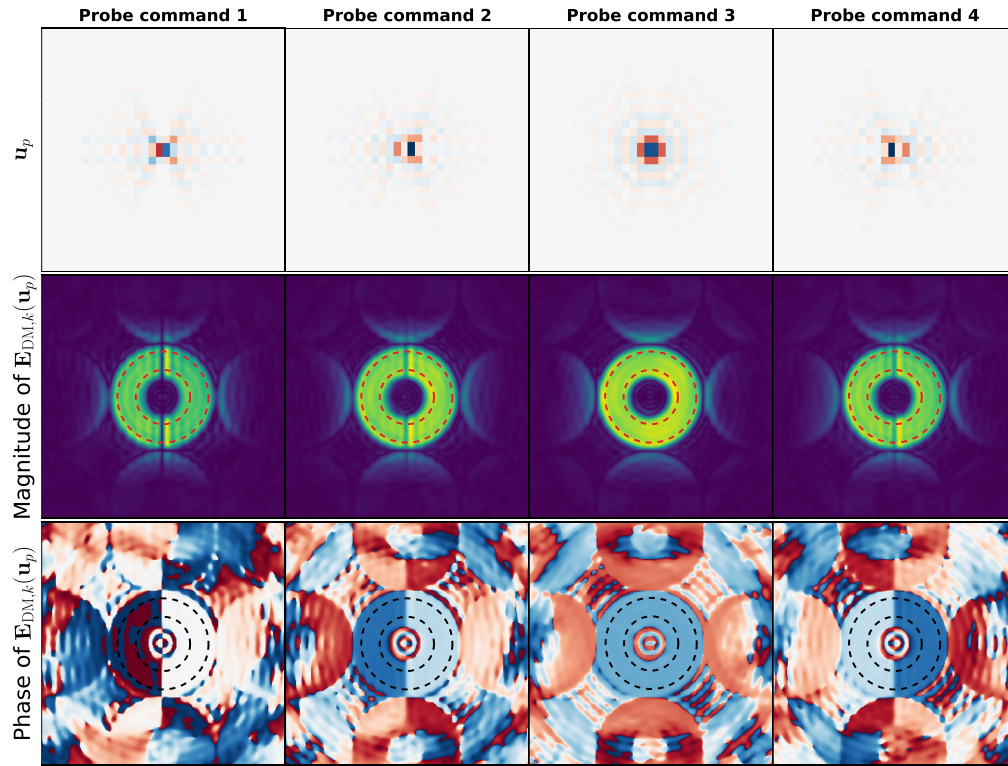


Fig. 8 The probe commands \mathbf{u}_p for pairwise estimation were optimized so that the resultant dark-zone electric field $\mathbf{E}_{\text{DM},k}(\mathbf{u}_p) = \text{sgn}\{\rho_x\} \exp\{i\pi(\rho - 1)\text{sgn}\{\rho_x\}/4\}$, as described in Sec. 3.3. The inner and outer edges of the dark zone are also shown for reference. The probe commands are close to the inverse Fourier transform of the dark zone geometry (an annulus), modulated by a horizontal sinusoid whose phase angle is proportional to the desired piston phase, and projected onto the DM actuator coordinates. Only the pupil-plane DM is modulated.

4 Experimental Results and Discussion

We conducted a series of experiments to compare the contrast performance of AD-PSM and AD-EFC relative to SM and EFC, respectively. All experiments used an annular control region extending from $5.8 \lambda_0/D_{\text{LS}}$ to $9.8 \lambda_0/D_{\text{LS}}$, where $\lambda_0 = 638 \text{ nm}$ is the central wavelength of the Thorlabs MCLS1 laser diode and D_{LS} is the Lyot stop diameter. Each experiment consisted of 80 WFS&C iterations. For each experiment, we computed the median and 10th and 90th percentile values of the spatially averaged dark-zone contrast values for the final 50 iterations, which captured the steady-state performance of the WFS&C loop with each algorithm after converging to deepest-possible contrast. Figures 9 and 10 show example contrast vs. iteration time series with AD-PSM/SM and AD-EFC/EFC, respectively, as well as the PSF from the iteration with deepest contrast using AD-PSM and AD-EFC. In both cases, the convergence properties of AD-PSM and AD-EFC are nearly identical to their Jacobian-based counterparts, validating our proposed approach.

For SM and AD-PSM, we chose $I_{\text{target},k} = 0.5 \|\hat{\mathbf{E}}_{\text{ab},k}\|^2$ as the contrast target, i.e., a factor of two improvement iteration-over-iteration in spatially integrated dark zone contrast. For the Lagrange multiplier line search in SM, we let $\mu_k^{n+1} = 1.3\mu_k^n$. For AD-PSM, we tested combinations of the penalty parameter $\eta_{00} \in \{10, 100, 1000\}$ and the nonlinear optimization convergence tolerance $\varepsilon \in \{10^{-2}, 10^{-3}, 10^{-4}\}$. We determined these values based on simulations and prior WFS&C experiments on HiCAT.

For EFC and AD-EFC, we selected the Tikhonov regularization matrix as $\mathbf{\Gamma}_k = \alpha_k \mathbb{I}$, with $\alpha_k^2 \in \{10^{-2}, 10^{-3}, 10^{-4}\}$ for the first 30 WFS&C iterations and $\alpha_k^2 = 10^{-1}$ thereafter. We determined through previous experiments with EFC that maintaining an aggressive α_k value throughout the experiment was detrimental to the stability of the control loop after reaching the steady-state regime. For AD-EFC, we used the same set of ε values as AD-PSM.

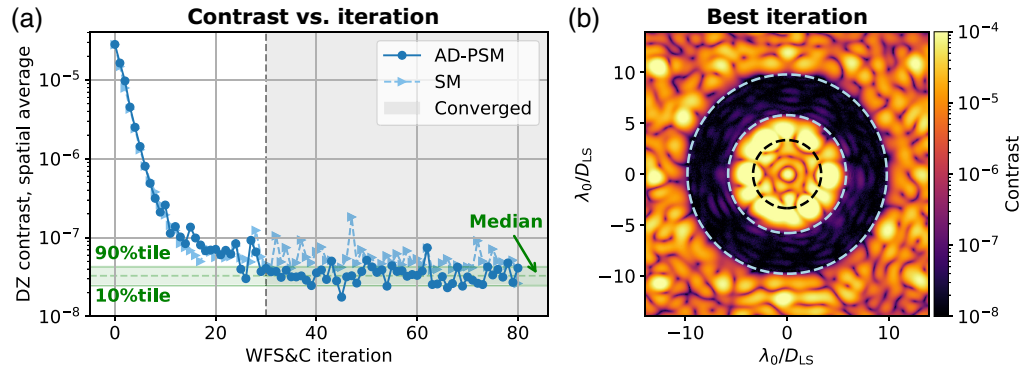


Fig. 9 (a) Spatially averaged dark-zone contrast versus WFS&C iteration using AD-PSM with $\eta_{00} = 10$ and $\varepsilon = 10^{-4}$, compared to an experiment with SM. The median, 10th, and 90th percentile of the final 50 iterations, representing the steady-state behavior after convergence, are shown in green. (b) The on-axis PSF corresponding to the iteration with deepest contrast using AD-PSM.

