Journal of Astronomical Telescopes, Instruments, and Systems

AstronomicalTelescopes.SPIEDigitalLibrary.org

Recursive starlight and bias estimation for high-contrast imaging with an extended Kalman filter

A. J. Eldorado Riggs
N. Jeremy Kasdin
Tyler D. Groff

Recursive starlight and bias estimation for high-contrast imaging with an extended Kalman filter

A. J. Eldorado Riggs,* N. Jeremy Kasdin, and Tyler D. Groff

Princeton University, Department of Mechanical and Aerospace Engineering, Engineering Quadrangle, Olden Street, Princeton, New Jersey 08544, United States

Abstract. For imaging faint exoplanets and disks, a coronagraph-equipped observatory needs focal plane wavefront correction to recover high contrast. The most efficient correction methods iteratively estimate the stellar electric field and suppress it with active optics. The estimation requires several images from the science camera per iteration. To maximize the science yield, it is desirable both to have fast wavefront correction and to utilize all the correction images for science target detection. Exoplanets and disks are incoherent with their stars, so a nonlinear estimator is required to estimate both the incoherent intensity and the stellar electric field. Such techniques assume a high level of stability found only on space-based observatories and possibly ground-based telescopes with extreme adaptive optics. In this paper, we implement a nonlinear estimator, the iterated extended Kalman filter (IEKF), to enable fast wavefront correction and a recursive, nearly-optimal estimate of the incoherent light. In Princeton's High Contrast Imaging Laboratory, we demonstrate that the IEKF allows wavefront correction at least as fast as with a Kalman filter and provides the most accurate detection of a faint companion. The nonlinear IEKF formalism allows us to pursue other strategies such as parameter estimation to improve wavefront correction. © The Authors. Published by SPIE under a Creative Commons Attribution 3.0 Unported 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.2.1.011017](https://doi.org/10.1117/1.JATIS.2.1.011017)]

Keywords: wavefront sensing; wavefront control; high-contrast imaging; coronagraph; exoplanet.

Paper 15039SSP received Jun. 1, 2015; accepted for publication Jan. 12, 2016; published online Feb. 5, 2016.

1 Introduction

Radial velocity and transit surveys have transformed our understanding of the universe by detecting thousands of planets outside our solar system. Dozens of these known exoplanets are close enough to image directly, which would allow us to obtain their spectra and fully determine their orbital parameters. Direct imaging requires high contrast in the image, factors of 10^{10} or more for Earth-size planets or 10^9 for Neptune and Jupiter analogs. Atmospheric turbulence precludes obtaining such high contrast from the ground, so a space-based observatory is necessary. The proposed Coronagraph Instrument (CGI) on the Wide-Field Infrared Survey Telescope-Astrophysics Focused Telescope Asset (WFIRST-AFTA) is expected to image about 16 known cool gas-giant exoplanets and spectrally characterize about six of them.¹

A coronagraph uses a series of apodizers, masks, and stops in the optical train to modify or remove the point spread function (PSF) of the telescope and create image plane regions of high contrast where an exoplanet can be seen. The optics for a coronagraph cannot be manufactured to the smoothness and reflectivity requirements to obtain 10^{-9} or better planet-to-star contrast passively.² A set of deformable mirrors (DMs) is necessary to mitigate these aberrations and recover a high-contrast region called a dark hole. The CGI on WFIRST-AFTA will be equipped with two coronagraph types and two DMs, making it the first high-contrast coronagraphic mission in space with high-actuator-count wavefront control.

Wavefront correction for high-contrast coronagraphy differs from the method regularly used in astronomy. In traditional adaptive optics, the wavefront phase is measured at a pupil and conjugated by a DM. That approach is inadequate for generating high contrast; a nonflat pupil phase is acceptable as long as the starlight destructively interferes in the image. In addition, uncontrollable, high-spatial frequencies in the pupil phase can mix and move light into the dark hole, even if all correctable spatial frequencies are eliminated by the DMs. Finally, amplitude aberrations degrade contrast and cannot be mitigated with just phase conjugation. As a result, the wavefront control approach being planned for extreme high contrast in space relies on correction in the focal plane. This requires forming an estimate of the focal plane electric field and then finding a DM command to improve the contrast.

The main challenge of focal plane wavefront correction with a high-contrast coronagraph is to quickly sense the wavefront. The most efficient correction routines need an estimate of the stellar electric field. A wavefront sensor cannot be used (alone) because it estimates only the phase and introduces noncommon-path aberrations. The science camera is the only common-path sensor available, but exposure times can become very long as the contrast gets higher and the signal becomes fainter.

Several techniques, all of which require one or more extra images for the estimator, exist for creating focal plane intensity diversity to calculate the electric field. The self-coherent camera (SCC)³ is an estimation technique for coronagraphs with focal plane phase masks. Pinholes outside the nominal beam radius at the Lyot stop produce an interference pattern in the image that is used to calculate the electric field. Another technique is coronagraphic focal-plane wave-front estimation for exoplanet

*Address all correspondence to: A. J. Eldorado Riggs, E-mail: ariggs@princeton.edu

detection (COFFEE),⁴ which utilizes a maximum a posteriori approach to estimate pupil plane aberrations and the bias signal. COFFEE introduces large phase aberrations at the DM to create diversity in the image plane. In this paper, we use the most tested and widely applicable estimation technique in which small, pair-wise probes are actuated on the DM to create image plane diversity.⁵

High-contrast wavefront correction is an iterative process. With each electric field estimate, the controller suppresses as much starlight as possible in the dark hole. The contrast is then remeasured and more correction iterations are used until sufficiently high contrast is reached. Our prior work utilized a Kalman filter (KF) to estimate the stellar electric field recursively during wavefront correction. In this paper, we explore the use of a nonlinear filter, the extended Kalman filter (EKF), to estimate recursively both the stellar electric field and the intensity bias during wavefront correction. Science targets such as exoplanets and disks will be incoherent with a star, so they will appear in the bias estimate. Therefore, nonlinear, recursive wavefront correction lets us build the best possible real-time estimate of our potential science targets while recovering high contrast. The KF and EKF formalisms used in this paper are equally applicable to the SCC, and COFFEE could be easily modified to allow recursive estimation.

The recursive estimation techniques in this paper are discussed in the context of space-based observatories but may also apply to some ground telescopes, in particular, those with extreme adaptive optics. If an observatory and the wavefront are stable enough for focal plane wavefront correction to function, then Kalman filtering should still be able to improve wavefront estimation by accounting for the model uncertainty of the system. Nonlinear estimation of the wavefront and bias is less robust to large uncertainties, however, and is most likely better suited for ultrastable space telescopes.

2 Review of Focal Plane Wavefront Correction

In this section, we describe the current progress in focal plane wavefront estimation and control. A longer discussion can be found in the paper by Groff et al.⁶ in this issue. We rederive important results to pose the problem and to establish the mathematical framework for the EKF derivation in Sec. 3.

2.1 Linear Focal Plane Wavefront Control

The first successful controller for focal plane wavefront correction was speckle nulling.⁷ In this estimation-free scheme, sinusoids with different phases are applied to the DM to suppress stellar speckles at the targeted spatial frequency. Speckle nulling requires many hundreds or thousands of correction iterations, too many for use in a space mission.

