

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

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

^aNASA Goddard Space Flight Center, Greenbelt, MD 20771, USA

^bThe Institute of Optics, University of Rochester, Rochester, NY 14627, USA

^cSpace Telescope Science Institute, 3700 San Martin Drive, Baltimore MD 21218, USA

^dSigma Space, Lanham, MD 20706, USA

^eLESIA, Observatoire de Paris, Université PSL, Sorbonne Université, Université Paris Cité, CNRS, 5 place Jules Janssen, 92195 Meudon, France,

^fDepartment of Mechanical and Aerospace Engineering, Princeton University, Princeton, NJ 08540, USA

Abstract. Future space-based coronagraphs will rely critically on focal-plane wavefront sensing and control with deformable mirrors 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 deformable mirror 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 deformable mirror 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 (LUVOIR), 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 (HiCAT) testbed at the Space Telescope Science Institute (STScI) and present initial experimental results, demonstrating contrast suppression capabilities equivalent to Jacobian-based methods.

Keywords: coronagraphy, high-contrast imaging, wavefront control, optimization.

*Scott D. Will, scott.d.will@nasa.gov

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 the star and evolves slowly over time in response to miniscule drifts in the thermal and mechanical state of the

35 observatory. If uncorrected, the speckle floor overwhelms the faint image of the orbiting planet,
36 which is expected to be 10^{10} times fainter than the host star at 0.1 arcseconds of separation or less at
37 visible wavelengths.¹

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

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

52 As direct imaging missions demand DMs with ever-higher actuator density to enable wider
53 and wider search areas, computational aspects will inevitably become a point of concern from a
54 systems engineering standpoint. In on-orbit WFS&C, all sensing and control computations are
55 processed by the flight computer; conversely, in ground-in-the-loop (GITL) scenarios, raw data
56 is communicated to a ground-based computing node that calculates the DM correction and relays
57 the correction back to the observatory. Though each approach has tradeoffs, a major advantage of

58 on-orbit WFS&C is the ability to update DM commands more frequently without relying on the
59 continuous availability of communication links with the ground stations. On-orbit WFS&C can
60 help to relax observatory-level wavefront stability requirements by enabling the high-contrast dark
61 zone to be maintained over shorter time intervals. Successful deployment of an on-orbit architecture
62 is predicated on the availability of sufficient computational resources; however, radiation-hard,
63 space-qualified computing hardware lags behind conventional hardware by decades and is extremely
64 resource-limited. This poses a substantial computation capability gap if the current algorithms
65 are expected to be deployed on-orbit on a future direct imaging mission such as NASA’s planned
66 Habitable Worlds Observatory (HWO).

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

81 correspond to an increase by a factor of 32 (Habex) and 512 (LUVOIR-A) of the worst-case Jacobian
82 dimensionality, for each control wavelength, compared to DST.

83 In recent work, we formulated an alternative wavefront control framework that iteratively
84 compute DM updates using gradient-based optimization techniques, which eliminates the need to
85 calculate and manipulate the Jacobian.⁷ Our approach is based in part on a numerical technique
86 known as algorithmic differentiation (AD),⁸ to calculate the gradients of the cost function for
87 optimal control accurately and efficiently. We described an AD-based counterpart to SM, which we
88 termed algorithmic differentiation penalty stroke minimization (AD-PSM), and compared it to SM
89 in simulations with a small-angle LUVOIR design. Our results indicated superior computational
90 efficiency with AD-PSM and comparable starlight suppression performance. While the CPU time
91 and memory consumption of SM grew superlinearly with actuator count, the increase in both
92 for AD-PSM was negligible (e.g., with 128×128 actuators, AD-PSM utilized 95% less memory
93 and CPU time), suggesting that iterative methods are a promising alternative to Jacobian-based
94 techniques for on-orbit WFS&C with high actuator counts and large dark zones.

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

101 This paper is structured as follows. In Section 2, we review concepts from our earlier work,
102 including algorithmic differentiation and the mathematical formulation of AD-PSM and AD-EFC.
103 In Section 3, we provide an overview of HiCAT and discuss our experimental setup, including

104 algorithm implementation details that are pertinent to our demonstration. In Section 4, we present
 105 and discuss our experimental results. Finally, in Section 5, we provide our conclusions and briefly
 106 describe our planned future work.

107 *1.1 Notation*

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

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

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

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

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

118 **2 Wavefront control using algorithmic differentiation**

119 The goal of the WFS&C loop in coronagraphy is to drive starlight within the dark zone toward
 120 zero over time so that a faint orbiting companion becomes detectable against the reduced starlight
 121 background. Each iteration of the WFS&C loop, indexed by the integer k , consists of two steps: (1)
 122 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
 123 within the dark zone from focal-plane intensity measurements, and (2) a control step, in which the
 124 DM correction is updated to reduce the energy in $\mathbf{E}_{\text{ab},k}$, as shown in Figure 1. In this paper, we
 125 focus principally on the control step.

126 Modern model-based wavefront control algorithms find the DM correction update \mathbf{a}_k by mini-
 127 mizing 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
 128 between minimizing starlight and minimizing the size of the correction, which helps to regularize
 129 the problem and stabilize the solution. In general, the true relationship between the starlight in the
 130 dark zone and the DM correction is highly nonlinear and nonconvex, owing to the fact that the DMs
 131 impart phase-only corrections of the form $\exp\{i\phi_{\text{DM}}\}$ at or near the coronagraph entrance pupil.

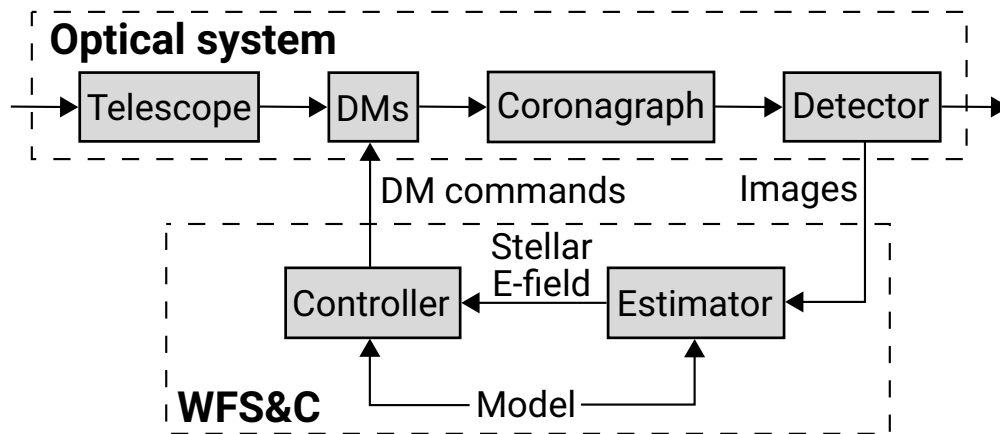


Fig 1 Closed-loop coronagraphic WFS&C uses images from the science camera, rather than an external wavefront sensing instrument, to estimate the electric field from the host star within the dark zone and drive it toward zero. Model-based WFS&C algorithms use a numerical model of the coronagraph to solve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

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

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

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

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

138 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
 139 of dark-zone pixels, and N_{act} is the total number of controllable DM actuators. The intensity from
 140 the corrected electric field, integrated over the dark zone, can be written in terms of the Jacobian as

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

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

143 SM and EFC utilize the relationship in Eq. (5) to derive closed-form expressions for this optimal
 144 correction in terms of \mathbf{G}_k that can be evaluated by solving a linear system of dimension $N_{\text{act}} \times N_{\text{act}}$,
 145 as illustrated in Figure 2. As we discuss in Appendix B, this is equivalent to minimizing the cost
 146 function using Newton's method with an exact Hessian matrix. Alternatively, it is possible to find

147 the DM correction by minimizing the cost function J_k with respect to \mathbf{a}_k iteratively, rather than
 148 analytically, using gradient-based nonlinear optimization as shown in Figure 3. To do so eliminates
 149 the need to calculate the Jacobian matrix, but requires a way of calculating the gradient vector
 150 $\partial J_k / \partial \mathbf{a}_k$. RMAD provides a way of doing so that is both computationally efficient and accurate,
 151 in the sense that the derivatives computed by RMAD are accurate to machine precision and do
 152 not utilize finite-difference approximations.⁸ The basic principle of RMAD is that any function
 153 that can be written down as a sequence of differentiable operations, called the forward model,
 154 can be transformed mechanistically to construct a related function, called the adjoint model, that
 155 evaluates the derivative of the forward model with respect to any of the intermediate variables
 156 encountered during its evaluation, as well as its inputs. Figure 4 illustrates this procedure for the
 157 wavefront control cost function J_k . We refer the reader to our earlier work for a more comprehensive
 158 discussion of the principles of RMAD and its application to wavefront control.⁷

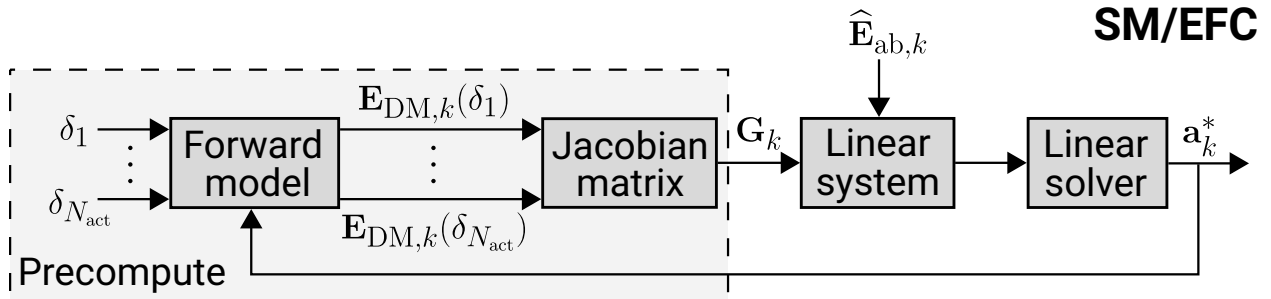


Fig 2 In EFC and SM, the Jacobian matrix is precomputed by using a computer model of the coronagraph to predict the effect of an update to each of the N_{act} DM actuators individually on the dark-zone electric field, here represented by the Kronecker δ functions δ_n , where $\delta_n[i] = 1$ if $i = n$ 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.