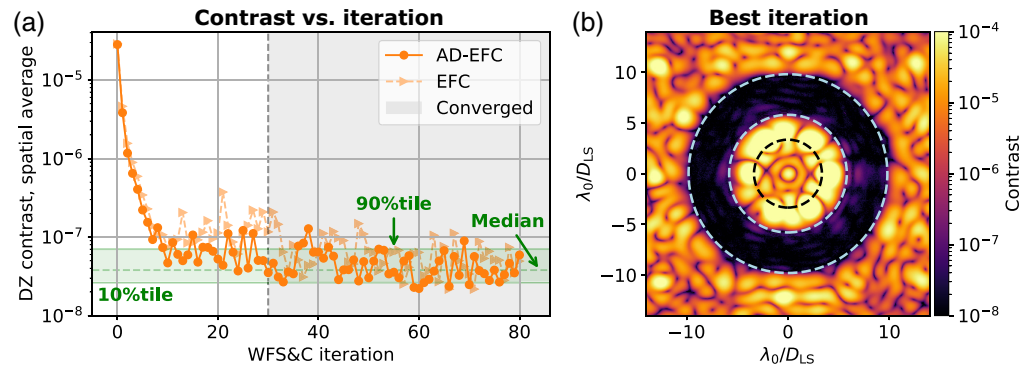


Fig. 10 (a) Spatially averaged dark-zone contrast versus WFS&C iteration using AD-EFC with $\alpha_k^2 = 10^{-2}$ and $\varepsilon = 10^{-4}$, compared to an experiment with EFC using the same value for α_k . The median, 10th, and 90th percentile of the converged datapoints are shown in green. (b) The on-axis PSF corresponding to the iteration with deepest contrast using AD-EFC.

For AD-EFC and AD-PSM, we compared the value of the cost function for two different starting guesses for the DM correction: $\mathbf{a}_k^0 = \mathbf{a}_{k-1}^*$, i.e., the solution of the previous WFS&C iteration, and $\mathbf{a}_k^0 = 0$. The starting guess with the lower of the two cost function values was then selected. In all cases $\mathbf{a}_k^0 = 0$ was ultimately chosen as the starting guess.

Figures 11 and 12 show the statistics of the post-convergence spatially averaged dark-zone contrast achieved with AD-PSM and AD-EFC for each combination of regularization (η_{00} or α_k) and optimization tolerance ε , respectively, compared to reference experiments with SM and EFC. For all parameter combinations, AD-PSM and AD-EFC equaled the contrast performance of SM and EFC, respectively. Moreover, we observed no strongly identifiable trends in achievable contrast as a function of the algorithm parameters, indicating that in all cases, AD-PSM and AD-EFC reached the contrast floor imposed by environmental instabilities. For full contrast versus iteration curves for all experiments, we refer to Fig. 13 in Appendix A.

4.1 Discussion

Our experiments were aimed at exploring a relevant subset of the space of free parameters for each algorithm, namely the nonlinear optimization convergence tolerance ε , the Tikhonov regularization α_k for AD-EFC, and the penalty parameter η_{00} for AD-PSM. In principle, each parameter affects the attainable contrast of the WFS&C loop, but in subtly different ways, which we discuss here.

As discussed in Sec. 3.2.1, ε determines the effort that the nonlinear optimization algorithm will expend to find a solution close to the true minimum of the cost function. A smaller value of ε corresponds to a larger number of optimization iterations before termination, and ultimately a




Fig. 11 Median, 10th percentile, and 90th percentile spatially averaged contrast values achieved by AD-PSM (dark blue) as a function of optimizer tolerance, for three different values of the penalty parameter η_{00} , compared to SM (light blue). The contrast vs. iteration time series for the rightmost result in the left pane ($\eta_{00} = 10$, $\varepsilon = 10^{-4}$) is illustrated in Fig. 9. In all cases, the contrast performance of AD-PSM was equivalent to that of SM within statistical uncertainty.

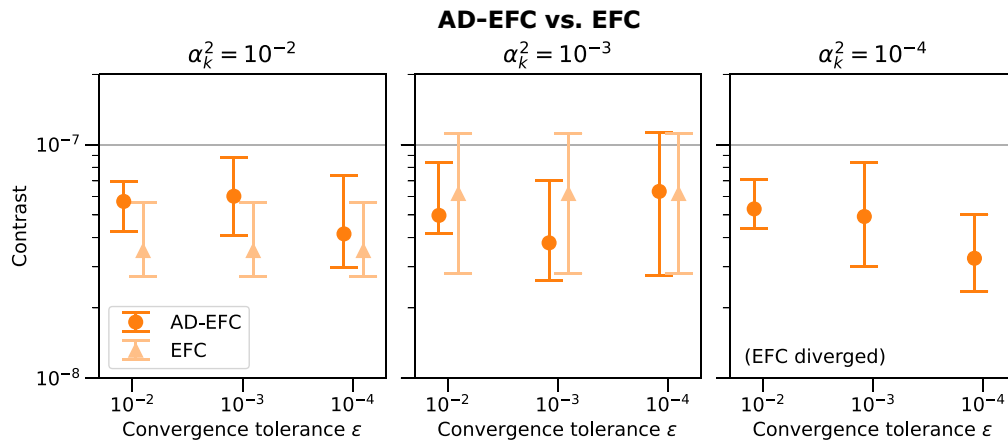


Fig. 12 Median, 10th percentile, and 90th percentile spatially averaged contrast values achieved by AD-EFC (dark orange) as a function of optimizer tolerance, for three different values of the Tikhonov regularization parameter α_k . For each α_k , we also ran a reference experiment with EFC (light orange); the EFC experiment with $\alpha_k^2 = 10^{-4}$ diverged and is not shown. The contrast versus iteration time series for the rightmost result in the left pane ($\alpha_k^2 = 10^{-2}$, $\varepsilon = 10^{-4}$) is illustrated in Fig. 10. The performance of AD-EFC was statistically equivalent to EFC for $\alpha_k^2 = 10^{-2}$ and $\alpha_k^2 = 10^{-3}$. As discussed in Sec. 4.1, the optimizer tolerance prevented AD-EFC from diverging in this case.

larger time interval between DM updates. In principle, on a system, such as HiCAT, where environmental disturbances cause the aberrated electric field to evolve over time scales on the order of seconds or faster, an excessive delay between the estimation step and application of the DM correction can cause a degradation in achievable contrast. However, we observed no meaningful degradation for smaller values of ε . In a real space-borne system, this is unlikely to be a significant consideration because of the much greater electric field stability, and because the total duration of each WFS&C iteration will be dominated by the exposure times needed for the estimation step.