Model-based estimation and control enables faster correction. The first model-based controller was Electric Field Conjugation (EFC).⁵ EFC minimizes energy in the dark hole and uses Tikhonov regularization to prevent too large of a command from being sent to the DMs. Because the electric field varies with wavelength, field estimates at several wavelengths within a larger bandpass are required to achieve broadband correction. An alternative model-based controller, stroke minimization,⁸ minimizes DM actuation subject to a constraint on contrast. Stroke minimization has the same mathematical formula as EFC but provides a logical means of choosing the actuator regularization value at each correction iteration.⁶

Here, we derive the linearized electric field at the DM for use with the controller and estimator. Let $\tilde{E}_0(x, y)$ be the initial complex electric field at the DM including the incident field and the nominal complex aberrations on the DM, where (x, y) are coordinates in the plane of the DM. Let $\phi_{k-1}(x, y)$ be the total phase contribution of the DM at correction iteration $k - 1$, and let $\Delta\phi_k(x, y)$ be the perturbation of the DM phase at correction iteration k such that

$$\phi_{k-1}(x, y) = \sum_{j=1}^{k-1} \Delta\phi_j(x, y). \quad (1)$$

The phase at the DM is twice the surface height of the DM and scales inversely with wavelength λ (in meters). Since for small deformations we can approximate the DM surface as the sum of the normalized actuator influence function $f(x, y)$ times the displacement command $\Delta u_{k,q}$ (in meters) at each actuator q 's center location (x_q, y_q) , the perturbation phase at the DM is given by

$$\Delta\phi_k(x, y) = \frac{2}{\lambda} \sum_q^{N_{\text{act}}} \Delta u_{k,q} f(x - x_q, y - y_q), \quad (2)$$

where N_{act} is the number of DM actuators.

Assuming small perturbation commands to the DM, we can approximate the electric field leaving the DM, $\tilde{E}_k(x, y)$, with a first order Taylor series expansion about the most recent DM perturbation

$$\tilde{E}_k(x, y) = \tilde{E}_0(x, y) e^{i[\phi_{k-1}(x, y) + \Delta\phi_k(x, y)]}, \quad (3)$$

$$\approx \tilde{E}_0(x, y) e^{i\phi_{k-1}(x, y)} [1 + i\Delta\phi_k(x, y)]. \quad (4)$$

Nearly all coronagraphs can be modeled as a series of linear operators such as Fourier transforms, Fresnel propagations, and mask multiplications. Since the control and estimation methods presented here are general to all coronagraphs, we represent the propagation from the DM to the science camera by the linear operator $\mathcal{C}\{\cdot\}$ to obtain the focal plane electric field $E_k(\xi, \eta)$

$$\begin{aligned} E_k(\xi, \eta) &= \mathcal{C}\{\tilde{E}_k(x, y)\} \\ &\approx \mathcal{C}\{\tilde{E}_0(x, y) e^{i\phi_{k-1}(x, y)}\} \\ &\quad + \mathcal{C}\{i\tilde{E}_0(x, y) e^{i\phi_{k-1}(x, y)} \Delta\phi_k(x, y)\} \\ &= E_{k-1}(\xi, \eta) + \sum_q^{N_{\text{act}}} \Delta u_{k,q} \mathcal{C}\{i\tilde{E}_0(x, y) \\ &\quad \times e^{i\phi_{k-1}(x, y)} f(x - x_q, y - y_q)\} \\ &= E_{k-1}(\xi, \eta) + \sum_q^{N_{\text{act}}} \Delta u_{k,q} B_{k-1,q}(\xi, \eta), \end{aligned} \quad (5)$$

where (ξ, η) are coordinates in the image. The aberrated focal plane electric field before the new command at correction iteration k is

$$E_{k-1}(\xi, \eta) = \mathcal{C}\{\tilde{E}_0(x, y) e^{i\phi_{k-1}(x, y)}\}, \quad (6)$$

and the Jacobian of each actuator at the image is given by the function

$$B_{k-1,q}(\xi, \eta) = \mathcal{C}\{i\tilde{E}_0(x, y)e^{i\phi_{k-1}(x,y)}f(x-x_q, y-y_q)\}. \quad (7)$$

Detectors measure the intensity in finite-sized pixels, so the focal plane field is converted from the continuous coordinates (ξ, η) to the discrete indices (m, n) . The detector integrates the intensity over the whole pixel whereas the model just samples discretely at each pixel. With greater than Nyquist discretization (≥ 2 pixels per λ/D) of the PSF already required for wavefront correction, the effect of sampling the PSF at each pixel instead of integrating over that area is small. The discretized focal plane electric field is thus

$$E_{k,m,n} = E_{k-1,m,n} + \sum_q^{N_{\text{act}}} \Delta u_{k,q} B_{k-1,q,m,n}, \quad (8)$$

where we have implicitly defined the region as being only within the dark hole. To perform matrix operations on the discretized field, it is convenient to reshape the field into a vector of length N_{pix} , the number of dark hole pixels, such that

$$E_k = E_{k-1} + G_{k-1} \Delta u_k. \quad (9)$$

Both E_k and E_{k-1} have dimensions $N_{\text{pix}} \times 1$, the control Jacobian G_{k-1} has dimension $N_{\text{pix}} \times N_{\text{act}}$, and the vector of control commands Δu_k has dimension $N_{\text{act}} \times 1$.

Setting Eq. (9) equal to zero and solving for Δu_k , we obtain the command to minimize the dark hole electric field,

$$\Delta u_k = -\mathcal{R}\{G_{k-1}^L E_{k-1}\}, \quad (10)$$

where the superscript L gives the left pseudoinverse and $\mathcal{R}\{\cdot\}$ returns the real part. In both EFC and stroke minimization, the actuator command is damped to avoid singularities and becomes

$$\Delta u_k = -(G_{k-1}^* G_{k-1} + \alpha \mathbb{I})^{-1} \mathcal{R}\{G_{k-1}^* E_{k-1}\}, \quad (11)$$

where $*$ gives the conjugate transpose, α is the damping (i.e., regularization) value, and $\mathbb{I}$ is the identity matrix. The control Jacobian G_{k-1} is usually not updated (G_0 is used instead) to save computation time at the expense of somewhat slower correction. From here on we will use the notation G instead of G_0 and u_k instead of Δu_k for convenience. The errors in the estimate and model, ignored nonlinearities of the DM phase contribution, and the use of regularization cause the new, corrected field to have only a slight improvement in contrast after each control step. The correction is thus iterative, with a new DM command calculated and applied after each new estimate of the electric field.

2.2 Batch Process Pair-Wise Estimation

The model-based control techniques in the previous section require knowledge of the electric field in the dark hole. An estimation approach is thus needed to determine the field from intensity measurements in the science camera. Currently, the baseline estimation method for a coronagraphic space mission is pair-wise difference imaging as developed by Give'on et al.,⁵ which probes the image via small DM perturbations. It is the only model-based estimation scheme that has attained better than 10^{-8} contrast in laboratory experiments,⁹⁻¹³ all of which have been in the High Contrast Imaging Testbed (HCIT) at the Jet Propulsion Laboratory (JPL). Pair-wise estimation is

characterized by two notable features: it can be used with any coronagraph and it is fairly robust to model uncertainty. This method is described in detail in several papers,^{5,6,14} but we revisit the derivation to provide the mathematical foundation for our new work.

In this paper, we focus on monochromatic wavefront estimation. At both Princeton's High Contrast Imaging Laboratory (HCIL) and JPL's HCIT, broadband wavefront correction is accomplished by taking images at several smaller bandpasses within the whole bandpass, creating separate electric field estimates for each one, and then weighting the bandpasses equally within the controller.

In pair-wise estimation, shapes are actuated on the DM to probe the electric field in the dark hole. Give'on et al.¹⁴ explain one such method for choosing probe sets to modulate sufficiently the real and imaginary parts of the field. A separate image is taken for the positive and negative of each probe shape applied on the DM. Let u_i be the differential control signal for the i 'th positive probe shape. Then, the change in the focal plane electric field from a positive or negative probe is defined as $p_{i\pm} = \pm G u_i$. For convenience we will not explicitly write the dark hole pixel index for the probe field p , the electric field E , and the intensity I . The focal plane intensity at each dark hole for a given positive or negative probe shape is then

$$\begin{aligned} I_{k,i\pm} &= |E_k + p_{k,i\pm}|^2 + n_{k,i\pm} \\ &= |E_k|^2 + |p_{k,i}|^2 \pm 2\mathcal{R}\{E_k^* p_{k,i}\} + n_{k,i\pm}, \end{aligned} \quad (12)$$

where $n_{k,i\pm}$ is the zero-mean, Gaussian measurement noise. The difference of the positive and negative probed images is equal to twice the cross term,

$$\Delta I_{k,i} = I_{k,i+} - I_{k,i-} = 4\mathcal{R}\{E_k^* p_{k,i}\} + n_{k,i}, \quad (13)$$