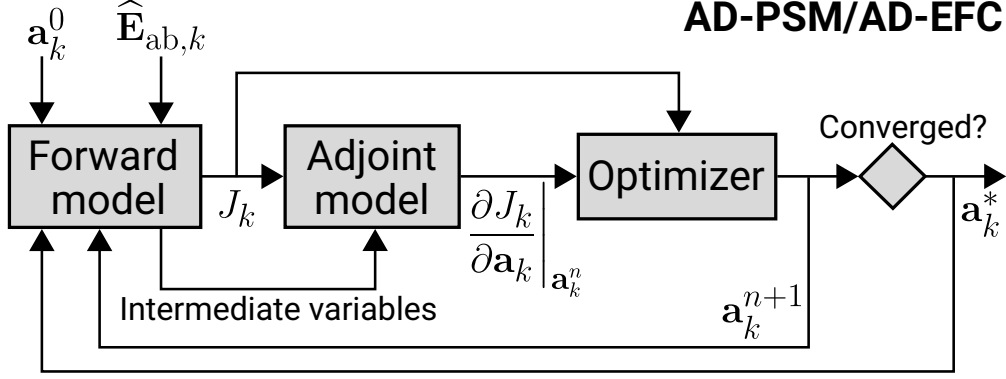


Fig 3 Our algorithmic differentiation-based wavefront controllers AD-PSM and AD-EFC use RMAD to differentiate the wavefront control cost function J_k with respect to the DM correction update \mathbf{a}_k , yielding the gradient vector $\partial J_k / \partial \mathbf{a}_k$ evaluated at the current iterate \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 $\widehat{\mathbf{E}}_{ab,k}$ are the input parameters.

159 **2.1 Stroke minimization: from Lagrange multipliers to penalty method**

160 SM finds the smallest DM correction that achieves a desired level of stellar intensity integrated over

161 the dark zone, denoted by $I_{target,k}$.² It is solved by finding the stationary point of the Lagrangian

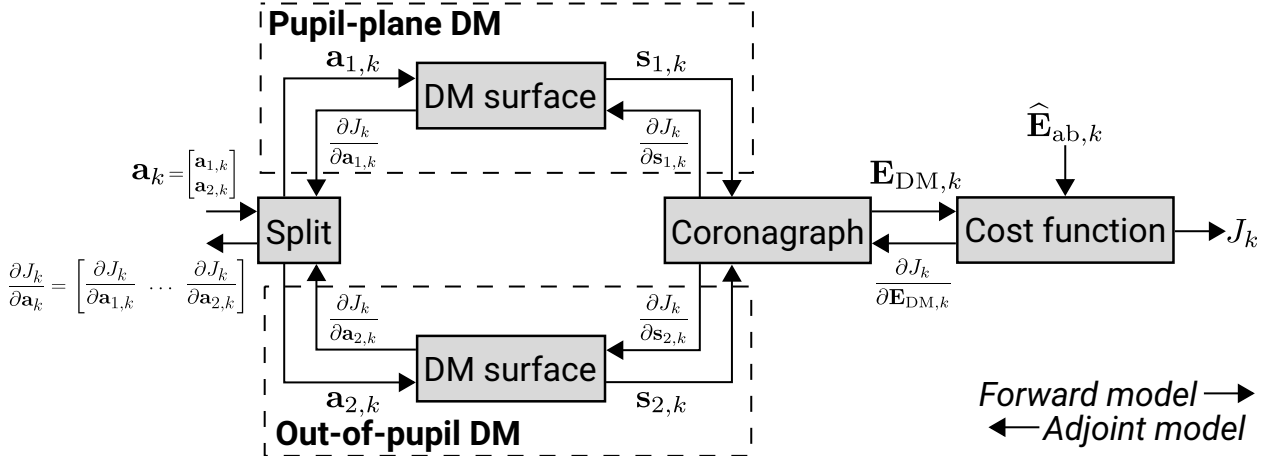


Fig 4 The forward model for the wavefront control problem maps DM command updates \mathbf{a}_k to a scalar cost function J_k . The DM command vector is split into independent command vectors $\mathbf{a}_{1,k}$ and $\mathbf{a}_{2,k}$ for the pupil-plane and out-of-pupil DM, respectively. These are mapped 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 dark-zone electric field $\mathbf{E}_{DM,k}$. Reverse-mode algorithmic differentiation 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.

162 function

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

163 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
 164 the corrected electric field in the dark zone, $\mathbf{a}'_k \triangleq \begin{bmatrix} \mathbf{a}_k^T & \mu_k \end{bmatrix}^T$, and μ_k is the Lagrange multiplier.
 165 Because this stationary point is a saddle point, it cannot be reached by minimizing $\mathcal{L}_{\text{SM},k}$ with
 166 respect to \mathbf{a}'_k . Instead, one chooses a fixed starting value for the Lagrange multiplier, μ_k^0 , and
 167 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
 168 $\|\mathbf{E}_{\text{DZ},k}(\mathbf{a}_k^0)\|^2 \leq I_{\text{target},k}$ is not satisfied, a larger value $\mu_k^1 > \mu_k^0$ is selected and this procedure is
 169 repeated. For any fixed value of μ_k^n , the minimum of Eq. (6) is given in closed form in terms of the
 170 Jacobian matrix \mathbf{G}_k :

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

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

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

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

173 where η_k is a parameter that encodes the relative importance of minimizing actuator stroke and
 174 driving the integrated intensity toward the target value. The minimum of $J_{\text{PSM},k}$ with respect to \mathbf{a}_k
 175 is coincident with the stationary point of $\mathcal{L}_{\text{SM},k}$,⁹ meaning that the DM solutions obtained by SM

176 and AD-PSM are, in principle, identical. In this paper, we choose η_k in each WFS&C iteration as

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

177 where η_{00} , known as the penalty parameter, is a constant set by the experimenter. The denominator
 178 $\left(\|\widehat{\mathbf{E}}_{\text{ab},k}\|^2 - I_{\text{target},k}\right)^2$ scales the cost function to be invariant to the energy in the dark-zone electric
 179 field, which helps in practice to obtain good solutions in all WFS&C iterations as $\widehat{\mathbf{E}}_{\text{ab},k}$ gradually
 180 converges toward zero.

181 2.2 Electric field conjugation

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

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

185 where $\mathbf{\Gamma}_k$ is the Tikhonov regularization matrix. In the most common case, one chooses $\mathbf{\Gamma}_k = \alpha_k \mathbb{I}$
 186 and $\mathbf{E}_{\text{target},k} = 0$, making EFC identical to SM with a fixed Lagrange multiplier $\mu_k = 1/\alpha_k^2$. For
 187 this case, its solution can be obtained using Eq. (7). The general solution for the Jacobian-based
 188 formulation of EFC is given by

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

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

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

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

193 3 Experimental setup

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

197 3.1 The HiCAT testbed

198 HiCAT is a testbed dedicated to technology demonstrations for coronagraphy on segmented-aperture
199 space observatories, with the intent of being directly traceable to a future LUVOIR-like mission.¹⁰⁻¹⁷
200 These technologies include Lyot coronagraphy, high-order WFS&C for generating and stabilizing
201 dark zones, and low-order wavefront sensing (LOWFS).¹⁸ HiCAT operates in a mid-contrast regime
202 (10^{-7} to 10^{-8}) which approaches the limit achievable outside of a vacuum environment. The testbed
203 is equipped with two Boston Micromachines Kilo-DMs for high-order sensing and control, with
204 952 actuators each, making it suitable for our proof-of-concept demonstrations. One DM is placed
205 in a plane conjugate to the HiCAT entrance pupil, while the second DM is placed approximately
206 300 mm farther along the optical axis, corresponding to a Fresnel number $N_F \approx 98$ at a wavelength
207 of 638 nm. This configuration enables simultaneous control of amplitude and phase aberrations over

208 a dark zone that extends over both halves of the image plane. We conducted our experiments using
 209 a Thorlabs MCLS1 laser diode source, which emits monochromatic light with a central wavelength
 210 $\lambda_0 = 638$ nm.

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

216 Our experiments on HiCAT utilized a classical Lyot coronagraph (CLC) design with a hexagonal
 217 entrance pupil mask designed to mask the extreme edges of the IrisAO, a circular focal-plane mask,
 218 and a circular Lyot stop. Figure 6 shows an overlay of the CLC pupil masks along with a simulated