The value ε can also serve as an auxiliary form of regularization, by terminating the optimization algorithm before it reaches an overly aggressive DM correction caused by a noisy electric field estimate, insufficient regularization using α_k or η_{00} , or both. For instance, in Fig. 12, with $\alpha_k^2 = 10^{-4}$ (rightmost panel), EFC diverged altogether whereas AD-EFC did not. On the other hand, choosing ε too large can impose an effective contrast floor by limiting the ability of

the optimization algorithm to converge to appropriately strong corrections. We determined in simulation that this was the case for $\varepsilon > 10^{-2}$.

For EFC and AD-EFC, smaller values of the Tikhonov regularization α_k correspond to more aggressive correction of the electric field, with the tradeoff of increased sensitivity to small perturbations in the estimated electric field, which reduces the stability of the WFS&C loop. For AD-PSM and SM, the aggressiveness of the WFS&C control loop is set first and foremost by the target contrast level $I_{\text{target},k}$. For SM, a smaller $I_{\text{target},k}$ corresponds to a larger value of the Lagrange multiplier μ_k , and therefore more aggressive correction (recalling from Sec. 2.2 that $\mu_k = 1/\alpha_k^2$). For AD-PSM, as η_{00} tends toward infinity, the goal of achieving the target contrast is enforced more strongly; η_{00} is interpretable as tuning the aggressiveness of the control loop up to a maximum level imposed by $I_{\text{target},k}$. Our experiments indicated that over the range of values considered, the performance of the WFS&C loop was insensitive to the value of η_{00} .

In principle, because α_k and ε have similar effects on the performance of AD-EFC, there potentially exists a single combination of the two parameters that is optimal in terms of contrast, or perhaps a continuum of combinations with similar performance, with a change in α_k compensated by a change to ε in the opposite direction. The same holds true for AD-PSM.

5 Conclusions and Future Work

In this paper, we reported the first experimental demonstrations of two AD-based wavefront control algorithms, AD-PSM and AD-EFC, using the HiCAT testbed. To within statistical uncertainty, AD-PSM and AD-EFC equaled their Jacobian-based counterparts in dark-zone contrast for all combinations of parameters that we tested. These demonstrations pave the way for future experimental validation at higher contrast.

The analysis in our earlier work indicated that the largest computational gains are realized for DMs with more than 64×64 actuators. The DMs currently in use on HiCAT have 34 actuators across the diameter of the active region or 952 actuators per DM in total, which is comparatively low. Therefore, our goal was to validate the fundamental capability of AD-PSM and AD-EFC to reach deep contrast, rather than to demonstrate improved computational efficiency.

In this work and in our earlier work, we utilized the L-BFGS optimization algorithm to minimize the wavefront control cost function because of its desirable convergence properties compared to first-order optimization methods and low storage requirements. Despite this, L-BFGS is known to converge slowly for poorly conditioned problems compared to methods that utilize the exact Hessian matrix.⁹ However, there exist alternative methods, such as truncated Newton algorithms, with excellent convergence properties for quadratic or nearly-quadratic cost functions, that require only the ability to evaluate Hessian-vector products, rather than the Hessian matrix itself.⁹ Hessian-vector products can be evaluated using AD in a similar fashion to gradients. Future work will explore such methods as a potentially faster approach than L-BFGS.

6 Appendix A: Convergence Data for All Experiments

In Sec. 4, we showed the time series of spatially averaged contrast vs. WFS&C iteration for an AD-PSM with $\eta_{00} = 10$ and $\varepsilon = 10^{-4}$ compared to an experiment with SM (Fig. 9), as well as an AD-EFC experiment with $\alpha_k^2 = 10^{-2}$ and $\varepsilon = 10^{-4}$ compared to an EFC experiment with the same value of α_k^2 (Fig. 10). We then summarized the steady-state time series statistics (median, 10th percentile, and 90th percentile) for a series of nine AD-PSM and nine AD-EFC runs with different combinations of (η_{00}, ε) and $(\alpha_k^2, \varepsilon)$, respectively (Figs. 11 and 12). In this section, Fig. 13 shows the full time series of contrast versus iteration for all 18 AD-PSM and AD-EFC experiments compared to their respective SM and EFC reference experiments.