where $n_{k,i} = n_{k,i+} - n_{k,i-}$ is the total noise having twice the variance of a single probed image. For a set of measurements from N_{pp} probe pairs, the measurement equation is

$$\begin{aligned} \begin{bmatrix} \Delta I_{k,1} \\ \vdots \\ \Delta I_{k,N_{pp}} \end{bmatrix} &= 4 \begin{bmatrix} \mathcal{R}\{p_{k,1}\} & \mathcal{I}\{p_{k,1}\} \\ \vdots & \vdots \\ \mathcal{R}\{p_{k,N_{pp}}\} & \mathcal{I}\{p_{k,N_{pp}}\} \end{bmatrix} \begin{bmatrix} \mathcal{R}\{E_k\} \\ \mathcal{I}\{E_k\} \end{bmatrix} \\ &+ \begin{bmatrix} n_{k,1} \\ \vdots \\ n_{k,N_{pp}} \end{bmatrix}, \end{aligned} \quad (14)$$

where $\mathcal{I}\{\cdot\}$ takes the imaginary part of the complex value. We rewrite Eq. (14) as

$$z_k = H_k x_k + n_k, \quad (15)$$

where the set of measurements is

$$z_k = \begin{bmatrix} \Delta I_{k,1} \\ \vdots \\ \Delta I_{k,N_{pp}} \end{bmatrix}, \quad (16)$$

the linear observation matrix is

$$H_k = 4 \begin{bmatrix} \mathcal{R}\{p_{k,1}\} & \mathcal{I}\{p_{k,1}\} \\ \vdots & \vdots \\ \mathcal{R}\{p_{k,N_{pp}}\} & \mathcal{I}\{p_{k,N_{pp}}\} \end{bmatrix}, \quad (17)$$

the state vector is

$$x_k = \begin{bmatrix} \mathcal{R}\{E_k\} \\ \mathcal{I}\{E_k\} \end{bmatrix}, \quad (18)$$

and the measurement noise vector is

$$n_k = \begin{bmatrix} n_{k,1} \\ \vdots \\ n_{k,N_{pp}} \end{bmatrix}. \quad (19)$$

The best estimate $\hat{x}_k$ of the field's real and imaginary parts is found by taking the left pseudoinverse of H_k

$$\hat{x}_k = H_k^L z_k, \quad (20)$$

which requires two probe pairs to be invertible and at least three probe pairs for a least-squares estimate to reduce error from measurement noise.

In addition to the pairs of probed images, we always take an unprobed image I_k to measure the current contrast level. The more critical role of the unprobed image is in the empirically-based estimate of the probe amplitude

$$|p_{k,i}| = \sqrt{\frac{I_{k,i+} + I_{k,i-}}{2}} - I_k. \quad (21)$$

As described by Give' on et al.,¹⁴ this technique mitigates several types of model error to enable faster and deeper correction. Even with a good model of the laboratory, the measured and modeled probe amplitudes can have different morphologies and differ by several times in magnitude. The phase of the probe is still calculated using the model.

In this batch process estimation, there is an implicit assumption that the wavefront is static. The estimator and controller can still create a dark hole as long as the electric field is static at the level of the contrast target over the course of a few correction iterations.

2.3 Recursive Pair-Wise Estimation with a Kalman Filter

Groff and Kasdin¹⁵ incorporated pair-wise estimation into a KF for better accuracy and robustness. The KF optimally utilizes the previous estimate, the expected model uncertainty, the control signal, and new measurements in the estimate calculation. Since the KF knows the previous estimate, it does not require a full, invertible set of new measurements. Therefore, the wavefront estimate can be updated with just three new images: one unprobed image and one pair of probed images. If the estimate can be as accurate with fewer images of the same exposure time, this technique can further increase the speed of wavefront correction.^{13,15,16} The KF formulation also permits a dynamic state, so the KF can correct a dynamic wavefront faster and more robustly than a batch process estimator.

2.4 Batch Estimation of Incoherent Light

Both formulations of pair-wise estimation (batch and recursive) can yield a batch estimate of the incoherent light intensity at each correction iteration k . The incoherent intensity estimate $\hat{I}_{\text{inco},k}$ at each pixel is

$$\hat{I}_{\text{inco},k} = I_k - |\hat{E}_k|^2, \quad (22)$$

where $\hat{E}_k$ is the estimated stellar electric field. We will not derive the variance for the starlight intensity here, but from Eq. (22) we can see that the incoherent intensity batch estimate has a higher variance than a single image. Both terms in Eq. (22) are susceptible to noise sources (shot, readout, and dark current noise), and the estimated starlight intensity is susceptible to model errors. To mitigate both model errors and measurement noise, it would be better to estimate the incoherent light recursively with a filter that can use previous data and appropriate weights on the error sources.

3 Recursive Estimation of Both Incoherent Light and the Stellar Wavefront

The batch and recursive pair-wise estimators in Secs. 2.2 and 2.3 have been tested and proven to work at high contrast for suppressing coherent, on-axis light in JPL's HCIT.^{13,14} The ultimate goal of wavefront correction, however, is to image faint sources that are incoherent with a star such as exoplanets and disks. During wavefront correction, the starlight speckles change as they are suppressed while the exoplanets and disks remain unchanged, so a recursive filter can form a better estimate of the incoherent light with each new set of images. By implementing Bayesian techniques to locate any exoplanets or disks in the incoherent image,¹⁷ we can better detect and characterize our science targets with just the correction images.

One possible approach to recursive incoherent estimation is to use another KF on the batch incoherent estimates from Eq. (22). While this method would let us use two linear estimators, it is inefficient because the estimates of the incoherent light and starlight are interdependent. To produce the best estimate of the incoherent light with all available data, we need to estimate the stellar electric field and incoherent intensity simultaneously in a nonlinear estimator.

With a nonlinear filter, one could attempt to utilize the true nonlinear phase dependence of the electric field on the DM surface,

$$E_k(\xi, \eta) = \mathcal{C}\{\tilde{E}_0(x, y)e^{i\phi_{k-1}(x, y)}e^{i\Delta\phi_k(x, y)}\}, \quad (23)$$

in the modeled propagation of the electric field, but we do not. Because opaque coronagraphic masks and field stops between the DM and camera block light, we cannot directly back-propagate our electric field estimate from the focal plane to the DM-plane. Maximum *a posteriori* methods such as COFFEE exist to solve this problem, but they are currently too slow computationally to implement in real time. The other reason for not utilizing the full nonlinear model and for not including more terms in the Taylor expansion in Eq. (4) is that the errors in our knowledge of $\mathcal{C}\{\cdot\}$, $\tilde{E}_0(x, y)$, $\phi_{k-1}(x, y)$, and $\Delta\phi_k(x, y)$ might outweigh the better accuracy from a higher-order model.

As a first step into nonlinear focal plane wavefront estimators, we derive an EKF as first shown by Riggs et al.¹⁸ In this paper, the EKF utilizes the same probing strategy as the KF does for easier performance comparison. The EKF has the advantage that it can utilize the unprobed image recursively as well.