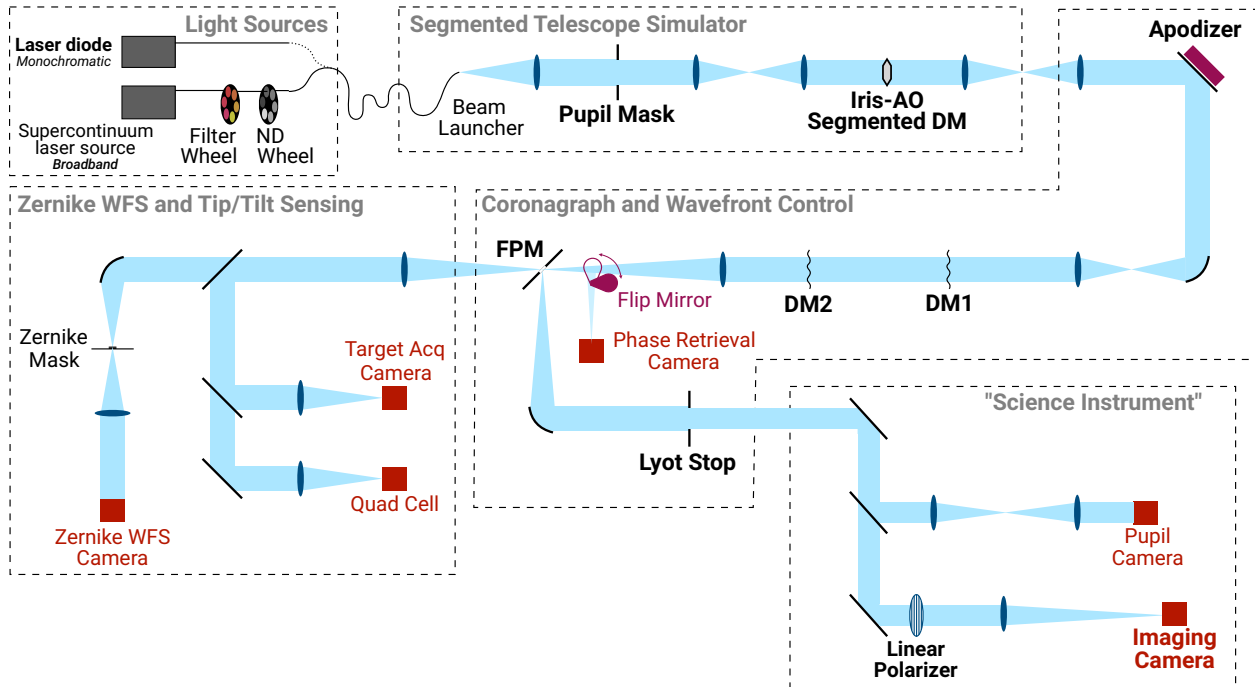


Fig 5 Simplified, partially-unfolded layout of the HiCAT testbed. The elements encountered by the primary imaging beam path are highlighted in bold. DM1 and DM2 indicate the in-pupil and out-of-pupil DMs, respectively. Our experiments utilized a coronagraph configuration in which the apodizer is replaced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

219 stellar PSF. Figure 7 shows example experimental PSFs obtained before and after closed-loop
220 WFS&C using SM along with the corresponding DM commands.

221 *3.2 Algorithm implementation*

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

- 225 1. A model to compute the DM surface resulting from a given set of actuator commands using a
226 fast convolutional representation. This is described in further detail in Appendix C.
- 227 2. A model to apply the phase corrections imparted by the in-pupil DM (DM1) and out-of-pupil
228 DM (DM2), including the free-space propagation between the two DMs.
- 229 3. A model to propagate the electric field after DM correction through the HiCAT CLC. We use
230 the semi-analytical Lyot coronagraph model originally described in Ref. 21.

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

234 *3.2.1 Optimization algorithm*

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




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.

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

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

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

253 **3.3 Estimation algorithm**

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

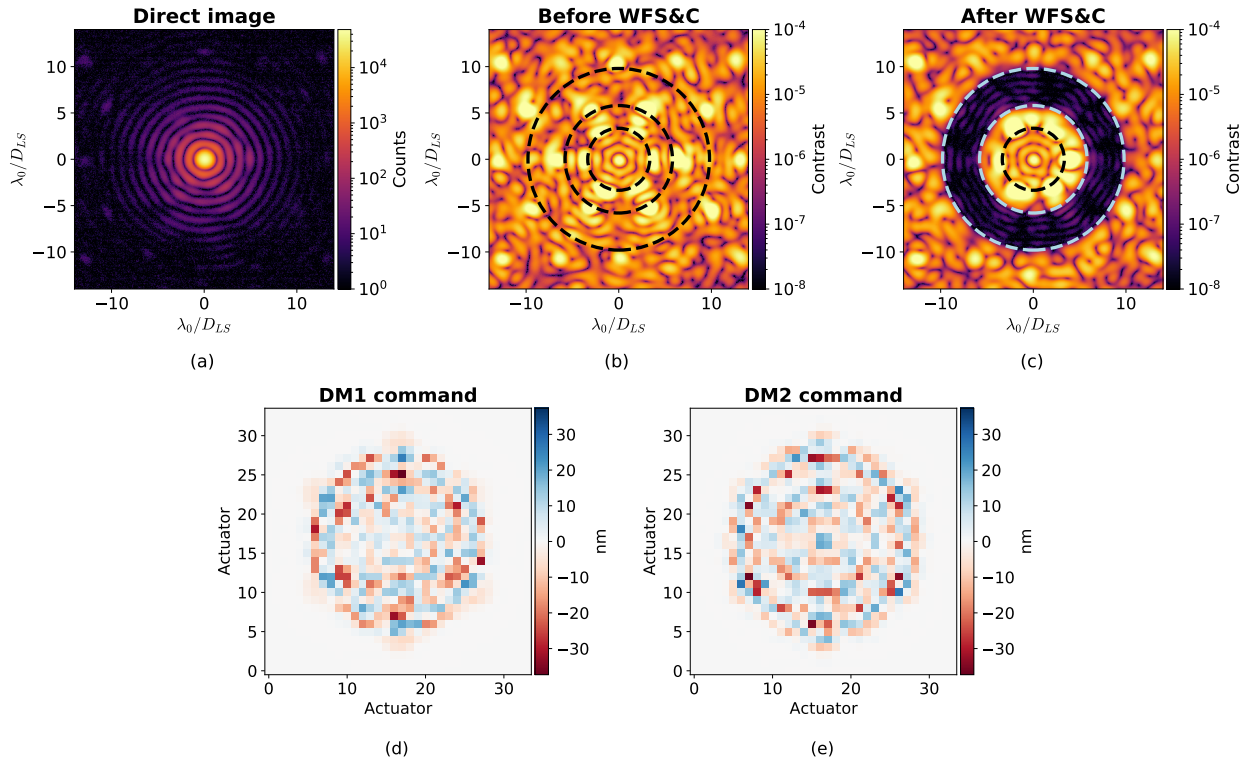


Fig 7 Experimental on-axis images from HiCAT. (a) Non-coronagraphic image, (b) coronagraphic image prior to the WFS&C loop, (c) coronagraphic image after 80 iterations of SM, and corresponding actuator commands for the in-pupil DM (d) and out-of-pupil DM (e). In (b) and (c), the geometrical edge of the 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.

259 found by finding the least-squares solution of the system

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

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

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

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

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

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

266 where α_{probe}^2 is the Tikhonov regularization term for the probe calculation and $\mathbf{G}_{1,k}$ is the Jacobian
267 for the in-pupil DM. This is identical to the EFC problem in Eq. (10) but with the probing field
268 $\mathbf{E}_{\text{DM},k}^p$ on the right-hand side instead of $-\widehat{\mathbf{E}}_{\text{ab},k}$, and has the solution

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

269 For the experiments with SM and EFC, we applied Eq. (16) to compute probe commands for
 270 pairwise probe estimation. For experiments with AD-PSM and AD-EFC, we iteratively solved
 271 Eq. (15) using the AD-EFC framework. In all cases, we set the Tikhonov regularization parameter
 272 for probe generation as $\alpha_{\text{probe}}^2 = 0.7$; for the iterative probes, we set the optimizer tolerance to
 273 $\varepsilon = 10^{-5}$. Figure 8 shows the probe commands \mathbf{u}_p along with the corresponding magnitude and
 274 phase of $\mathbf{E}_{\text{DM},k}(\mathbf{u}_p)$.

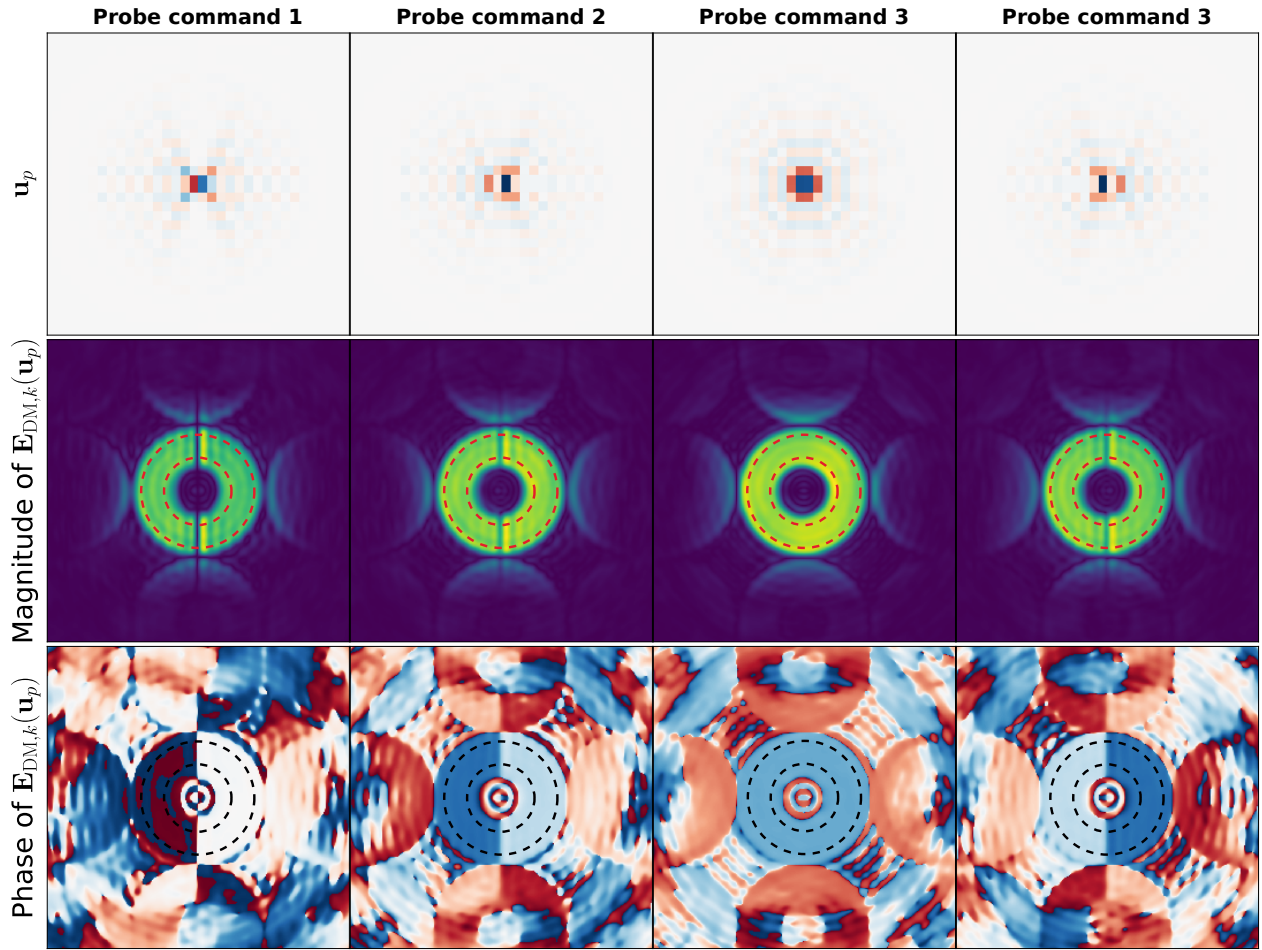


Fig 8 The probe commands \mathbf{u}_p for pairwise estimation were optimized so that the resultant dark-zone electric field $\mathbf{E}_{\text{DM},k}(\mathbf{u}_p) = \text{sgn}\{\rho_x\} \exp\{i\pi(p-1) \text{sgn}\{\rho_x\}/4\}$, as described in Section 3.3. The inner and outer edges of the dark zone 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.