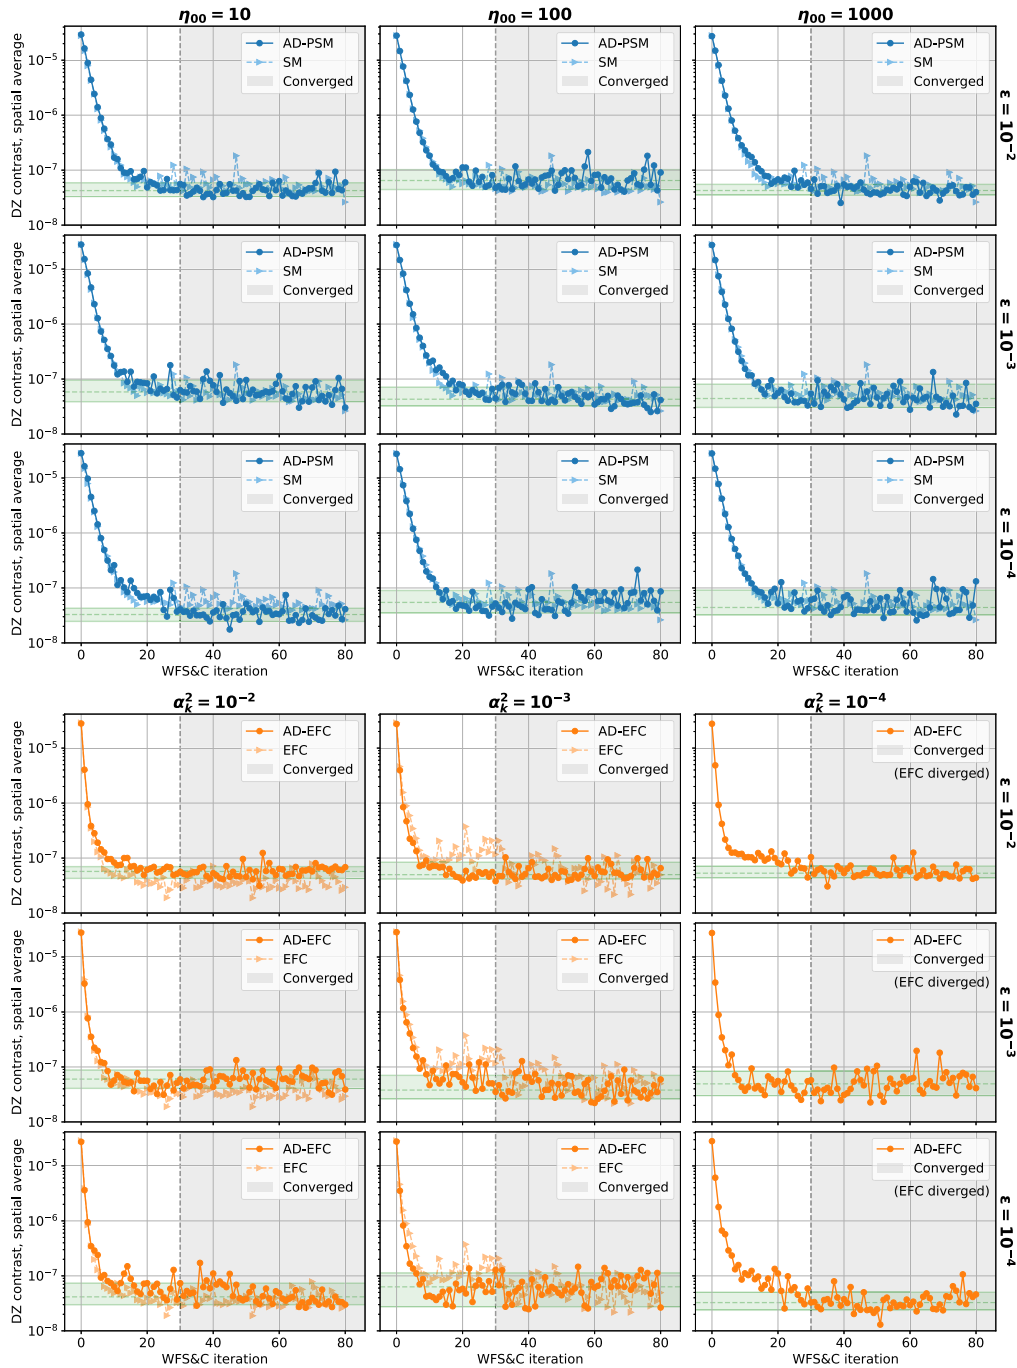


Fig. 13 Spatially averaged dark-zone contrast versus iteration for all experiments discussed in Sec. 4 and summarized in Figs. 11 and 12. The 10th percentile, median, and 90th percentile of the steady-state contrast values (gray region) for AD-PSM and AD-EFC are shown in green. The AD-PSM experiment with ($\eta_{00} = 10, \epsilon = 10^{-4}$) and the AD-EFC experiment with ($\alpha_k^2 = 10^{-2}, \epsilon = 10^{-4}$) are also shown in Figs. 9 and 10, respectively.

7 Appendix B: Equivalence of Jacobian-Based Solutions and Newton's Method

Finding solutions for EFC and SM using the Jacobian matrix is equivalent to minimizing their respective cost functions with respect to \mathbf{a}_k using Newton's method. Newton's method is a second-order optimization technique that utilizes second-derivative information about the cost

function, given by the local Hessian matrix $\mathbf{H}_k(\mathbf{a}_k^n)$ at any point \mathbf{a}_k^n in the DM command parameter space. Given an initial guess for the solution \mathbf{a}_k^0 , the full Newton update is given as⁹

$$\mathbf{a}_k^1 = \mathbf{a}_k^0 - \mathbf{H}_k^{-1}(\mathbf{a}_k^0) \left. \frac{\partial J_k}{\partial \mathbf{a}_k^T} \right|_{\mathbf{a}_k = \mathbf{a}_k^0}. \quad (17)$$

For cost functions that are exactly quadratic, including EFC and SM, Newton's method converges in a single iteration.

For general numerical optimization problems, Newton's method is rarely used in practice because forming the Hessian matrix explicitly is expensive. On the other hand, quasi-Newton methods, such as the BFGS algorithm or the limited-memory BFGS (L-BFGS) variant, can approximate \mathbf{H}_k^{-1} using changes in $\partial J_k / \partial \mathbf{a}_k^T$ over successive optimization iterations. As a consequence, they are substantially less computationally expensive. Although quasi-Newton algorithms do not converge as rapidly as Newton's method, and in particular can converge slowly for poorly conditioned problems, they are nonetheless superior to purely first-order methods such as steepest descent.⁹

As we show below, the Hessian matrix for EFC and SM has an analytic expression in terms of the Jacobian \mathbf{G}_k given by $\mathbf{H}_k = 2(\text{Re}\{\mathbf{G}_k^\dagger \mathbf{G}_k\} + \mathbf{C})$, where \mathbf{C} is a symmetric, positive-definite matrix given by $\mathbf{\Gamma}_k^T \mathbf{\Gamma}_k$ for EFC and \mathbb{I} / μ_k^n for SM. Our approach is to replace this full Newton iteration, requiring computation of the Jacobian, by a series of cheaper quasi-Newton iterations instead, requiring only computation of the gradient $\partial J_k / \partial \mathbf{a}_k^T$, which we achieve using AD.

The cost function for the EFC algorithm described in Sec. 2 can be written in the form

$$J_k(\mathbf{a}_k) = \|\mathbf{E}_{\text{DZ},k}(\mathbf{a}_k)\|^2 + \|\mathbf{\Gamma}_k \mathbf{a}_k\|^2, \quad (18)$$

where $\mathbf{\Gamma}_k$ is a regularization matrix; for SM, $\mathbf{\Gamma}_k = \mathbb{I} / \sqrt{\mu_k}$. For now, we will restrict our attention to the EFC algorithm, but note that the result derived here applies equally as well to minimizing the Lagrangian function for SM with respect to the DM correction \mathbf{a}_k .

Expanding Eq. (18) and recalling that $\mathbf{E}_{\text{DZ},k} = \mathbf{G}_k \mathbf{a}_k + \hat{\mathbf{E}}_{\text{ab},k}$, the cost function has the form

$$J_{\text{EFC},k}(\mathbf{a}_k) = \mathbf{a}_k^T (\text{Re}\{\mathbf{G}_k^\dagger \mathbf{G}_k\} + \mathbf{\Gamma}_k^T \mathbf{\Gamma}_k) \mathbf{a}_k + 2\mathbf{a}_k^T \text{Re}\{\mathbf{G}_k^\dagger \hat{\mathbf{E}}_{\text{ab},k}\} + \hat{\mathbf{E}}_{\text{ab},k}^\dagger \hat{\mathbf{E}}_{\text{ab},k}, \quad (19)$$

where we use the fact that \mathbf{a}_k is purely real to discard $\text{Im}\{\mathbf{G}_k^\dagger \mathbf{G}_k\}$. The Jacobian-based solution is found by finding \mathbf{a}_k such that $\partial J_{\text{EFC},k} / \partial \mathbf{a}_k$ vanishes. We therefore begin by writing down the gradient:

$$\frac{\partial J_{\text{EFC},k}}{\partial \mathbf{a}_k^T} = 2(\text{Re}\{\mathbf{G}_k^\dagger \mathbf{G}_k\} + \mathbf{\Gamma}_k^T \mathbf{\Gamma}_k) \mathbf{a}_k + 2 \text{Re}\{\mathbf{G}_k^\dagger \hat{\mathbf{E}}_{\text{ab},k}\} = 0. \quad (20)$$