3.1 Constructing the Extended Kalman Filter

We augment the original state vector in Eq. (18) to include the incoherent intensity at each pixel, $I_{\text{inco},k}$,

$$x_k = \begin{bmatrix} \mathcal{R}\{E_k\} \\ \mathcal{I}\{E_k\} \\ I_{\text{inco},k} \end{bmatrix}. \quad (24)$$

The most general EKF measurement vector is the actual set of images taken. This formulation allows the use of unpaired probes or multi-DM probes, which we leave for future work. Here we use the same set of images as in pair-wise estimation such that z at each dark hole pixel consists of the unprobed image I_k and the $2 \times N_{pp}$ probe images,

$$z_k = \begin{bmatrix} I_k \\ I_{k,1+} \\ I_{k,1-} \\ \vdots \\ I_{k,N_{pp}+} \\ I_{k,N_{pp}-} \end{bmatrix} = h(x_k) + n_k + \mathcal{O}\{\Delta\phi_k^2\}, \quad (25)$$

where $h(x_k)$ is the nonlinear measurement function and $\mathcal{O}\{\Delta\phi_k^2\}$ is the model error from ignored higher-order terms of the DM-phase Taylor series. The additive measurement noise vector n_k consists of readout noise, photon shot noise, and dark current. By not performing pair-wise differencing of the probed images, the terms of order $\Delta\phi_k^2$ ignored in Eq. (4) no longer cancel and appear in the measurement. The quadratic, approximate measurement function is

$$\begin{aligned} h(x_k) &= \begin{bmatrix} |E_k|^2 + I_{\text{inco},k} \\ |E_k + Gu_1|^2 + I_{\text{inco},k} \\ |E_k - Gu_1|^2 + I_{\text{inco},k} \\ \vdots \\ |E_k + Gu_{N_{pp}}|^2 + I_{\text{inco},k} \\ |E_k - Gu_{N_{pp}}|^2 + I_{\text{inco},k} \end{bmatrix} \\ &= \begin{bmatrix} (\mathcal{R}\{E_k\})^2 + (\mathcal{I}\{E_k\})^2 + I_{\text{inco},k} \\ (\mathcal{R}\{E_k + Gu_1\})^2 + (\mathcal{I}\{E_k + Gu_1\})^2 + I_{\text{inco},k} \\ (\mathcal{R}\{E_k - Gu_1\})^2 + (\mathcal{I}\{E_k - Gu_1\})^2 + I_{\text{inco},k} \\ \vdots \\ (\mathcal{R}\{E_k + Gu_{N_{pp}}\})^2 + (\mathcal{I}\{E_k + Gu_{N_{pp}}\})^2 + I_{\text{inco},k} \\ (\mathcal{R}\{E_k - Gu_{N_{pp}}\})^2 + (\mathcal{I}\{E_k - Gu_{N_{pp}}\})^2 + I_{\text{inco},k} \end{bmatrix} \\ &= \begin{bmatrix} (x_k[1])^2 + (x_k[2])^2 + x_k[3] \\ (x_k[1] + \mathcal{R}\{G\}u_1)^2 + (x_k[2] + \mathcal{I}\{G\}u_1)^2 + x_k[3] \\ (x_k[1] - \mathcal{R}\{G\}u_1)^2 + (x_k[2] - \mathcal{I}\{G\}u_1)^2 + x_k[3] \\ \vdots \\ (x_k[1] + \mathcal{R}\{G\}u_{N_{pp}})^2 + (x_k[2] + \mathcal{I}\{G\}u_{N_{pp}})^2 + x_k[3] \\ (x_k[1] - \mathcal{R}\{G\}u_{N_{pp}})^2 + (x_k[2] - \mathcal{I}\{G\}u_{N_{pp}})^2 + x_k[3] \end{bmatrix}, \quad (26) \end{aligned}$$

where $x_k[m]$ represents the m 'th element of vector x_k , and u_i is the additive DM control signal for the positive probe shape i .

To derive the EKF, we first redefine the true dynamic state equation at each image plane pixel as

$$x_k = \Phi x_{k-1} + \Gamma u_{k-1} + \Lambda w_{k-1}, \quad (27)$$

where Φ is the state transition matrix, Γ is the real-valued control Jacobian, and Λ is the disturbance Jacobian. The variable w_{k-1} is random process noise; it is included in the model to accommodate model errors and random, unknown disturbances. We treat our system as static, so Φ is just the identity matrix. The only source of change is from the DMs, which means the only source of model error is in our knowledge of the DM response. Thus, $\Lambda = \Gamma_{\text{true}} - \Gamma_{\text{model}}$ and $w_{k-1} = u_{k-1}$. From here on Γ will mean Γ_{model} . The third row of Γ is zeroes because the incoherent light is not modulated by the DMs. (Only the PSF core is observable for faint incoherent sources, and high-order wavefront correction primarily changes the wings of the PSF.)

The EKF minimizes both the error of the state estimate and the state covariance estimate. The state covariance P is defined as the expectation value $E[\cdot]$ of the outer product of the error in the state estimate,

$$P_k = E[(x_k - \hat{x}_k)(x_k - \hat{x}_k)^T]. \quad (28)$$

In the first two equations of the EKF, the dynamics of the system are used to propagate the previous estimates of the state and state covariance to the current time step. Following the derivation and notation by Stengel,¹⁹ the state estimate time update is

$$\hat{x}_k(-) = \Phi \hat{x}_{k-1}(+) + \Gamma u_{k-1}, \quad (29)$$

and the covariance estimate time update is

$$P_k(-) = \Phi P_{k-1}(+) \Phi^T + Q_{k-1}, \quad (30)$$

where Q_{k-1} is the process noise matrix. The signifier $(-)$ means $\hat{x}_k$ or P_k is the model-based time update, and the signifier $(+)$ means it is the measurement-updated estimate for that correction iteration. We assume that the unknown disturbance Λw_{k-1} is Gaussian and zero mean so it should not change the expected value $\hat{x}_k$. The process noise covariance matrix is given by

$$Q_{k-1} = \Lambda E[w_{k-1} w_{k-1}^T] \Lambda^T. \quad (31)$$

The last stage of the EKF is to improve the estimates with new data in the measurement update equations,

$$\hat{x}_k(+) = \hat{x}_k(-) + K_k \{z_k - h[\hat{x}_k(-)]\}, \quad (32)$$

$$P_k(+) = [I - K_k H_k] P_k(-), \quad (33)$$

where the Kalman gain K_k optimally balances the weighting of the model error and old data versus the new measurements. The Kalman gain is defined as

$$K_k = P_k(-) H_k^T [H_k P_k(-) H_k^T + R_k]^{-1}, \quad (34)$$

where R_k is the measurement noise covariance matrix, which we discuss in more detail in Sec. 3.2. H_k is the observation matrix linearized about the state time update,

Table 1 Dimensions of variables for the EKF. $N_z = 1 + 2N_{pp}$.

Variable	Representation	Dimension
State estimate	$\hat{x}_k = \begin{bmatrix} \mathcal{R}\{E_k\} \\ \mathcal{I}\{E_k\} \\ I_{\text{Inco},k} \end{bmatrix}$	3×1
Intensity measurements	z_k	$N_z \times 1$
Sensor noise	n_k	$N_z \times 1$
DM commands	u_k	$(N_{\text{DMs}} \times N_{\text{act}}) \times 1$
Process noise	w_k	$(N_{\text{DMs}} \times N_{\text{act}}) \times 1$

$$H_k = \left. \frac{\partial h(x)}{\partial x} \right|_{x=\hat{x}_k(-)} = \begin{bmatrix} 2x[1] & 2x[2] & 1 \\ 2(x[1] + \mathcal{R}\{G\}u_1) & 2(x[2] + \mathcal{I}\{G\}u_1) & 1 \\ \vdots & \vdots & \vdots \\ 2(x[1] - \mathcal{R}\{G\}u_{N_{pp}}) & 2(x[2] - \mathcal{I}\{G\}u_{N_{pp}}) & 1 \end{bmatrix} \bigg|_{x=\hat{x}_k(-)} \quad (35)$$

In summary, the five EKF equations for our formulation are

$$\hat{x}_k(-) = \hat{x}_{k-1}(+) + \Gamma u_{k-1}, \quad (36)$$

$$P_k(-) = P_{k-1}(+) + Q_{k-1}, \quad (37)$$

$$K_k = P_k(-)H_k^T[H_k P_k(-)H_k^T + R_k]^{-1}, \quad (38)$$

$$\hat{x}_k(+) = \hat{x}_k(-) + K_k\{z_k - h[\hat{x}_k(-)]\}, \quad (39)$$

$$P_k(+) = [\mathbb{I} - K_k H_k] P_k(-). \quad (40)$$

We summarize the variables used in these equations in Table 1, and we list the matrices and their definitions in Table 2. The estimate is performed separately at each dark hole pixel to avoid the use of extremely large matrices.