275 4 Experimental results and discussion

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

287 For SM and AD-PSM, we chose $I_{\text{target},k} = 0.5 \|\widehat{\mathbf{E}}_{\text{ab},k}\|^2$ as the contrast target, i.e., a factor of two
288 improvement iteration-over-iteration in spatially-integrated dark zone contrast. For the Lagrange
289 multiplier line search in SM, we let $\mu_k^{n+1} = 1.3 \mu_k^n$. For AD-PSM, we tested combinations of the
290 penalty parameter $\eta_{00} \in \{10, 100, 1000\}$ and the nonlinear optimization convergence tolerance
291 $\varepsilon \in \{10^{-2}, 10^{-3}, 10^{-4}\}$. We determined these values based on simulations and prior WFS&C
292 experiments on HiCAT.

293 For EFC and AD-EFC, we selected the Tikhonov regularization matrix as $\mathbf{\Gamma}_k = \alpha_k \mathbb{I}$, with $\alpha_k^2 \in$
294 $\{10^{-2}, 10^{-3}, 10^{-4}\}$ for the first 30 WFS&C iterations and $\alpha_k^2 = 10^{-1}$ thereafter. We determined
295 through previous experiments with EFC that maintaining an aggressive α_k value throughout the
296 experiment was detrimental to the stability of the control loop after reaching the steady-state regime.

297 For AD-EFC, we used the same set of ε values as AD-PSM.

298 For AD-EFC and AD-PSM, we compared the value of the cost function for two different starting
 299 guesses for the DM correction: $\mathbf{a}_k^0 = \mathbf{a}_{k-1}^*$, i.e., the solution of the previous WFS&C iteration, and
 300 $\mathbf{a}_k^0 = 0$. The starting guess with the lower of the two cost function values was then selected. In all
 301 cases $\mathbf{a}_k^0 = 0$ was ultimately chosen as the starting guess.

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

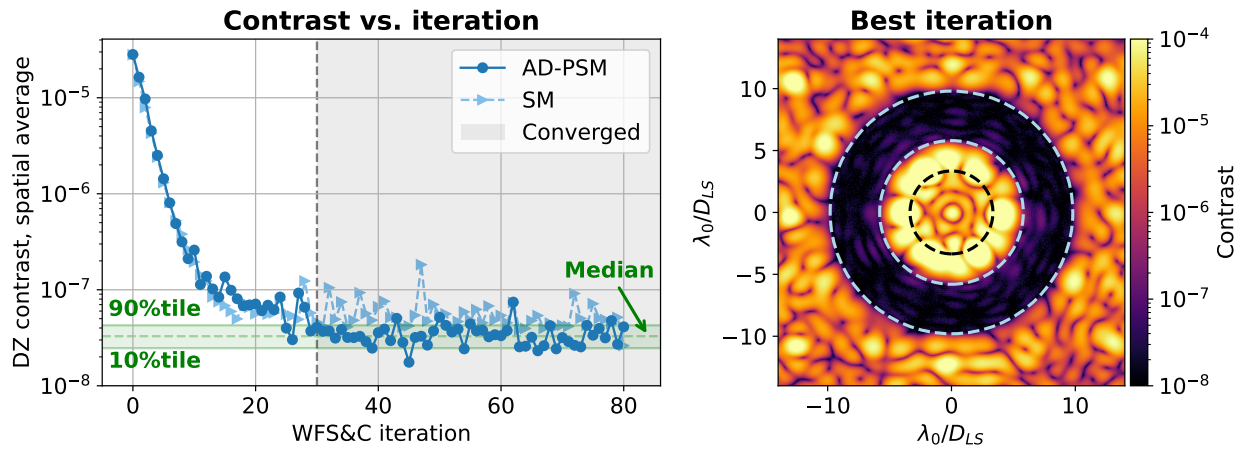


Fig 9 Left: spatially averaged dark-zone contrast vs. WFS&C iteration using AD-PSM with $\eta_{00} = 10$ and $\varepsilon = 10^{-4}$, compared to an experiment with SM. The median, 10th, and 90th percentile of the final 50 iterations, representing the steady-state behavior after convergence, are shown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

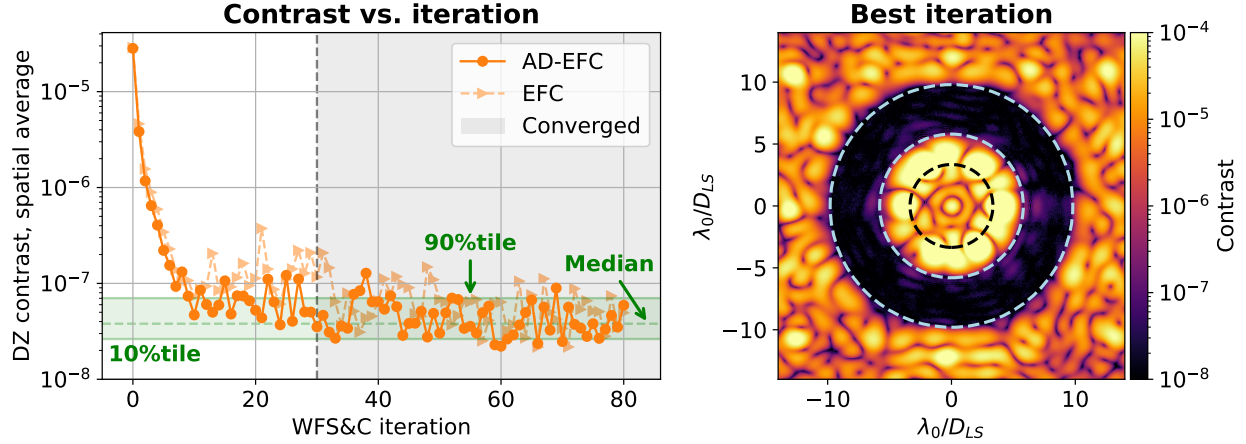


Fig 10 Left: spatially averaged dark-zone contrast vs. WFS&C iteration using AD-EFC with $\alpha_k^2 = 10^{-2}$ and $\varepsilon = 10^{-4}$, compared to an experiment with EFC using the same value for α_k . The median, 10th, and 90th percentile of the converged datapoints are shown in green. Right: the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

310 **4.1 Discussion**

311 Our experiments were aimed at exploring a relevant subset of the space of free parameters for each
 312 algorithm, namely, the nonlinear optimization convergence tolerance ε , the Tikhonov regularization
 313 α_k for AD-EFC, and the penalty parameter η_{00} for AD-PSM. In principle, each parameter affects

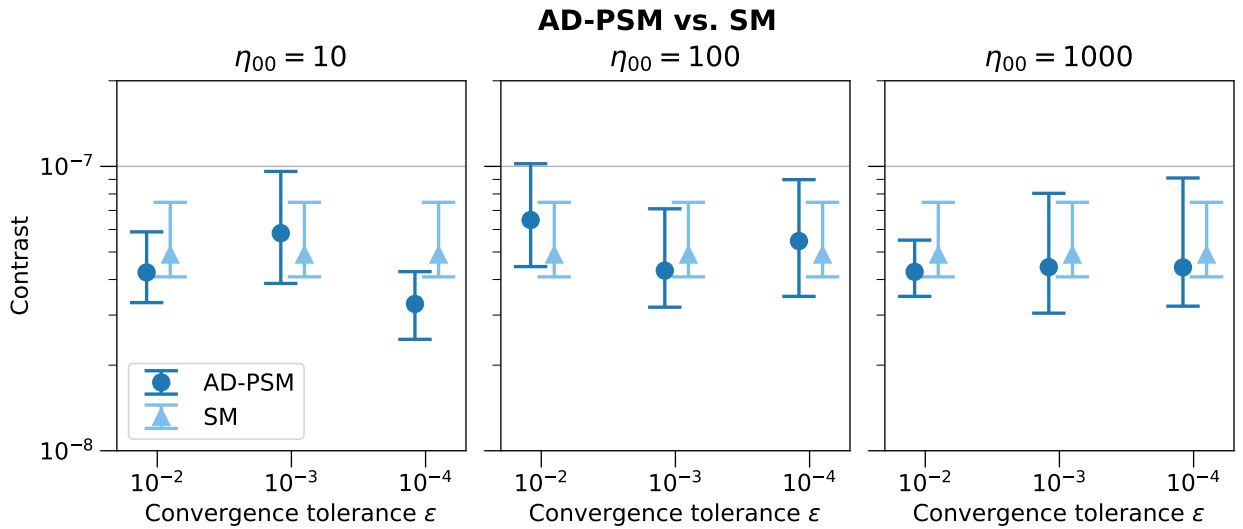


Fig 11 Median, 10th percentile and 90th percentile spatially averaged contrast values achieved by AD-PSM (dark blue) as a function of optimizer tolerance, for three different values of the penalty parameter η_{00} , compared to SM (light blue). The contrast vs. iteration time series for the rightmost result in the left pane ($\eta_{00} = 10, \varepsilon = 10^{-4}$) is illustrated in Figure 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 vs. iteration time series for the rightmost result in the left pane ($\alpha_k^2 = 10^{-2}$, $\epsilon = 10^{-4}$) is illustrated in Figure 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 Section 4.1, the optimizer tolerance prevented AD-EFC from diverging in this case.