This is a linear system of equations that we can solve for the optimal correction \mathbf{a}_k^* :

$$\mathbf{a}_k^* = -(\text{Re}\{\mathbf{G}_k^\dagger \mathbf{G}_k\} + \mathbf{\Gamma}_k^T \mathbf{\Gamma}_k)^{-1} \text{Re}\{\mathbf{G}_k^\dagger \hat{\mathbf{E}}_{\text{ab},k}\}. \quad (21)$$

The Hessian matrix is given as

$$\mathbf{H}_k = \frac{\partial^2 J_{\text{EFC},k}}{\partial \mathbf{a}_k \partial \mathbf{a}_k^T} = 2(\text{Re}\{\mathbf{G}_k^\dagger \mathbf{G}_k\} + \mathbf{\Gamma}_k^T \mathbf{\Gamma}_k). \quad (22)$$

Since both terms in \mathbf{H}_k are positive definite, \mathbf{H}_k is positive definite as well, confirming that the solution is a minimum of the cost function.

We will now show that the solution \mathbf{a}_k^* obtained above is the same as the solution obtained by applying a single iteration of Newton's method to the EFC cost function. Let \mathbf{a}_k^0 be an initial guess for the solution. Newton's method produces an iterate of the form⁹

$$\mathbf{a}_k^1 = \mathbf{a}_k^0 - \mathbf{H}_k^{-1} \left. \frac{\partial J_{\text{EFC},k}}{\partial \mathbf{a}_k^T} \right|_{\mathbf{a}_k = \mathbf{a}_k^0}. \quad (23)$$

Combining Eqs. (20) and (22)

$$\left. \frac{\partial J_{\text{EFC},k}}{\partial \mathbf{a}_k^T} \right|_{\mathbf{a}_k = \mathbf{a}_k^0} = \mathbf{H}_k \mathbf{a}_k^0 + 2 \text{Re}\{\mathbf{G}_k^\dagger \hat{\mathbf{E}}_{\text{ab},k}\}. \quad (24)$$

Inserting back into Eq. (23)

$$\mathbf{a}_k^1 = \mathbf{a}_k^0 - \mathbf{H}_k^{-1}(\mathbf{H}_k \mathbf{a}_k^0 + 2 \text{Re}\{\mathbf{G}_k^\dagger \hat{\mathbf{E}}_{\text{ab},k}\}), \quad (25a)$$

$$= \mathbf{a}_k^0 - \mathbf{H}_k^{-1} \mathbf{H}_k \mathbf{a}_k^0 - 2 \mathbf{H}_k^{-1} \text{Re}\{\mathbf{G}_k^\dagger \hat{\mathbf{E}}_{\text{ab},k}\}, \quad (25b)$$

$$= \mathbf{a}_k^0 - \mathbf{a}_k^0 - 2 \mathbf{H}_k^{-1} \text{Re}\{\mathbf{G}_k^\dagger \hat{\mathbf{E}}_{\text{ab},k}\}, \quad (25c)$$

$$= -2 \mathbf{H}_k^{-1} \text{Re}\{\mathbf{G}_k^\dagger \hat{\mathbf{E}}_{\text{ab},k}\}. \quad (25d)$$

Inserting the definition of the Hessian matrix from Eq. (22), we see that the Newton iterate \mathbf{a}_k^1 is identical to the analytical solution in Eq. (21)

$$\mathbf{a}_k^1 = -2[2(\text{Re}\{\mathbf{G}_k^\dagger \mathbf{G}_k\} + \mathbf{\Gamma}_k^T \mathbf{\Gamma}_k)]^{-1} \text{Re}\{\mathbf{G}_k^\dagger \hat{\mathbf{E}}_{\text{ab},k}\}, \quad (26a)$$

$$= -(\text{Re}\{\mathbf{G}_k^\dagger \mathbf{G}_k\} + \mathbf{\Gamma}_k^T \mathbf{\Gamma}_k)^{-1} \text{Re}\{\mathbf{G}_k^\dagger \hat{\mathbf{E}}_{\text{ab},k}\}, \quad (26b)$$

$$= \mathbf{a}_k^*. \quad (26c)$$

As we described earlier, the same result holds if minimizing $\mathcal{L}_{\text{SM},k}$ in Sec. 2.1 with respect to \mathbf{a}_k .

8 Appendix C: Fast Convolutional DM Model

Consider a DM with N_A actuators along each side (i.e., $N_{\text{act}} = N_A^2$) whose surface $s(x, y)$ can be modeled as a linear superposition of identical influence functions $f(x, y)$

$$s(x, y) = \sum_{m=1}^{N_A} \sum_{n=1}^{N_A} a_{m,n} f(x - x_m, y - y_n). \quad (27)$$

For fixed actuator spacing along the horizontal and vertical directions, we can rewrite the above summation as a convolution between a weighted Dirac comb function and the influence function

$$s(x, y) = f(x, y) * \sum_{m=1}^{N_A} \sum_{n=1}^{N_A} a_{m,n} \delta(x - x_m, y - y_n). \quad (28)$$

Fourier transforming both sides transforms the convolution operation into a multiplication:

$$\mathcal{F}\{s(x, y)\} = \mathcal{F}\{f(x, y)\} \sum_{m=1}^{N_A} \sum_{n=1}^{N_A} a_{m,n} \mathcal{F}\{\delta(x - x_m, y - y_n)\}, \quad (29a)$$

$$= \mathcal{F}\{f(x, y)\} \sum_{m=1}^{N_A} \sum_{n=1}^{N_A} a_{m,n} \exp\{-i2\pi(x_m f_x + y_n f_y)\}. \quad (29b)$$

We define $\tilde{s}(f_x, f_y) \triangleq \mathcal{F}\{s(x, y)\}$ and $\tilde{f}(f_x, f_y) \triangleq \mathcal{F}\{f(x, y)\}$:

$$\tilde{s}(f_x, f_y) = \tilde{f}(f_x, f_y) \sum_{m=1}^{N_A} \sum_{n=1}^{N_A} a_{m,n} \exp\{-i2\pi(x_m f_x + y_n f_y)\}. \quad (30)$$