3.2 Sensor and Process Noise

In any KF, R_k and Q_{k-1} are the tuning parameters. The sensor noise matrix is defined as the expectation value of the outer product of the measurement noise vector,

$$R_k = E[n_k n_k^T]. \quad (41)$$

We can calculate the value of R_k based on camera measurements, so only Q_k needs to be tuned. The main sources of measurement noise are dark current, readout noise, and photon shot noise. The total variance in analog-digital units (ADU, or counts) expected at each pixel is

$$\sigma_{\text{total}}^2 = [c_k f_{\text{star}} t_{\text{exp}} + \sigma_{\text{ron}}^2 + (s_{\text{dark}} t_{\text{exp}})] / (g n_{\text{exp}}), \quad (42)$$

where g is the gain of the detector in photoelectrons/ADU, c_k is the average measured contrast in the dark hole, f_{star} is the peak flux of the starlight in photoelectrons/second, t_{exp} is the exposure time per frame, σ_{ron}^2 is the variance of the readout noise in photoelectrons, s_{dark} is the dark current rate in photoelectrons/second, and n_{exp} is the number of exposures averaged to make an image. The contrast across the dark hole in either the probed or unprobed images are relatively uniform, so we use the same matrix R_k at each pixel. We still use separate values for probed or unprobed images since the probed images have more light.

Table 2 Dimensions of matrices for the EKF.

Matrix	Representation	Dimension
Linearized state response	$\Phi = \mathbb{I}$	3×3
Nonlinear observation	$h(x)$	$N_z \times 1$
Linearized observation	$H_k = \left. \frac{\partial h(x)}{\partial x} \right _{x=\hat{x}_k(-)}$	$N_z \times 3$
Linearized complex response of probing DM	G	$1 \times N_{\text{act}}$
Linearized response of probing DM	$\Gamma = \begin{bmatrix} \text{Re}\{G[1]\} & \cdots & \text{Re}\{G[N_{\text{act}}]\} \\ \text{Im}\{G[1]\} & \cdots & \text{Im}\{G[N_{\text{act}}]\} \\ 0 & \cdots & 0 \end{bmatrix}$	$3 \times N_{\text{act}}$
Disturbance response	$\Lambda = \Gamma$	$3 \times N_{\text{act}}$
State covariance (time update)	$P_k(-) = E\{[x_k - \hat{x}_k(-)][x_k - \hat{x}_k(-)]^T\}$	3×3
State covariance (measurement update)	$P_k(+) = E\{[x_k - \hat{x}_k(+)] [x_k - \hat{x}_k(+)]^T\}$	3×3
Process noise	$Q_k = \Lambda E[w_k w_k^T] \Lambda^T$	3×3
Sensor noise	$R_k = E[n_k n_k^T]$	$N_z \times N_z$
Kalman gain	K_k is computed	$3 \times N_z$

The noise from image to image is uncorrelated, so R_k is a diagonal matrix. This means that each diagonal entry r_k in R_k is simply the variance in units of contrast

$$r_k = \sigma_{k,\text{total}}^2 / (f_{\text{star}} t_{\text{exp}}). \quad (43)$$

The sensor noise matrix is then

$$R_k = \begin{bmatrix} r_{k,\text{unpr}} & & & 0 \\ & r_{k,\text{pr}} & & \\ & & \ddots & \\ 0 & & & r_{k,\text{pr}} \end{bmatrix}, \quad (44)$$

where $r_{k,\text{unpr}}$ is for the unprobed image and $r_{k,\text{pr}}$ is for the probed images.

We must include a nonzero process noise Q_{k-1} in the covariance estimate extrapolation because of the uncertainty in the control step. Although we assume that the DM does not modulate the incoherent light, we must still include process noise to prevent the filter from converging quickly to an incorrect value. We scale the process noise for the third state with the average incoherent intensity estimate. Without location-specific information of the process noise, we assign the same Q_{k-1} matrix to each image plane pixel. Similarly, we have no way of distinguishing if the real or imaginary parts of the starlight should have more or less model error, so we set those covariance values as equal. For each pixel at correction iteration k , we thus use the process noise matrix

$$Q_{k-1} = \begin{bmatrix} q_0 |\hat{E}_{k-1}|_{\text{avg}}^2 & 0 & 0 \\ 0 & q_0 |\hat{E}_{k-1}|_{\text{avg}}^2 & 0 \\ 0 & 0 & q_3 (\hat{I}_{\text{inco},k-1})_{\text{avg}}^2 \end{bmatrix}, \quad (45)$$

where q_0 and q_3 are constants used to tune the relative values. There should also be off-diagonal elements in Q_k since the model errors of the states are cross-correlated and there is unmodeled interactuator coupling, but we nevertheless keep Q_{k-1} diagonal to avoid dropping rank before the inversion in the Kalman gain calculation. With good tuning, the values of $P_k(+)$ tend to show no cross-correlation (zero values) among the electric field and incoherent states but a slight cross-correlation (about 10% of the diagonal values) between the real and imaginary parts of the field. Including this nonzero off-diagonal term in Q_{k-1} did not change performance of the EKF, and was therefore not included for all the tests reported in Sec. 5.

3.3 Iterated Extended Kalman Filter

When using pair-wise differencing for the starlight measurements, the linear observation matrix H_k is correct to third order in the DM-phase Taylor series expansion. In our EKF formulation without differencing, the quadratic terms no longer cancel. The linearization of the observation $h(x)$ in Eq. (35) thus depends on the current state, so it is necessary to use the initial, inaccurate time update $\hat{x}_k(-)$ as the linearization point.

A large body of research already exists to address nonlinear errors when implementing an EKF. Here we try the simplest improvement, which is to iterate the EKF (known as an IEKF) to mitigate nonlinearities. The main error in H_k (and

subsequently in K_k , $\hat{x}_k(+)$, and $P_k(+)$) comes from the linearization about the model-based time update $\hat{x}_k(-)$, but after running the EKF [Eqs. (36) to (40)] we have a more accurate estimate of the state available. Using $\hat{x}_k(+)$ as the new linearization point for an updated H_k , the IEKF recomputes K_k , $\hat{x}_k(+)$, and $P_k(+)$. There is now an even better estimate of the state, and this process of iterating the EKF can be repeated until the state estimate converges on a solution. Defining the subscripts for the EKF iterations as $j = 0, 1, 2, \dots, N_{\text{it}}$, we follow the notation of Gelb²⁰ and Simon²¹ and write the IEKF equations as

$$H_{k,j} = \left. \frac{\partial h(x)}{\partial x} \right|_{x=\hat{x}_{k,j}(+)}, \quad (46)$$

$$K_{k,j} = P_k(-) H_{k,j}^T [H_{k,j} P_k(-) H_{k,j}^T + R_k]^{-1}, \quad (47)$$

$$\begin{aligned} \hat{x}_{k,j+1}(+) &= \hat{x}_k(-) + K_{k,j} \{z_k - h[\hat{x}_{k,j}(+)] \\ &\quad - H_{k,j} [\hat{x}_k(-) - \hat{x}_{k,j}(+)]\}, \end{aligned} \quad (48)$$

$$P_{k,j+1}(+) = [\mathbb{I} - K_{k,j} H_{k,j}] P_k(-). \quad (49)$$

We initialize the IEKF with

$$\hat{x}_{k,j=0}(+) = \hat{x}_k(-), \quad (50)$$

$$P_{k,j=0}(+) = P_k(-), \quad (51)$$

and then iterate the filter by updating Eqs. (46) to (49) to converge on a better state estimate at the k 'th control step. If $N_{\text{it}} = 0$, the IEKF simplifies back to the EKF. In Sec. 5.3.2, we test several values of N_{it} and determine the minimum value to converge on an estimate.

4 Computational Complexity

Space-based observatories have limited computing power, so it is important to compare the computational complexity of the estimators proposed. In pair-wise estimation, the state can be estimated separately at each image-plane pixel, which means that only small matrices are required but the calculations must be repeated thousands of times. Here, we derive estimates of the complexity for each estimator based on the number of matrix multiplications (and divisions) required per pixel. Because there are many possible methodological (such as taking another image during calculations) or algorithmic (such as using alternate forms of equations) approaches to reduce the effective complexity of the estimators, we perform just a simple analysis as a starting point for comparison. We leave out the recalculation of the observation matrix, H_k , for each estimator because calculating the new values of $\mathcal{R}\{G\}u_{k,i}$ and $\mathcal{I}\{G\}u_{k,i}$ is common to all the estimators and requires the square root calculation in Eq. (21). It should also be noted that the recursive estimators require more memory, but that is a separate issue from the processing speed considered here.