314 the attainable contrast of the WFS&C loop, but in subtly different ways, which we discuss here.

315 As discussed in Section 3.2.1, ϵ determines the effort that the nonlinear optimization algorithm
 316 will expend to find a solution close to the true minimum of the cost function. A smaller value of ϵ
 317 corresponds to a larger number of optimization iterations before termination, and ultimately a larger
 318 time interval between DM updates. In principle, on a system such as HiCAT where environmental
 319 disturbances cause the aberrated electric field to evolve over time scales on the order of seconds or
 320 faster, an excessive delay between the estimation step and application of the DM correction can
 321 cause a degradation in achievable contrast. However, we observed no meaningful degradation for
 322 smaller values of ϵ . In a real spaceborne system, this is unlikely to be a significant consideration
 323 because of the much greater electric field stability, and because the total duration of each WFS&C
 324 iteration will be dominated by the exposure times needed for the estimation step.

325 The value ϵ can also serve as an auxiliary form of regularization, by terminating the optimization

326 algorithm before it reaches an overly aggressive DM correction caused by a noisy electric field
 327 estimate, insufficient regularization using α_k or η_{00} , or both. For instance, in Figure 12, with
 328 $\alpha_k^2 = 10^{-4}$ (rightmost panel), EFC diverged altogether whereas AD-EFC did not. On the other hand,
 329 choosing ε too large can impose an effective contrast floor by limiting the ability of the optimization
 330 algorithm to converge to appropriately strong corrections. We determined in simulation that this
 331 was the case for $\varepsilon > 10^{-2}$.

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

342 In principle, because α_k and ε have similar effects on the performance of AD-EFC, there
 343 potentially exists a single combination of the two parameters that is optimal in terms of contrast, or
 344 perhaps a continuum of combinations with similar performance, with a change in α_k compensated
 345 by a change to ε in the opposite direction. The same holds true for AD-PSM. Our experiments
 346 exhibited no strongly identifiable trend over the range of values that we considered, which we
 347 attribute to the fact that for all parameter contributions, the WFS&C loop was able to reach the
 348 contrast floor imposed by environmental, rather than algorithmic, factors.

349 **5 Conclusions and future work**

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

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

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

369 **Appendix A: Convergence data for all experiments**

370 In Section 4, we showed the time series of spatially averaged contrast vs. WFS&C iteration for an
371 AD-PSM with $\eta_{00} = 10$ and $\varepsilon = 10^{-4}$ compared to an experiment with SM (Figure 9), as well as
372 an AD-EFC experiment with $\alpha_k^2 = 10^{-2}$ and $\varepsilon = 10^{-4}$ compared to an EFC experiment with the
373 same value of α_k^2 (Figure 10). We then summarized the steady-state time series statistics (median,
374 10th percentile, 90th percentile) for a series of nine AD-PSM and nine AD-EFC runs with different
375 combinations of (η_{00}, ε) and $(\alpha_k^2, \varepsilon)$, respectively (Figures 11 and 12). In this section, Figure 13
376 shows the full time series of contrast vs. iteration for all 18 AD-PSM and AD-EFC experiments
377 compared to their respective SM and EFC reference experiments.

378 **Appendix B: Equivalence of Jacobian-based solutions and Newton’s method**

379 Finding solutions for EFC and SM using the Jacobian matrix is equivalent to minimizing their
380 respective cost functions with respect to \mathbf{a}_k using Newton’s method. Newton’s method is a second-
381 order optimization technique that utilizes second-derivative information about the cost function,
382 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.
383 Given an initial guess for the solution \mathbf{a}_k^0 , the full Newton update is given by⁹

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

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

386 For general numerical optimization problems, Newton’s method is rarely used in practice because
387 forming the Hessian matrix explicitly is expensive. On the other hand, quasi-Newton methods

388 such as the Broyden-Fletcher-Goldfarb-Shanno (BFGS) algorithm, or the limited-memory BFGS
389 (L-BFGS) variant, can approximate \mathbf{H}_k^{-1} using changes in $\partial J_k / \partial \mathbf{a}_k^T$ over successive optimization
390 iterations. As a consequence, they are substantially less computationally expensive. Although
391 quasi-Newton algorithms do not converge as rapidly as Newton's method, and in particular can
392 converge slowly for poorly-conditioned problems, they are nonetheless superior to purely first-order
393 methods such as steepest descent.⁹

394 As we show below, the Hessian matrix for EFC and SM has an analytic expression in terms of
395 the Jacobian \mathbf{G}_k given by $\mathbf{H}_k = 2\left(\text{Re}\left\{\mathbf{G}_k^\dagger \mathbf{G}_k\right\} + \mathbf{C}\right)$, where \mathbf{C} is a symmetric, positive-definite
396 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
397 iteration, requiring computation of the Jacobian, by a series of cheaper quasi-Newton iterations
398 instead, requiring only computation of the gradient $\partial J_k / \partial \mathbf{a}_k^T$, which we achieve using algorithmic
399 differentiation.

400 The cost functions for the EFC algorithm described in Section 2 can be written in the form

$$J_k(\mathbf{a}_k) = \|\mathbf{b}_k\|^2 + \|\mathbf{\Gamma}_k \mathbf{a}_k\|^2, \quad (18)$$

401 where $\mathbf{\Gamma}_k$ is a regularization matrix. For EFC, $\mathbf{b}_k = \mathbf{E}_{\text{DZ},k}(\mathbf{a}_k) - \mathbf{E}_{\text{target},k}$, while for SM, $\mathbf{b}_k =$
402 $\mathbf{E}_{\text{DZ},k}(\mathbf{a}_k)$ and $\mathbf{\Gamma}_k = \mathbb{I} / \sqrt{\mu_k}$. For now, we will restrict our attention to the EFC algorithm, but note
403 that the result derived here applies equally as well to minimizing the Lagrangian function for SM
404 with respect to the DM correction \mathbf{a}_k .

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

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

406 where $\delta \mathbf{E}_k \triangleq \widehat{\mathbf{E}}_{\text{ab},k} - \mathbf{E}_{\text{target},k}$ and where we use the fact that \mathbf{a}_k is purely real to discard $\text{Im}\{\mathbf{G}_k^\dagger \mathbf{G}_k\}$.

407 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

408 begin by writing down the gradient:

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

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

$$\mathbf{a}_k^* = - \left(\text{Re}\{\mathbf{G}_k^\dagger \mathbf{G}_k\} + \mathbf{\Gamma}_k^T \mathbf{\Gamma}_k \right)^{-1} \text{Re}\{\mathbf{G}_k^\dagger \delta \mathbf{E}_k\}. \quad (21)$$

410 The Hessian matrix is given by

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

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

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

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

416 Combining Eqs. (20) and (22):

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

417 Inserting back into Eq. (23):

$$\mathbf{a}_k^1 = \mathbf{a}_k^0 - \mathbf{H}_k^{-1} \left(\mathbf{H}_k \mathbf{a}_k^0 + 2 \operatorname{Re} \left\{ \mathbf{G}_k^\dagger \delta \mathbf{E}_k \right\} \right) \quad (25a)$$

$$= \mathbf{a}_k^0 - \mathbf{H}_k^{-1} \mathbf{H}_k \mathbf{a}_k^0 - 2 \mathbf{H}_k^{-1} \operatorname{Re} \left\{ \mathbf{G}_k^\dagger \delta \mathbf{E}_k \right\} \quad (25b)$$

$$= \mathbf{a}_k^0 - \mathbf{a}_k^0 - 2 \mathbf{H}_k^{-1} \operatorname{Re} \left\{ \mathbf{G}_k^\dagger \delta \mathbf{E}_k \right\} \quad (25c)$$

$$= -2 \mathbf{H}_k^{-1} \operatorname{Re} \left\{ \mathbf{G}_k^\dagger \delta \mathbf{E}_k \right\}. \quad (25d)$$

418 Inserting the definition of the Hessian matrix from Eq. (22), we see that the Newton iterate \mathbf{a}_k^1 is

419 identical to the analytical solution in Eq. (21):

$$\mathbf{a}_k^1 = -2 \left[2 \left(\operatorname{Re} \left\{ \mathbf{G}_k^\dagger \mathbf{G}_k \right\} + \mathbf{\Gamma}_k^T \mathbf{\Gamma}_k \right) \right]^{-1} \operatorname{Re} \left\{ \mathbf{G}_k^\dagger \delta \mathbf{E}_k \right\} \quad (26a)$$

$$= - \left(\operatorname{Re} \left\{ \mathbf{G}_k^\dagger \mathbf{G}_k \right\} + \mathbf{\Gamma}_k^T \mathbf{\Gamma}_k \right)^{-1} \operatorname{Re} \left\{ \mathbf{G}_k^\dagger \delta \mathbf{E}_k \right\} \quad (26b)$$

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

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

421 **Appendix C: Fast convolutional deformable mirror model**

422 Consider a deformable mirror with N_A actuators along each side (i.e., $N_{\text{act}} = N_A^2$) whose surface
 423 $s(x, y)$ can be modeled as a linear superposition of identical influence functions $f(x, y)$

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

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

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

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

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

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

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

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

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

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

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

429 so that

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

430 where \circ denotes element-wise multiplication. Finally, we define the vectors of actuator center
 431 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
 432 commands \mathbf{A} for which $\mathbf{A}[m, n] = a_{m,n}$:

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

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

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

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

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

436 The term in the parentheses is more commonly referred to as the matrix Fourier transform²¹ or

437 matrix triple product Fourier transform²⁵ of \mathbf{A} , which we denote as follows:

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

438 yielding

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

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

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

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

444 *C.1 Adjoint model*

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

451

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

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

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

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

452 This leads to the following adjoint model, following the RMAD adjoint variable rules in Refs.

453 [7, 26]:

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

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

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

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

455 transform:

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

456 Combining these expressions, the adjoint model is then

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

457 **Appendix D: Adjoint model for EFC cost function**

458 In Section 2.2, we describe the cost function for the EFC algorithm for a single correction wavelength.

459 Here, we derive its RMAD adjoint model, which computes the derivative $\partial J_{\text{EFC},k}/\partial \mathbf{E}_{\text{DM},k}$.