We next define the discretized surface and influence function arrays $\tilde{\mathbf{s}}$ and $\tilde{\mathbf{f}}$ such that

$$\tilde{\mathbf{s}}[p, q] = \tilde{s}(p\Delta f_x, q\Delta f_y), \quad (31a)$$

$$\tilde{\mathbf{f}}[p, q] = \tilde{f}(p\Delta f_x, q\Delta f_y), \quad (31b)$$

so that

$$\tilde{\mathbf{s}} = \tilde{\mathbf{f}} \circ \sum_{m=1}^{N_A} \sum_{n=1}^{N_A} a_{m,n} \exp\{-i2\pi(x_m \mathbf{f}_x + y_n \mathbf{f}_y)\}, \quad (32)$$

where \circ denotes element-wise multiplication. Finally, we define the vectors of actuator center coordinates \mathbf{x}_c and \mathbf{y}_c such that $\mathbf{x}_c[m] = x_m$ and $\mathbf{y}_c[n] = y_n$, as well as the array of actuator commands \mathbf{A} for which $\mathbf{A}[m, n] = a_{m,n}$

$$\tilde{\mathbf{s}}[p, q] = \tilde{\mathbf{f}}[p, q] \circ \sum_{m=1}^{N_A} \sum_{n=1}^{N_A} \mathbf{A}[m, n] \exp\{-i2\pi(\mathbf{x}_c[m] \mathbf{f}_x[p] + \mathbf{y}_c[n] \mathbf{f}_y[q])\}, \quad (33a)$$

$$= \tilde{\mathbf{f}}[p, q] \circ \sum_{m=1}^{N_A} \exp\{-i2\pi \mathbf{x}_c[m] \mathbf{f}_x[p]\} \sum_{n=1}^{N_A} \mathbf{A}[m, n] \exp\{-i2\pi \mathbf{y}_c[n] \mathbf{f}_y[q]\}. \quad (33b)$$

We can write this more succinctly as the following sequence of element-wise and matrix products

$$\tilde{\mathbf{s}} = \tilde{\mathbf{f}} \circ (\exp\{-i2\pi \mathbf{f}_x \mathbf{x}_c^T\} \mathbf{A} \exp\{-i2\pi \mathbf{y}_c \mathbf{f}_y^T\}), \quad (34)$$

where exponentiation is performed element-wise and $\mathbf{a}\mathbf{b}^T$ denotes the outer product of the vectors \mathbf{a} and \mathbf{b} .

The term in the parentheses is more commonly referred to as the matrix Fourier transform²¹ or matrix triple product Fourier transform²⁵ of \mathbf{A} , which we denote as follows:

$$\text{MFT}\{\mathbf{A}; \mathbf{x}_c, \mathbf{y}_c, \mathbf{f}_x, \mathbf{f}_y\} \triangleq \exp\{-i2\pi \mathbf{f}_x \mathbf{x}_c^T\} \mathbf{A} \exp\{-i2\pi \mathbf{y}_c \mathbf{f}_y^T\}, \quad (35)$$

yielding

$$\tilde{\mathbf{s}} = \tilde{\mathbf{f}} \circ \text{MFT}\{\mathbf{A}; \mathbf{x}_c, \mathbf{y}_c, \mathbf{f}_x, \mathbf{f}_y\}. \quad (36)$$

The final step is to compute an inverse discrete Fourier transform to obtain the desired discrete DM surface \mathbf{s} , which is carried out most efficiently using the inverse fast Fourier transform, yielding the final result

$$\mathbf{s} = \text{IFFT}\{\tilde{\mathbf{f}} \circ \text{MFT}\{\mathbf{A}; \mathbf{x}_c, \mathbf{y}_c, \mathbf{f}_x, \mathbf{f}_y\}\}. \quad (37)$$

For DMs whose active actuators are a subset of the $N_A \times N_A$ grid modeled above, only the elements of \mathbf{A} corresponding to active actuators are set to nonzero values.

8.1 C.1 Adjoint Model

The algorithm described in the previous section computes the DM surface resulting from a two-dimensional array of actuator commands \mathbf{A} , under the assumptions that the influence function is identical across all actuators and that the surface can be approximated as a linear superposition of the actuator influence functions. In the context of gradient-based nonlinear optimization using RMAD, we can derive an adjoint model for this algorithm that computes the derivative $\bar{\mathbf{A}} \triangleq \partial J / \partial \mathbf{A}^T$ for some scalar cost function J , given the derivative $\bar{\mathbf{s}}$ with respect to the surface \mathbf{s} .

To begin, we break the forward model into the following sequence:

$$\tilde{\mathbf{A}} = \text{MFT}\{\mathbf{A}; \mathbf{x}_c, \mathbf{y}_c, \mathbf{f}_x, \mathbf{f}_y\}, \quad (38a)$$

$$\tilde{\mathbf{s}} = \tilde{\mathbf{f}} \circ \tilde{\mathbf{A}}, \quad (38b)$$

$$\mathbf{s} = \text{IFFT}\{\tilde{\mathbf{s}}\}. \quad (38c)$$

This leads to the following adjoint model, following the RMAD adjoint variable rules in Refs. 7 and 26:

$$\bar{\bar{\mathbf{s}}} = \text{FFT}\{\bar{\mathbf{s}}\}, \quad (39a)$$

$$\bar{\bar{\mathbf{A}}} = \bar{\bar{\mathbf{f}}}^* \circ \bar{\bar{\mathbf{s}}}, \quad (39b)$$

$$\bar{\mathbf{A}} = \text{IMFT}\{\bar{\bar{\mathbf{A}}}; \mathbf{f}_x, \mathbf{f}_y, \mathbf{x}_c, \mathbf{y}_c\}, \quad (39c)$$

where $*$ denotes element-wise complex conjugation and IMFT denotes the inverse matrix Fourier transform:

$$\text{IMFT}\{\bar{\mathbf{A}}; \mathbf{f}_x, \mathbf{f}_y, \mathbf{x}_c, \mathbf{y}_c\} \triangleq \exp\{i2\pi\mathbf{x}_c\mathbf{f}_x^T\}\bar{\mathbf{A}}\exp\{i2\pi\mathbf{f}_y\mathbf{y}_c^T\}. \quad (40)$$

Combining these expressions, the adjoint model is then

$$\bar{\mathbf{A}} = \text{IMFT}\{\bar{\mathbf{f}}^* \circ \text{FFT}\{\bar{\mathbf{s}}\}; \mathbf{f}_x, \mathbf{f}_y, \mathbf{x}_c, \mathbf{y}_c\}. \quad (41)$$

9 Appendix D: Adjoint Model for EFC Cost Function

In Sec. 2.2, we describe the cost function for the EFC algorithm for a single correction wavelength. Here, we derive its RMAD adjoint model, which computes the derivative $\partial J_{\text{EFC},k}/\partial \mathbf{E}_{\text{DM},k}$.