Table 3 shows the relative complexity of each estimator in terms of floating point multiplications required per pixel. The batch process calculation is based on Eq. (20) assuming that the minimum of two probe pairs is used. The total number of multiplications is 26.

The KF equations have the same form as Eqs. (36) to (40) except the state estimate update has a linear observation

Table 3 Number of scalar multiplications (and divisions) required per image-plane pixel for each estimator. The number of probe pairs used is shown in parentheses. BP stands for batch process, and N_{it} is the number of EKF iterations.

Estimator	Number of multiplications
BP (2p)	26
KF (1p)	19
KF (2p)	46
EKF (1p)	150
EKF (2p)	360
IEKF (1p)	$150 + 159N_{it}$
IEKF (2p)	$360 + 375N_{it}$

$$\hat{x}_k(+) = \hat{x}_k(-) + K_k[z_k - H_k\hat{x}_k(-)]. \quad (52)$$

The KF has only two states, so the dimensions of the KF matrices are the same as for the EKF in Table 2 except each 3 should be a 2. The most computationally expensive calculation in the KF is the multiplication of Γu_{k-1} in the time update of the state estimate, which requires $2N_{act}$ multiplications. For WFIRST-AFTA, this would be about three thousand. The controller should already have performed this calculation to choose the optimal DM command, so we assume that it adds no complexity to the estimator. We find that the 1-pair KF requires 19 multiplications, so it is slightly less computationally expensive than the batch process. The 2-pair KF, requiring 46 multiplications, needs fewer than twice the number of computations in the 2-pair batch process.

The EKF as listed in Eqs. (36) to (40) requires many more calculations than the KF because of the longer state and measurement vectors. We find that the 1-pair EKF requires 150 multiplications and the 2-pair EKF requires 360. These numbers could be reduced by exploiting the sparsity of the matrices and not performing a brute force matrix inversion in the Kalman gain calculation. Iterating the filter requires several more calculations per EKF iteration because of the extra multiplication in Eq. (48). The IEKF requires on the order of a thousand multiplications at each of the few thousand pixels. Since the inversion in Eq. (11) of the controller can be precomputed, the IEKF and controller would each require on the order of a million calculations per correction iteration. We conclude that the IEKF should still be feasible for the limited computing power available on a space observatory.

5 Experimental Results

In this section, we present experimental validation of our EKF and IEKF formulations in the HCIL at Princeton. We use the stroke minimization controller with fixed settings, so that any variation in performance should arise from estimator. We compare the EKF and IEKF results to those for the batch process estimator and KF with and without incoherent sources present. Finally, we identify the limitations in our lab to be addressed in future work.

5.1 High Contrast Imaging Laboratory at Princeton

In the HCIL, we utilize shaped pupil (SP) coronagraphs to generate high contrast. Our layout, as shown in Fig. 1, uses as few optics as possible to enable easier alignment and introduce fewer optical aberrations. We inject monochromatic, 635-nm laser light directly from a fiber (launch 1) as the simulated stellar point source in the nominal experimental configuration. The 60-in. focal length of the off-axis parabola (OAP) allows us to approximate the central part of the beam as uniform. The collimated beam reflects off two Boston Micromachines Kilo-DMs and a fold mirror in series before reaching a transmissive, 10-mm diameter SP. The SP used in this paper is the freestanding Ripple3 design described by Belikov et al.⁹ and Kasdin et al.²² and shown in Fig. 2(a). The apodized PSF has a theoretical contrast of 3×10^{-10} from $4 - 40\lambda/D$ over symmetric 90° sectors as shown in Fig. 2(b), and the empirical, uncorrected PSF in the HCIL is shown in Fig. 2(c). The second and final OAP focuses the beam onto a focal plane mask (FPM), which is used only as a field stop for better dynamic range on the detector. Two achromatic lenses then reimage the stopped-down PSF onto a CCD camera.

To test our estimators in the presence of incoherent sources, we inject additional laser light at either of two locations on our bench as shown in the dashed boxes in Fig. 1. To create an exoplanet, we insert a beamsplitter in front of the original fiber source and place fiber launch 2 to reflect into the same beam path. To eliminate any dispersion or path difference errors for the star, launch 1 becomes the planet and launch 2 becomes the star. This is the simplest configuration to add an off-axis source, but the beamsplitter creates additional aberrations and strong polarization dependence. We adjust the planet intensity by using a separate laser source, and we can reposition the planet and star by translating the fiber heads. To approximate a flat zodiacal signal, we place another fiber tip (launch 3) approximately half a meter from the camera. The core of the expanding Gaussian beam is approximately uniform over the detector from this distance.

We block most of the stellar PSF with a field stop and perform wavefront correction in a subset of the transmitted region. The FPM passes light in symmetric areas from a radius of

Fig. 1 Diagram of the HCIL configuration. Dashed boxes show the modified configurations with additional fiber launches to introduce incoherent sources. The beamsplitter and fiber launch 2 are used to inject an off-axis, planet-like source. Fiber launch 3 adds a zodiacal, flat background in the focal plane.

Fig. 2 (a) The Ripple3 shaped pupil used in the HCIL along with (b) its normalized design PSF on a log scale. The ideal average contrast is 3×10^{-10} from $4 - 40\lambda/D$ over symmetric 90 deg sectors. (c) The uncorrected PSF as measured in the HCIL is shown with the same spatial scaling but a shorter log stretch.

$5 - 11\lambda/D$ over 90 deg sectors as shown in Fig. 3(a). The nominal aberrations set an average starting contrast of 6.51×10^{-5} as shown in Fig. 3(b), in which correction quickly gets the contrast below 10^{-6} and then slowly approaches the final achievable contrast of about 10^{-7} . In these experiments, we corrected only small, rectangular dark holes in the image plane with $\xi \in [-10, -7; 7, 10]\lambda/D$ and $\eta \in [-2, 2]\lambda/D$ as shown in the corrected PSF of Fig. 3(c). When we correct a larger region, we cannot reach as high of a contrast value.

To compare the relative performance of different estimators, we need to distinguish testbed fluctuations from algorithmic performance. If we perform separate correction runs in our testbed on the same day, we can safely compare them. Otherwise, the optics can drift out of alignment and degrade the correction performance. Because each correction run takes approximately half an hour, we have time for only one or two trials with each estimator when performing comparisons.

Groff et al.⁶ derive the dependence of the electric field estimate variance on the different sources of noise, and we summarize that result here. If the pair-wise probe amplitudes, $|p_{k,i}|$, are

much greater than the nominal field amplitude, then the photon shot noise from the starlight can be eliminated. Large probe amplitudes also reduce, but do not eliminate, the variance in the E-field estimate from readout noise and incoherent-light shot noise. If the probe amplitudes are too large, estimate error is introduced from ignored nonlinearities and model error. We manually tuned the probe amplitudes to be as large as possible without slowing correction or limiting the achievable contrast, which gave a probe intensity of about 10^{-6} for measured contrast values around 10^{-7} .

5.2 Experimental Goals

We sought to determine the accuracy of the EKF and IEKF estimates in several scenarios. Because the true stellar electric field is unobtainable in experiment, we compared the different estimators with contrast correction speed, defined as measured starlight contrast versus total exposure time. We assumed that a better estimate enables a larger correction step. Since the incoherent signal is directly measurable, we quantified the accuracy

Fig. 3 An example of a typical correction run, in this case with 2 probe pairs and the IEKF. (a) The stopped-down, uncorrected initial image, shown on a log scale, had a contrast of 6.51×10^{-5} in the correction region. (c) The final, corrected image had a measured average contrast of 1.2×10^{-7} in the rectangular dark holes. (b) The contrast correction curve started fast before gradually approaching the highest achievable contrast. The average values are plotted for the measured contrast, estimated starlight contrast, and estimated incoherent contrast in the dark hole, as well as the average standard deviation from readout noise at each pixel, $\sigma_{\text{ron/pixel}}$.

of the incoherent estimate with the PSF correlation and estimated contrast of an injected exoplanet. Our goals were thus to:

1. Compare contrast correction speed and achievable contrast for the IEKF, EKF, KF, and batch process estimator.
2. Determine the minimum number of EKF iterations needed to get the IEKF state estimate to converge.
3. Compare the estimators' performance in the presence of zodi-like, incoherent light.
4. Determine the accuracy of the incoherent estimate by retrieving an injected planet signal.