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

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

$$\Delta \mathbf{E}_k = \mathbf{E}_{\text{DZ},k} - \mathbf{E}_{\text{target},k} , \quad (42\text{b})$$

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

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

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

$$J_{\text{EFC},k} = J_{\Delta \mathbf{E},k} + J_{\mathbf{c},k} . \quad (42\text{f})$$

461 We now apply the RMAD gradient rules^{7,26} to each step in reverse order to derive the adjoint model,

462 letting $\bar{\mathbf{x}} \triangleq \partial J_k / \partial \mathbf{x}^T$ for any variable \mathbf{x} :

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

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

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

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

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

$$\bar{\mathbf{E}}_{\text{DZ},k} = \overline{\Delta \mathbf{E}}_k, \quad (43\text{f})$$

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

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

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

464 This gradient is passed to the next block of the adjoint model, which in this case is the adjoint model
 465 for propagation through the coronagraph, which, referring to Figure 4, evaluates the derivatives
 466 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
 467 derivation of the coronagraph propagation adjoint model.⁷

468 *Disclosures*

469 The authors declare no conflicts of interest.

470 *Data and Code Availability*

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

473 *Acknowledgments*

474 The authors are especially thankful to the extended HiCAT team (over 50 people) who have
475 worked over the past several years to develop this testbed. This work was supported in part by the
476 National Aeronautics and Space Administration under Grant 80NSSC19K0120 issued through the
477 Strategic Astrophysics Technology / Technology Demonstration for Exoplanet Missions Program
478 (SAT-TDEM; PI: R. Soummer). E.H.P. was supported by the NASA Hubble Fellowship grant
479 HST-HF2-51467.001-A awarded by the Space Telescope Science Institute, which is operated by
480 the Association of Universities for Research in Astronomy, Incorporated, under NASA contract
481 NAS5-26555. I.L. acknowledges the support by a postdoctoral grant issued by the Centre National
482 d'Études Spatiales (CNES) in France.

483 This paper was adapted from an earlier SPIE conference proceedings.²⁷

484 *Data, Materials, and Code Availability*

485 HiCAT makes use of the NumPy,^{28,29} Matplotlib,^{30,31} AstroPy,³² SciPy,³³ scikit-image,³⁴ Pan-
486 das,^{35,36} ImageIO,³⁷ Photutils,³⁸ HCIPy,³⁹ POPPY^{22,23} and CatKit⁴⁰ packages.

487 *References*

- 488 1 M. Perryman, *The Exoplanet Handbook*, Cambridge University Press, 2nd ed. (2018).
- 489 2 L. Pueyo, J. Kay, N. J. Kasdin, *et al.*, “Optimal dark hole generation via two deformable
490 mirrors with stroke minimization.,” *Appl. Opt.* **48**, 6296–6312 (2009).

- 491 3 A. Give'on, R. Belikov, S. Shaklan, *et al.*, "Closed loop, DM diversity-based, wavefront
492 correction algorithm for high contrast imaging systems," *Opt. Express* **15**(19), 12338–12343
493 (2007).
- 494 4 B.-J. Seo, K. Patterson, K. Balasubramanian, *et al.*, "Testbed demonstration of high-contrast
495 coronagraph imaging in search for Earth-like exoplanets," *Proc. SPIE* **11117**, 111171V (2019).
- 496 5 The LUVOIR Study Team, "LUVOIR," tech. rep., National Aeronautics and Space Adminis-
497 tration (2019).
- 498 6 The Habitable Exoplanet Observatory Study Team, "Habitable Exoplanet Observatory Final
499 Report," tech. rep., Jet Propulsion Laboratory (2019).
- 500 7 S. D. Will, T. D. Groff, and J. R. Fienup, "Jacobian-free coronagraphic wavefront control
501 using nonlinear optimization,," *J. Astron. Telesc. Instrum. Syst.* **7**, 019002 (2021).
- 502 8 A. Griewank and A. Walther, *Evaluating Derivatives: Principles and Techniques of Algorithmic*
503 *Differentiation*, Society for Industrial and Applied Mathematics, 2nd ed. (2008).
- 504 9 J. Nocedal and S. J. Wright, *Numerical Optimization*, Springer-Verlag New York (1999).
- 505 10 M. N'Diaye, E. Choquet, L. Pueyo, *et al.*, "High-contrast imager for complex aperture
506 telescopes (HiCAT): 1. testbed design," *Proc. SPIE* **8864**, 88641K (2013).
- 507 11 M. N'Diaye, E. Choquet, S. Eggen, *et al.*, "High-contrast imager for complex aperture
508 telescopes (HiCAT): II. Design overview and first light results," *Proc. SPIE* **9143**, 914327
509 (2014).
- 510 12 M. N'Diaye, J. Mazoyer, E. Choquet, *et al.*, "High-contrast imager for complex aperture
511 telescopes (HiCAT): 3. first lab results with wavefront control," *Proc. SPIE* **9605**, 96050I–
512 96050I–12 (2015).

- 513 13 L. Leboulleux, M. N'Diaye, A. J. E. Riggs, *et al.*, “High-contrast imager for Complex Aperture
514 Telescopes (HiCAT). 4. Status and wavefront control development,” *Proc. SPIE* **9904**, 99043C
515 (2016).
- 516 14 R. Soummer, G. R. Brady, K. Brooks, *et al.*, “High-contrast imager for complex aperture
517 telescopes (HiCAT): 5. first results with segmented-aperture coronagraph and wavefront
518 control,” *Proc. SPIE* **10698**, 106981O (2018).
- 519 15 C. Moriarty, K. Brooks, R. Soummer, *et al.*, “High-contrast imager for complex aperture
520 telescopes (HiCAT): 6. software control infrastructure and calibration,” *Proc. SPIE* **10698**,
521 1069853 (2018).
- 522 16 R. Soummer, I. Luginja, S. Will, *et al.*, “High-contrast imager for complex aperture telescopes
523 (HiCAT): 6. Two deformable mirror wavefront control (Conference Presentation),” *Proc. SPIE*
524 **11117**, 111171Y (2019).
- 525 17 R. Soummer, I. Luginja, S. D. Will, *et al.*, “High-contrast imager for complex aperture
526 telescopes (HiCAT): 7. Dark zone demonstration with fully segmented aperture coronagraph,”
527 *Proc. SPIE* **11823**, 11823–27 (2021).
- 528 18 R. Pourcelot, M. N'Diaye, E. H. Por, *et al.*, “Experimental validation of active control of
529 low-order aberrations with a Zernike sensor through a Lyot coronagraph,” *Proc. SPIE* **11823**,
530 11823–46 (2021).
- 531 19 I. Luginja, J.-F. Sauvage, L. Mugnier, *et al.*, “Wavefront tolerances of space-based segmented
532 telescopes at very high contrast: Experimental validation,” *Astron. Astrophys.* **658**, A84 (2022).
- 533 20 S. M. Redmond, L. Pueyo, L. Pogorelyuk, *et al.*, “Dark hole maintenance results for segmented

- 534 aperture wavefront error drift in a high contrast space coronagraph,” *Proc. SPIE* **11823**, 11823–
535 48 (2021).
- 536 21 R. Soummer, L. Pueyo, A. Sivaramakrishnan, *et al.*, “Fast computation of Lyot-style corona-
537 graph propagation,” *Opt. Express* **15**, 15935–15951 (2007).
- 538 22 M. D. Perrin, R. Soummer, E. M. Elliot, *et al.*, “Simulating point spread functions for the
539 James Webb Space Telescope with WebbPSF,” *Proc. SPIE* **84423D**, 8442 – 8442 – 11 (2012).
- 540 23 M. Perrin, J. Long, E. Douglas, *et al.*, “POPPY: Physical Optics Propagation in PYthon,”
541 (2016).
- 542 24 P. Virtanen, R. Gommers, T. E. Oliphant, *et al.*, “SciPy 1.0: Fundamental algorithms for
543 scientific computing in Python,” *Nat. Methods* **17**, 261–272 (2020).
- 544 25 A. S. Jurling, M. D. Bergkoetter, and J. R. Fienup, “Techniques for arbitrary sampling in
545 two-dimensional Fourier transforms,” *J. Opt. Soc. Am. A* **A35**, 1784–1796 (2018).
- 546 26 A. S. Jurling and J. R. Fienup, “Applications of algorithmic differentiation to phase retrieval
547 algorithms.,” *J. Opt. Soc. Am. A* **31**(7), 1348–59 (2014).
- 548 27 S. D. Will, M. D. Perrin, E. H. Por, *et al.*, “Wavefront control with algorithmic differentiation
549 on the HiCAT testbed,” *Proc. SPIE* **11823**, 118230V (2021).
- 550 28 T. E. Oliphant, *A guide to NumPy*, vol. 1, Trelgol Publishing USA (2006).
- 551 29 S. van der Walt, S. C. Colbert, and G. Varoquaux, “The NumPy array: a structure for efficient
552 numerical computation,” *Computing in Science and Engineering* **13**, 22–30 (2011).
- 553 30 J. D. Hunter, “Matplotlib: A 2D Graphics Environment,” *Computing in Science and Engineer-*
554 *ing* **9**, 90–95 (2007).
- 555 31 T. A. Caswell, M. Droettboom, A. Lee, *et al.*, “matplotlib/matplotlib: Rel: v3.3.3,” (2020).