We begin by writing the cost function as a series of operations evaluated sequentially:

$$\mathbf{E}_{\text{DZ},k} = \mathbf{E}_{\text{DM},k} + \hat{\mathbf{E}}_{\text{ab},k}, \quad (42a)$$

$$J_{\mathbf{E}_{\text{DZ},k}} = \|\mathbf{E}_{\text{DZ},k}\|^2, \quad (42b)$$

$$\mathbf{c}_k = \mathbf{\Gamma}_k \mathbf{a}_k, \quad (42c)$$

$$J_{\mathbf{c},k} = \|\mathbf{c}_k\|^2, \quad (42d)$$

$$J_{\text{EFC},k} = J_{\mathbf{E}_{\text{DZ},k}} + J_{\mathbf{c},k}. \quad (42e)$$

We now apply the RMAD gradient rules^{7,26} to each step in reverse order to derive the adjoint model, letting $\bar{\mathbf{x}} \triangleq \partial J_k / \partial \mathbf{x}^T$ for any variable \mathbf{x} :

$$\bar{J}_{\text{EFC},k} = 1, \quad (43a)$$

$$\bar{J}_{\mathbf{E}_{\text{DZ},k}} = \bar{J}_{\mathbf{c},k} = \bar{J}_{\text{EFC},k}, \quad (43b)$$

$$\bar{\mathbf{c}}_k = 2\mathbf{c}_k \bar{J}_{\mathbf{c},k}, \quad (43c)$$

$$\bar{\mathbf{a}}_k = \mathbf{\Gamma}_k^T \bar{\mathbf{c}}_k, \quad (43d)$$

$$\bar{\mathbf{E}}_{\text{DZ},k} = 2\mathbf{E}_{\text{DZ},k} \bar{J}_{\mathbf{E}_{\text{DZ},k}}, \quad (43e)$$

$$\bar{\mathbf{E}}_{\text{DM},k} = \bar{\mathbf{E}}_{\text{DZ},k}. \quad (43f)$$

Combining and simplifying, we see that the desired gradient is given as

$$\bar{\mathbf{E}}_{\text{DM},k} = 2\Delta \mathbf{E}_k. \quad (44)$$

This gradient is passed to the next block of the adjoint model, which in this case is the adjoint model for propagation through the coronagraph, which, referring to Fig. 4, evaluates the derivatives of $J_{\text{EFC},k}$ with respect to the surfaces $\mathbf{s}_{1,k}$ and $\mathbf{s}_{2,k}$, respectively. We refer to our earlier work for a derivation of the coronagraph propagation adjoint model.⁷

Disclosures

The authors declare no conflicts of interest.

Code and Data Availability

NASA regulations govern the release of data and source code. Please contact Scott Will at scott.d.will@nasa.gov to request supporting materials.

HiCAT makes use of the NumPy,^{27,28} Matplotlib,^{29,30} AstroPy,³¹ SciPy,²⁴ scikit-image,³² Pandas,^{33,34} ImageIO,³⁵ Photutils,³⁶ HCIpy,³⁷ POPPY,^{22,23} and CatKit³⁸ packages.

Acknowledgments

The authors are especially thankful to the extended HiCAT team (over 50 people) who have worked over the past several years to develop this testbed. This work was supported in part by the National Aeronautics and Space Administration (Grant No. 80NSSC19K0120) issued through the Strategic Astrophysics Technology/Technology Demonstration for Exoplanet Missions Program (SAT-TDEM; PI: R. Soummer). E.H.P. was supported by the NASA Hubble Fellowship

(Grant No. HST-HF2-51467.001-A) awarded by the Space Telescope Science Institute, which is operated by the Association of Universities for Research in Astronomy, Incorporated, under NASA (Grant No. NAS5-26555). I.L. acknowledges the support by a postdoctoral grant issued by the Centre National d'Études Spatiales (CNES) in France. This paper was adapted from an earlier SPIE conference proceedings.³⁹

References

1. M. Perryman, *The Exoplanet Handbook*, 2nd ed., Cambridge University Press (2018).
2. L. Pueyo et al., “Optimal dark hole generation via two deformable mirrors with stroke minimization,” *Appl. Opt.* **48**, 6296–6312 (2009).
3. A. Give'ón et al., “Closed loop, DM diversity-based, wavefront correction algorithm for high contrast imaging systems,” *Opt. Express* **15**(19), 12338–12343 (2007).
4. B.-J. Seo et al., “Testbed demonstration of high-contrast coronagraph imaging in search for Earth-like exoplanets,” *Proc. SPIE* **11117**, 111171V (2019).
5. The LUVOIR Study Team, “LUVOIR,” tech. rep., National Aeronautics and Space Administration (2019).
6. The Habitable Exoplanet Observatory Study Team, “Habitable Exoplanet Observatory Final Report,” tech. rep., Jet Propulsion Laboratory (2019).
7. S. D. Will, T. D. Groff, and J. R. Fienup, “Jacobian-free coronagraphic wavefront control using nonlinear optimization,” *J. Astron. Telesc. Instrum. Syst.* **7**, 019002 (2021).
8. A. Griewank and A. Walther, *Evaluating Derivatives: Principles and Techniques of Algorithmic Differentiation*, 2nd ed., Society for Industrial and Applied Mathematics (2008).
9. J. Nocedal and S. J. Wright, *Numerical Optimization*, Springer-Verlag, New York (1999).
10. M. N'Diaye et al., “High-contrast imager for complex aperture telescopes (HiCAT): 1. testbed design,” *Proc. SPIE* **8864**, 88641K (2013).
11. M. N'Diaye et al., “High-contrast imager for complex aperture telescopes (HiCAT): II. Design overview and first light results,” *Proc. SPIE* **9143**, 914327 (2014).
12. M. N'Diaye et al., “High-contrast imager for complex aperture telescopes (HiCAT): 3. First lab results with wavefront control,” *Proc. SPIE* **9605**, 96050I (2015).
13. L. Leboulleux et al., “High-contrast imager for Complex Aperture Telescopes (HiCAT). 4. Status and wavefront control development,” *Proc. SPIE* **9904**, 99043C (2016).
14. R. Soummer et al., “High-contrast imager for complex aperture telescopes (HiCAT): 5. First results with segmented-aperture coronagraph and wavefront control,” *Proc. SPIE* **10698**, 106981O (2018).
15. C. Moriarty, et al., “High-contrast imager for complex aperture telescopes (HiCAT): 6. software control infrastructure and calibration,” *Proc. SPIE* **10698**, 1069853 (2018).
16. R. Soummer et al., “High-contrast imager for complex aperture telescopes (HiCAT): 6. Two deformable mirror wavefront control (Conference Presentation),” *Proc. SPIE* **11117**, 111171Y (2019).
17. R. Soummer et al., “High-contrast imager for complex aperture telescopes (HiCAT): 7. Dark zone demonstration with fully segmented aperture coronagraph,” *Proc. SPIE* **11823**, 118230T (2021).
18. R. Pourcelot et al., “Experimental validation of active control of low-order aberrations with a Zernike sensor through a Lyot coronagraph,” *Proc. SPIE* **11823**, 118231M (2021).
19. I. Lajinja et al., “Wavefront tolerances of space-based segmented telescopes at very high contrast: Experimental validation,” *Astron. Astrophys.* **658**, A84 (2022).
20. S. M. Redmond et al., “Dark hole maintenance results for segmented aperture wavefront error drift in a high contrast space coronagraph,” *Proc. SPIE* **11823**, 118231K (2021).
21. R. Soummer et al., “Fast computation of Lyot-style coronagraph propagation,” *Opt. Express* **15**, 15935–15951 (2007).
22. M. D. Perrin et al., “Simulating point spread functions for the James Webb Space Telescope with WebbPSF,” *Proc. SPIE* **8442**, 84423D (2012).
23. M. Perrin et al., “POPPY: Physical Optics Propagation in PYthon,” Astrophysics Source Code Library Record, <https://ascl.net/1602.018> (2016).
24. P. Virtanen et al., “SciPy 1.0: fundamental algorithms for scientific computing in Python,” *Nat. Methods* **17**, 261–272 (2020).
25. A. S. Jurling, M. D. Bergkoetter, and J. R. Fienup, “Techniques for arbitrary sampling in two-dimensional Fourier transforms,” *J. Opt. Soc. Am. A* **A35**, 1784–1796 (2018).
26. A. S. Jurling and J. R. Fienup, “Applications of algorithmic differentiation to phase retrieval algorithms,” *J. Opt. Soc. Am. A* **31**(7), 1348–59 (2014).
27. T. E. Oliphant, *A Guide to NumPy*, Vol. **1**, Trelgol Publishing, USA (2006).
28. S. van der Walt, S. C. Colbert, and G. Varoquaux, “The NumPy array: a structure for efficient numerical computation,” *Comput. Sci. Eng.* **13**, 22–30 (2011).
29. J. D. Hunter, “Matplotlib: a 2D graphics environment,” *Comput. Sci. Eng.* **9**, 90–95 (2007).
30. T. A. Caswell et al., “matplotlib/matplotlib: Rel. v3.3.3,” (2020).