5.3 Experiments without Additional Incoherent Sources

5.3.1 Correction speed comparisons

In this experiment, we compared the contrast correction curves for the various estimators as shown in Fig. 4. The 2-pair estimators used five images per correction iteration (1 unprobed and 2 pairs of probed images) and the 1-pair estimators used three images per correction iteration (1 unprobed and 1 pair of probed images). The same probe shapes were applied every correction iteration for the 2-pair estimators; for the 1-pair estimators the probes alternated after every correction iteration to modulate the field sufficiently. The IEKF performed five EKF iterations.

We measured correction speed as the total number of images since all exposures had equal length. The exposure time was the longest possible without saturating any pixels on the detector and gave a contrast resolution of 1.8×10^{-8} per count. The batch estimator provided the slowest overall correction. It achieved the worst final contrast and took more correction iterations to reach it. The 1-pair KF was second slowest but did eventually reach as high a contrast as the other recursive estimators. The 2-pair KF and EKF performed equally well, and the 2-pair IEKF performed slightly better than those two after 100 total images. The 1-pair EKF and IEKF started off slightly faster than the others and reached their best achievable contrast in less than 80% the number of images required for the 2-pair recursive estimators, thereby showing some benefit from fewer probes

Fig. 4 Comparison of contrast correction speed for several estimators without an intentionally injected incoherent source. Exposure time for each image was constant.

per correction iteration. Without repeated trials to average out variability, it is important not to draw too many conclusions from these results. In another run (not shown), for example, the 2-pair IEKF performed the same as the 2-pair KF and EKF. Nevertheless, we have observed that the batch process and 1-pair KF are always slower than the other methods and that the 1-pair EKF and IEKF are always slightly faster than all the 2-pair versions. The higher computational complexity of the recursive estimators is partially offset since they require fewer correction iterations to reach a given contrast level. We have also demonstrated the viability of EKF and IEKF formulations that do not require image differencing.

After reaching the best achievable contrast, each correction curve started to worsen slightly (but did stay at or below about 2×10^{-7} contrast). This might have been from the controller being too aggressive or from the modeled control Jacobian not matching the true system well. In a space mission, the system would be better characterized and the controller could be made less aggressive near the ultimate contrast value to prevent divergence.

Although the 1-pair EKF and IEKF slightly outperformed the 2-pair recursive estimators, in later tests we used the 2-pair versions of the estimators. We found in various trials (not shown) that the 2-pair estimates were slightly less sensitive to optical misalignments than the 1-pair estimates. A space-based coronagraph would not suffer from the same problem because of sturdier mounts, better initial alignment, and greater stability.

5.3.2 Relative accuracy of the estimators

Here, we compare the average contrast of the estimated starlight or incoherent light for each estimator. This approach can reveal a net bias but averages out the Gaussian noise of individual pixels. To eliminate variations in images between different correction runs, we ran the estimators on a set of saved images from a 2-pair IEKF correction run. There was no intentionally injected incoherent source. The only difference from using stored data instead of real-time data to feed the estimator is that the control signal was predetermined, but this did not change the accuracy of the estimator. By using the same images for all estimators, we decoupled correction speed from estimator accuracy. While we cannot know the true state values, we infer that if most of the estimators are close to a value, then it is most likely the true value.

All the estimators except the EKF gave almost exactly the same average starlight contrast values in Fig. 5(a). The EKF exhibited a large bias and mistook much of the starlight for incoherent light, as shown in Fig. 5(b). The IEKF eliminated the starlight estimate bias with just one EKF iteration. The batch estimate started to exhibit more fluctuations than the other estimates during the last half of the correction, probably because the batch estimator did not utilize previous estimates to average out read noise.

The EKF estimates in Fig. 5 were significantly biased, and yet the contrast correction speed in Fig. 4 clearly shows that the EKF was no slower or less robust at wavefront correction than the IEKF or KF. Somehow a net bias of the estimate at all the pixels did not degrade performance, whereas random errors from noise on batch estimates do slow the correction. This result contradicts our premise that a more accurate estimate yields faster wavefront correction.

The starlight intensity estimate in Fig. 5(a) was approximately 1×10^{-7} below the measured contrast, and the incoherent

Fig. 5 (a) Comparison of the average starlight intensity estimate for each 2-pair estimator. (b) Comparison of the average incoherent intensity estimate for each estimator. N_{it} is the number of EKF iterations performed in the IEKF. All estimators in this case used the same saved set of correction run images to allow a direct comparison of their quality with the exact same data. The measured intensity is included for reference.

estimate converged to this level in Fig. 5(b). The incoherent estimate's structure in Fig. 6(b) matched neither that of the coherent estimate in Fig. 6(a) nor that of the nearly flat probe signal, so it was not an artifact of pair-wise estimation. The incoherent signal was also not random, meaning it was not attributable to read noise. We conclude that the incoherent estimate was a true signal composed of stray light.

For the incoherent estimates shown in Fig. 5(b), the batch process and KF initially gave more accurate values when the measured contrast was still several times above the true value. In later correction iterations as the contrast curve started to level off, the recursive estimate of the IEKF was much less variable and ended up lower than for the batch process and KF. This suggests that the IEKF adequately filters out read noise. We see in Fig. 5(b) that the first EKF iteration helped the most in improving the IEKF incoherent estimate, a second EKF iteration helped a little more in the early iterations, and more than two EKF iterations produced no discernible change. It appears that two EKF iterations are sufficient for convergence of the IEKF estimates. From here on, we stop testing the un-iterated EKF in comparisons because of its heavily biased estimates.

5.4 Experiments in the Presence of a Flat, Incoherent Background

In this experiment, we compared the 2-pair batch process, KF, and IEKF estimators in the presence of a bright, zodi-like

Fig. 6 Estimates of (a) the starlight intensity at 5×10^{-8} average contrast and (b) the incoherent intensity at 6×10^{-8} average contrast at the end of the correction run in Fig. 3. The image plane is cropped to just the dark holes.

background as shown in Fig. 7(a). We chose a bright incoherent background at 2.45×10^{-5} contrast to give a noise floor worse than the previously achievable contrast floor. The average standard deviation at each pixel was 5.2×10^{-7} contrast from readout noise and incoherent-light shot noise, with readout noise alone giving a standard deviation of 8×10^{-8} . To verify the estimated starlight intensity, we acquired another image with the incoherent source turned off after each correction iteration, as in Fig. 7(b).

We compared the contrast correction speed for the batch process, KF, and IEKF in the presence of the flat background. The nominal incoherent signal at about 1×10^{-7} was still present with the zodi turned off, so the estimated starlight intensity should have been below the measured, background-off intensity by that amount. We define the nominal incoherent signal as the stray light when only the starlight laser is on. The KF and IEKF starlight estimates in Figs. 8(b) and 8(c), respectively, were indeed about 1×10^{-7} below the measured, zodi-less intensities, while the batch starlight intensity estimate in Fig. 8(a) was too bright by about 1×10^{-7} . Directly comparing the measured, zodi-less contrast curves in Fig. 8(d), we did not observe any definitive differences in correction speed or achievable contrast because of the large fluctuations between iterations.

Compared with the contrast correction curves without zodi in Fig. 4, the curves with bright background in Fig. 8 were much more erratic in their convergence. One possible explanation is that the much larger standard deviation in each measurement was corrupting the estimate quality. Another possibility is that the drifting laser power of the bright background was invalidating the assumption of a static state. Over the course of a correction run, we observed the laser power drift by 1.5%, corresponding to about 4×10^{-7} in contrast. Such a large drift in the allocation of intensity could explain the nonsmooth correction curves and the reduced achievable contrast for the bright background case.