- 556 32 Astropy Collaboration, A. M. Price-Whelan, B. M. Sipőcz, *et al.*, “The Astropy Project:
557 Building an Open-science Project and Status of the v2.0 Core Package,” *Astrophys. J.* **156**,
558 123 (2018).
- 559 33 P. Virtanen, R. Gommers, T. E. Oliphant, *et al.*, “SciPy 1.0: Fundamental algorithms for
560 scientific computing in Python,” *Nature Methods* **17**, 261–272 (2020).
- 561 34 S. van der Walt, J. L. Schönberger, J. Nunez-Iglesias, *et al.*, “scikit-image: Image processing
562 in Python,” *PeerJ* (2014).
- 563 35 Wes McKinney, “Data structures for statistical computing in python,” in *Proceedings of the*
564 *9th Python in Science Conference*, Stéfan van der Walt and Jarrod Millman, Eds., 56 – 61
565 (2010).
- 566 36 J. Reback, W. McKinney, jbrockmendel, *et al.*, “pandas-dev/pandas: Pandas 1.1.4,” (2020).
- 567 37 S. Silvester, A. Tanbakuchi, P. Müller, *et al.*, “imageio/imageio v2.8.0,” (2020).
- 568 38 L. Bradley, B. Sipőcz, T. Robitaille, *et al.*, “astropy/photutils: 1.0.1,” (2020).
- 569 39 E. H. Por, S. Y. Haffert, V. M. Radhakrishnan, *et al.*, “High Contrast Imaging for Python
570 (HCIPy): an open-source adaptive optics and coronagraph simulator,” *Proc. SPIE* **10703**,
571 1070342 (2018).
- 572 40 J. Noss, J. Fowler, C. Moriarty, *et al.*, “spacetelescope/catkit: v0.36.1,” (2021).

573 **Scott D. Will** is an optical engineer at NASA Goddard Space Flight Center in Greenbelt, Maryland,
574 USA. He received a B.S. in electrical engineering from the University at Buffalo in 2015, and
575 a Ph.D. in optics from the University of Rochester in 2021. As a graduate student, he was also
576 a member of the Russell B. Makidon Optics Laboratory team at the Space Telescope Science

577 Institute in Baltimore, Maryland, USA. His research interests include wavefront sensing and control,
578 computational imaging, and astronomical instrumentation.

579 **Emiel H. Por** is an NHFP Sagan fellow at the Russell B. Makidon Optics Laboratory at the Space
580 Telescope Science Institute in Baltimore, USA. He received his B.Sc. degree in Astronomy and
581 Physics, his M.Sc. degree in Astronomical Instrumentation, and his Ph.D. in Astronomy from
582 Leiden University, The Netherlands. His research interests include coronagraphy, wavefront sensing
583 and control, and high-contrast imaging.

584 **Ananya Sahoo** is a Staff Scientist at STScI. She currently splits her time between infrastructure
585 development for the HiCAT testbed and studying wavefront-error induced by thermal disturbances
586 for future segmented space telescopes. She received her PhD degree in 2020 from The Graduate
587 University for Advanced Studies, SOKENDAI in alliance with National Astronomical Observatory
588 of Japan (NAOJ). She previously worked on the photometric and astrometric calibration for the
589 SCExAO instrument at the Subaru Telescope.

590 **Iva Laginja** is a postdoctoral fellow with CNES, the French space agency and previous member, now
591 collaborator, of the Russell B. Makidon Optics Laboratory at the Space Telescope Science Institute
592 in Baltimore, USA. She received a M.Sc. degree in Astronomy and Instrumentation at Leiden
593 University, Netherlands and her doctorate from Paris Observatory in 2021. Her research interests
594 include wavefront sensing and control, high contrast imaging and coronagraphy for exoplanet
595 detection.

596 **Raphaël Pourcelot** is a postdoc at STScI within the Russell B. Makidon Optics Laboratory. He
597 received a M.Sc. in astronomy from Aix-Marseille Université and an engineering degree at Institut
598 d'Optique Graduate School in Université Paris-Saclay. He received his Ph.D. from at Université

599 Côte d'Azur in Nice, France in 2022. He is mainly working on wavefront sensing and control
600 for high-contrast imaging for exoplanet detection and Zernike wavefront sensor applications in
601 particular.

602 **Susan F. Redmond** is a Ph.D. student at Princeton University in Princeton, USA and collaborates
603 with the Russell B. Makidon Optics Laboratory. She previously obtained her B.Eng. from Memorial
604 University of Newfoundland (2015) and her M.Eng., M.ASc degrees in Aerospace Engineering
605 from the University of Toronto (2016, 2018). Her current research is split between the optical and
606 thermal design for balloon-borne telescopes and developing focal-plane wavefront sensing and
607 control algorithms for high-contrast imaging.

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

613 **Tyler D. Groff** received his B.S. in mechanical engineering and astrophysics from Tufts University
614 and his Ph.D. in mechanical and aerospace engineering from Princeton University under an NESSF
615 fellowship. He was the lab manager of the Princeton High Contrast Imaging Laboratory and the
616 CHARIS instrument at Subaru Telescope. He is currently the lead engineer at Goddard Space Flight
617 Center for the spectroscopy and polarization modes on the Roman Space Telescope Coronagraph
618 Instrument.

619 **James R. Fienup** received his A.B. from Holy Cross College, and M.S. and Ph.D. degrees in
620 Applied Physics from Stanford University, where he was a National Science Foundation graduate

621 fellow. After performing research at ERIM, he became the Robert E. Hopkins Professor of Optics at
622 the University of Rochester. He is a fellow of SPIE and OSA, a member of the National Academy
623 of Engineering, and a recipient of SPIE's Rudolf Kingslake Medal and Prize.

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

629 Biographies and photographs of the other authors are not available.

630 **List of Figures**

631 1 Closed-loop coronagraphic WFS&C uses images from the science camera, rather
632 than an external wavefront sensing instrument, to estimate the electric field from
633 the host star within the dark zone and drive it toward zero. Model-based WFS&C
634 algorithms use a numerical model of the coronagraph to solve an inverse problem
635 for the unknown electric field and corresponding DM correction, respectively.

636 2 In EFC and SM, the Jacobian matrix is precomputed by using a computer model of
637 the coronagraph to predict the effect of an update to each of the N_{act} DM actuators
638 individually on the dark-zone electric field, here represented by the Kronecker δ
639 functions δ_n , where $\delta_n[i] = 1$ if $i = n$ and zero otherwise. The Jacobian \mathbf{G}_k and
640 the estimate of the aberrated dark-zone electric field $\widehat{\mathbf{E}}_{\text{ab},k}$ together are used to
641 construct a linear system of equations whose solution is the desired DM update \mathbf{a}_k^* .
642 We show in Appendix B that this is equivalent to minimizing the wavefront control
643 cost function using Newton's method.

644 3 Our algorithmic differentiation-based wavefront controllers AD-PSM and AD-EFC
645 use RMAD to differentiate the wavefront control cost function J_k with respect to
646 the DM correction update \mathbf{a}_k , yielding the gradient vector $\partial J_k / \partial \mathbf{a}_k$ evaluated at the
647 current iterate \mathbf{a}_k^n . A nonlinear optimization algorithm calculates a new iterate \mathbf{a}_k^{n+1}
648 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
649 is repeated until the gradient becomes sufficiently small, indicating that the solution
650 \mathbf{a}_k^* is at or near a local minimum of the cost function. A starting guess for the
651 solution \mathbf{a}_k^0 as well as the aberrated electric field $\widehat{\mathbf{E}}_{\text{ab},k}$ are the input parameters.

652 4 The forward model for the wavefront control problem maps DM command updates
653 \mathbf{a}_k to a scalar cost function J_k . The DM command vector is split into independent
654 command vectors $\mathbf{a}_{1,k}$ and $\mathbf{a}_{2,k}$ for the pupil-plane and out-of-pupil DM, respectively.
655 These are mapped onto DM surfaces $\mathbf{s}_{1,k}$ and $\mathbf{s}_{2,k}$ using the model in Appendix C,
656 and propagated through an end-to-end coronagraph model to predict the resulting
657 dark-zone electric field $\mathbf{E}_{\text{DM},k}$. Reverse-mode algorithmic differentiation transforms
658 the forward model into an adjoint model that backpropagates the partial derivatives
659 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 $\mathbf{E}_{\text{DM},k}$ in
660 reverse order, starting from the output on the right. The derivatives with respect to
661 the individual DM command vectors $\mathbf{a}_{1,k}$ and $\mathbf{a}_{2,k}$ are concatenated to form the full
662 gradient vector for optimization.

663 5 Simplified, partially-unfolded layout of the HiCAT testbed. The elements encoun-
664 tered by the primary imaging beam path are highlighted in bold. DM1 and DM2
665 indicate the in-pupil and out-of-pupil DMs, respectively. Our experiments utilized a
666 coronagraph configuration in which the apodizer is replaced with a flat mirror (see
667 3.1)

668 6 (a) Overlay of HiCAT pupil masks projected onto the in-pupil DM plane, including
669 the reflective area of the DM, the IrisAO segmented aperture, the entrance pupil
670 mask (with the same geometry, but slightly undersized relative to the IrisAO) and
671 the Lyot stop. (b) Simulated and (c) experimental coronagraphic PSFs with outline
672 of the geometrical edge of the focal-plane mask.

673 7 Experimental on-axis images from HiCAT. (a) Non-coronagraphic image, (b) coro-
674 nagraphic image prior to the WFS&C loop, (c) coronagraphic image after 80
675 iterations of SM, and corresponding actuator commands for the in-pupil DM (d)
676 and out-of-pupil DM (e). In (b) and (c), the geometrical edge of the focal-plane
677 mask (FPM) with radius $3.34 \lambda_0/D_{LS}$ is shown as well as the inner and outer edges
678 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
679 stop diameter and $\lambda_0 = 638$ nm.