31. Astropy Collaboration et al., “The astropy project: building an open-science project and status of the v2.0 core package,” *Astrophys. J.* **156**, 123 (2018).
32. S. van der Walt et al., “scikit-image: image processing in Python,” *PeerJ* **2**, e453 (2014).
33. Wes McKinney, “Data structures for statistical computing in python,” in *Proc. 9th Python in Sci. Conf.*, S. van der Walt and J. Millman, Eds., pp. 56–61 (2010).
34. J. Reback et al., “pandas-dev/pandas: Pandas 1.1.4,” (2020).
35. S. Silvester et al., “imageio/imageio v2.8.0,” (2020).
36. L. Bradley et al., “astropy/photutils: 1.0.1,” (2020).
37. E. H. Por et al., “High Contrast Imaging for Python (HCIPy): an open-source adaptive optics and coronagraph simulator,” *Proc. SPIE* **10703**, 1070342 (2018).
38. J. Noss et al., “spacetelescope/catkit: v0.36.1,” (2021).
39. S. D. Will et al., “Wavefront control with algorithmic differentiation on the HiCAT testbed,” *Proc. SPIE* **11823**, 118230V (2021).

Scott D. Will is an optical engineer at NASA Goddard Space Flight Center in Greenbelt, Maryland, USA. He received a BS degree in electrical engineering from the University at Buffalo in 2015 and a PhD in optics from the University of Rochester in 2021. As a graduate student, he was also a member of the Russell B. Makidon Optics Laboratory team at the Space Telescope Science Institute in Baltimore, Maryland, United States. His research interests include wavefront sensing and control, computational imaging, and astronomical instrumentation.

Emiel H. Por is an NHFP Sagan fellow at the Russell B. Makidon Optics Laboratory at the Space Telescope Science Institute in Baltimore, United States. He received his BSc degree in astronomy and physics, his MSc degree in astronomical instrumentation, and his PhD in astronomy from Leiden University, The Netherlands. His research interests include coronagraphy, wavefront sensing and control, and high-contrast imaging.

Ananya Sahoo is a staff scientist at STScI. She currently splits her time between infrastructure development for the HiCAT testbed and studying wavefront-error induced by thermal disturbances for future segmented space telescopes. She received her PhD in 2020 from The Graduate University for Advanced Studies, SOKENDAI in alliance with National Astronomical Observatory of Japan. She previously worked on the photometric and astrometric calibration for the SCEXAO instrument at the Subaru Telescope.

Iva Laginja is a postdoctoral fellow with CNES, the French space agency and previous member, now collaborator, of the Russell B. Makidon Optics Laboratory at the Space Telescope Science Institute in Baltimore, United States. She received an MSc degree in astronomy and instrumentation at Leiden University, Netherlands and her doctorate from Paris Observatory in 2021. Her research interests include wavefront sensing and control, high contrast imaging, and coronagraphy for exoplanet detection.

Raphaël Pourcelot is a postdoc at STScI within the Russell B. Makidon Optics Laboratory. He received an MSc degree in astronomy from Aix-Marseille Université and an engineering degree at Institut d’Optique Graduate School in Université Paris-Saclay. He received his PhD from Université Côte d’Azur in Nice, France, in 2022. He is mainly working on wavefront sensing and control for high-contrast imaging for exoplanet detection and Zernike wavefront sensor applications in particular.

Susan F. Redmond is a PhD student at Princeton University in Princeton, United States, and collaborates with the Russell B. Makidon Optics Laboratory. She previously obtained her BEng degree from Memorial University of Newfoundland (2015) and her MEng and MASc degrees in aerospace engineering from the University of Toronto in 2016 and 2018, respectively. Her current research is split between the optical and thermal design for balloon-borne telescopes and developing focal-plane wavefront sensing and control algorithms for high-contrast imaging.

Laurent Pueyo is an associate astronomer at STScI. He received his doctorate from Princeton University in 2008 and previously worked as a NASA Fellow at JPL, California and as a Sagan Fellow at JHU. His research focuses on imaging faint planets around nearby stars. He has pioneered optical technologies that allow astronomers to take images of other planetary systems, and has developed data analysis methods now standardly used to study extrasolar planets.

Tyler D. Groff received his BS degree in mechanical engineering and astrophysics from Tufts University and his PhD in mechanical and aerospace engineering from Princeton University under an NESSF fellowship. He was the lab manager of the Princeton High Contrast Imaging Laboratory and the CHARIS instrument at Subaru Telescope. He is currently the lead engineer at Goddard Space Flight Center for the spectroscopy and polarization modes on the Roman Space Telescope Coronagraph Instrument.

James R. Fienup received his AB degree from Holy Cross College and MS and PhD degrees in applied physics from Stanford University, where he was a National Science Foundation graduate fellow. After performing research at ERIM, he became the Robert E. Hopkins Professor of Optics at the University of Rochester. He is a fellow of SPIE and OSA, a member of the National Academy of Engineering, and a recipient of SPIE's Rudolf Kingslake Medal and Prize.

Remi Soummer is an associate astronomer at STScI. He received his doctorate from the University of Nice in 2002 and has been working in the field of high contrast imaging and instrumentation for the detection and characterization of exoplanets ever since. He is currently the head of the Russell B. Makidon Optics Laboratory and working on a coronagraph demonstration for future large segmented aperture space telescopes.

Biographies of the other authors are not available.