680 8 The probe commands \mathbf{u}_p for pairwise estimation were optimized so that the resul-
681 tant dark-zone electric field $\mathbf{E}_{\text{DM},k}(\mathbf{u}_p) = \text{sgn}\{\boldsymbol{\rho}_x\} \exp\{i\pi(p-1) \text{sgn}\{\boldsymbol{\rho}_x\}/4\}$, as
682 described in Section 3.3. The inner and outer edges of the dark zone are also shown
683 for reference. The probe commands are close to the inverse Fourier transform of
684 the dark zone geometry (an annulus), modulated by a horizontal sinusoid whose
685 phase angle is proportional to the desired piston phase, and projected onto the DM
686 actuator coordinates. Only the pupil-plane DM is modulated.

687 9 Left: spatially averaged dark-zone contrast vs. WFS&C iteration using AD-PSM
688 with $\eta_{00} = 10$ and $\varepsilon = 10^{-4}$, compared to an experiment with SM. The median, 10th,
689 and 90th percentile of the final 50 iterations, representing the steady-state behavior
690 after convergence, are shown in green. Right: the on-axis PSF corresponding to the
691 iteration with deepest contrast using AD-PSM.

692 10 Left: spatially averaged dark-zone contrast vs. WFS&C iteration using AD-EFC
693 with $\alpha_k^2 = 10^{-2}$ and $\varepsilon = 10^{-4}$, compared to an experiment with EFC using the same
694 value for α_k . The median, 10th, and 90th percentile of the converged datapoints are
695 shown in green. Right: the on-axis PSF corresponding to the iteration with deepest
696 contrast using AD-EFC experiment.

697 11 Median, 10th percentile and 90th percentile spatially averaged contrast values
698 achieved by AD-PSM (dark blue) as a function of optimizer tolerance, for three
699 different values of the penalty parameter η_{00} , compared to SM (light blue). The
700 contrast vs. iteration time series for the rightmost result in the left pane ($\eta_{00} = 10$,
701 $\varepsilon = 10^{-4}$) is illustrated in Figure 9. In all cases, the contrast performance of
702 AD-PSM was equivalent to that of SM within statistical uncertainty.

703 12 Median, 10th percentile and 90th percentile spatially averaged contrast values
704 achieved by AD-EFC (dark orange) as a function of optimizer tolerance, for three
705 different values of the Tikhonov regularization parameter α_k . For each α_k , we
706 also ran a reference experiment with EFC (light orange); the EFC experiment with
707 $\alpha_k^2 = 10^{-4}$ diverged and is not shown. The contrast vs. iteration time series for the
708 rightmost result in the left pane ($\alpha_k^2 = 10^{-2}$, $\varepsilon = 10^{-4}$) is illustrated in Figure 10.
709 The performance of AD-EFC was statistically equivalent to EFC for $\alpha_k^2 = 10^{-2}$ and
710 $\alpha_k^2 = 10^{-3}$. As discussed in Section 4.1, the optimizer tolerance prevented AD-EFC
711 from diverging in this case.

712 13 Spatially averaged dark-zone contrast vs. iteration for all experiments discussed in
713 Section 4 and summarized in Figures 11 and 12. The 10th percentile, median, and
714 90th percentile of the steady-state contrast values (gray region) for AD-PSM and
715 AD-EFC are shown in green. The AD-PSM experiment with ($\eta_{00} = 10, \varepsilon = 10^{-4}$)
716 and the AD-EFC experiment with ($\alpha_k^2 = 10^{-2}, \varepsilon = 10^{-4}$) are also shown in Figures
717 9 and 10, respectively.

718 **List of Tables**

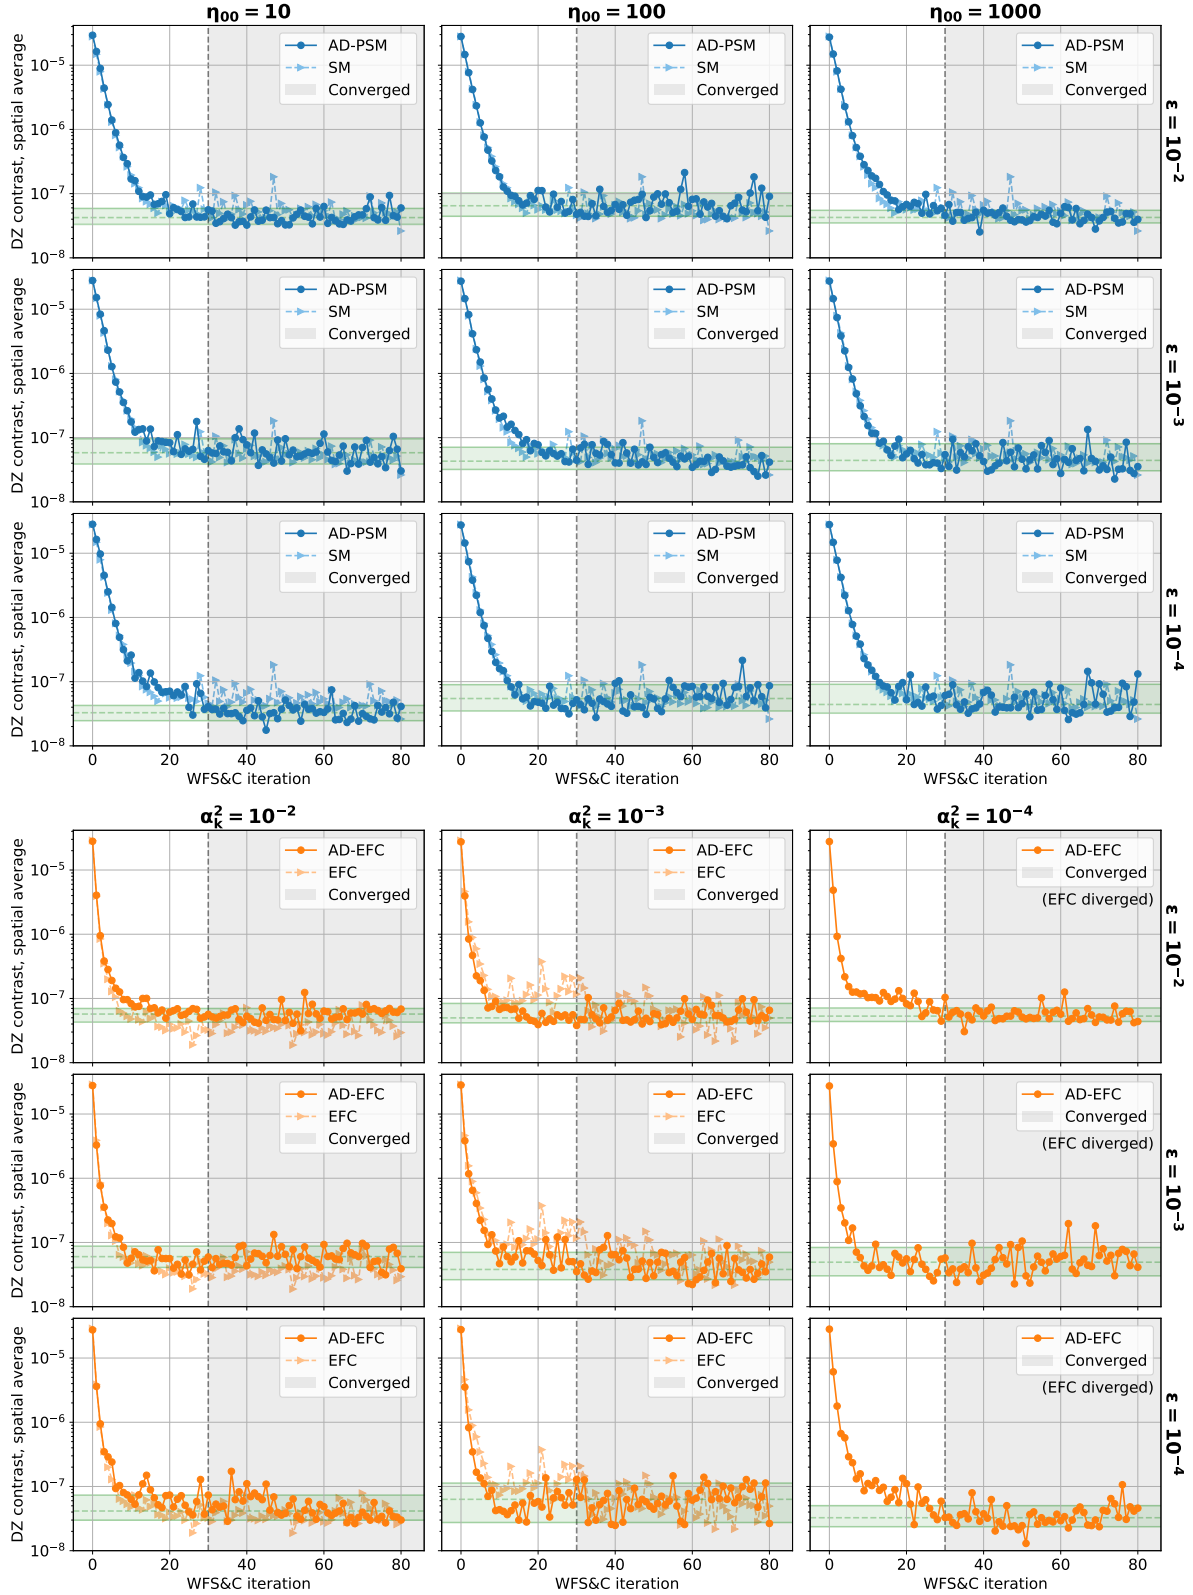


Fig 13 Spatially averaged dark-zone contrast vs. iteration for all experiments discussed in Section 4 and summarized in Figures 11 and 12. The 10th percentile, median, and 90th percentile of the steady-state contrast values (gray region) for AD-PSM and AD-EFC are shown in green. The AD-PSM experiment 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 Figures 9 and 10, respectively.