5.5 Experiments to Recover a Faint Companion

In this set of experiments, we injected a faint, off-axis point source to mimic an exoplanet, and then we attempted to recover

Fig. 7 Final PSFs for the IEKF correction run when (a) the incoherent background is still on and (b) is temporarily turned off. The incoherent signal appears behind the field stop since it was placed downstream in the system. We used the signal behind the center of the mask ($\xi \in [-1.5, 1.5]\lambda/D$) to determine the net standard deviation at each pixel and the true average incoherent intensity.

Fig. 8 Correction runs with different estimators in the presence of 2.45×10^{-5} incoherent background light. The mean standard deviation per pixel from zodi shot noise and readout noise was at 5.2×10^{-7} contrast. In all measurements, there was always also an incoherent signal from scattered light at about 1×10^{-7} average contrast. (a) The batch estimator did well until reaching the noise floor. (b) The KF had slower initial correction speed but allowed correction below the noise floor. (c) The IEKF allowed the fastest correction to the noise floor and below. (d) A direct comparison of the measured intensity with the incoherent background temporarily turned off.

its signal. We inserted the star-planet simulator (the beamsplitter and fiber launch 2) into the testbed and used the 2-pair IEKF for each correction run. We first ran wavefront correction without injecting a planet to determine the differences in nominal performance. The initial PSF is shown in Fig. 9(a), and the final, corrected PSF is shown in Fig. 9(c). The estimated starlight correction curve in Fig. 9(b) was approximately as fast as the case

without the star-planet simulator in Fig. 3(b), and the final starlight estimate in Fig. 9(d) was comparable to the one in Fig. 6(a). The beamsplitter introduced a much larger nominal incoherent signal at 3.6×10^{-7} average contrast as shown in Fig. 9(e).

Next, we performed wavefront correction trials with an injected planet at four contrast levels, starting below the average incoherent background level and ending slightly above it. The

Fig. 9 Results from a 2-probe-pair IEKF correction run with the star-planet simulator installed but no planet injected. (a) The stopped-down, uncorrected initial image, shown on a log scale, had a contrast of 5.87×10^{-5} in the correction region. (c) The final, corrected image had a measured average contrast of 4.4×10^{-7} in the rectangular dark holes (3.1×10^{-7} in just the right-side dark hole). (b) The measured and estimated contrast correction curves. (d) The estimated coherent light after correction. (e) The final estimated incoherent signal, larger from the beamsplitter being placed in the system.

injected planet used a separate laser channel and was centered at approximately $(\xi, \eta) = (8.0, -0.6)$. We used the IEKF during real-time correction, saved those images, and reused them for the KF trials for a more direct estimator comparison. This strategy eliminated variations between trials from noise, hardware, or the controller.

We compared three different techniques to recover the planet signal from the images. The first was simple PSF subtraction (PS). After a correction run, we took one image with the planet laser on and another with it off. Subtracting these two images yielded the PS estimate. The second technique defined the planet signal as the batch process incoherent estimate (BPIE) from Eq. (22), where the KF supplied the estimated starlight intensity, $|E_k|^2$. We included the BPIE because it utilized the concept of coherence diversity (i.e., modulating the stellar electric field to distinguish the incoherent signal) without requiring the IEKF developed in this paper. The final method used the recursive incoherent estimate (RIE) from the IEKF as the planet signal. To isolate the planet signal for this analysis, we subtracted the IEKF's best estimate of the nominal incoherent signal, as shown in Fig. 9(e), from the BPIEs and RIEs. Because the incoherent background is unlikely to be fully subtractable in this manner during a space mission, the analysis in this section represents only a best-case scenario. Any nonuniformities or asymmetries in the zodiacal or exozodiacal light would make the background more difficult to subtract.

We quantified the quality of the planet signal with two metrics: the accuracy of the planet's contrast estimate and the two-dimensional (2-D) correlation of the planet signal to the expected PSF. The contrast estimate was calculated by translating the normalized, on-axis PSF from Fig. 2(c) to the planet's location, as shown in Fig. 10, and then scaling the template

PSF's contrast to fit the planet signal. The 2-D correlation, $\mathcal{C}$, which quantitatively compares the morphology of the signals, was calculated by correlating the template PSF, I_{temp} , to the extracted planet signal, $\hat{I}_{\text{planet}}$

$$\mathcal{C} = \frac{\sum_{s=1}^{N_{\text{FWHM}}} [I_{\text{temp}}(s) \hat{I}_{\text{planet}}(s)]}{\sqrt{\left[\sum_{s=1}^{N_{\text{FWHM}}} I_{\text{temp}}(s)^2 \right] \left[\sum_{s=1}^{N_{\text{FWHM}}} \hat{I}_{\text{planet}}(s)^2 \right]}}. \quad (53)$$

To avoid fitting to noise for both of these metrics, we used only the N_{FWHM} pixels located within the full width at half maximum (FWHM) of the template PSF.

Fig. 10 Normalized, on-axis PSF shifted to the planet location for use as a PSF template for the planet. Only the region within the full-width at half maximum (shown as the dotted line) was used to avoid fitting to noise.

Fig. 11 Planet signals obtained from the different techniques at the end of correction. Only the right dark hole region is shown. Each column is for a different planet contrast level. The first row is the planet PSF measured by averaging 10 images with the starlight laser off.

Figure 11 shows the planet signal from each routine at each contrast level. The first row displays the planet-only images, which are the averages of 10 exposures with the star laser off. The contrast values reported above this row yielded the least-squares fit of the normalized template PSF to the given signal. The average readout noise per pixel in these averaged images was 2.8×10^{-8} contrast, so the contrast values of 8×10^{-8} , 2.0×10^{-7} , 3.8×10^{-7} , and 6.6×10^{-7} had signal-to-noise ratios (SNRs) of approximately 2.9, 7.3, 13.9, and 24.1, respectively. The second, third, and fourth rows show the planet signals

obtained from the PS, BPIE, and RIE methods, respectively. Upon visual inspection, the RIE produced the least noisy planet signals and the best chance of a detection for the faintest planet setting.

For the planet signals in Fig. 11, we show the corresponding planet contrast estimates in Fig. 12(a) and correlation values in Fig. 12(b). There were not enough data points to determine if PS or the BPIE was more accurate than the other for either metric. The RIE was the only method to be within 5% of the measured contrast for each planet brightness; higher noise in the other methods yielded much higher error. The correlation of the RIE was always higher than for the other methods at a given planet contrast. The RIE performed well at isolating the incoherent signal, averaging out noise over many iterations, and producing an un-biased contrast estimate.

The previous analysis used only the final estimates and images from the correction runs. That is the optimal strategy for PSF subtraction, which needs a dimmer dark hole to reduce stellar shot noise in the planet signal, but it might be unnecessary for the BPIE and RIE. Therefore, we calculated the BPIE and RIE after each correction iteration to determine how early they could estimate the planet accurately. We did not take an extra image with the planet laser turned off after each correction iteration, so PS was not included in this comparison.

Figures 13(a) and 13(b) show the BPIE and RIE correlation values, respectively. Both methods reached their final values by about the fifth correction iteration, which corresponded to the dark hole reaching approximately 2×10^{-6} contrast. The exception was for the faintest planet intensity, for which the correlation values showed much higher variability. For each of the four planet settings, the average RIE correlation was significantly higher and had less variability over time. The faintest two planets merited the most attention as the most difficult ones to detect in a space mission. For the 2.0×10^{-7} contrast planet in correction iterations 5–50, the mean correlation was 77% for the BPIE and 92% for the RIE. Similarly for the faintest planet, the mean correlation was 37% for the BPIE and 70% for the RIE. The RIE was thus a much better tool for detecting faint planets than the BPIE.

Figures 14(a) and 14(b) show the BPIE and RIE planet contrast values, respectively. The measured contrast levels are

Fig. 12 (a) Planet contrast estimates and (b) correlation values corresponding to the estimated planet signals in Fig. 11.

Fig. 13 Correlation between the planet signal and the template PSF after each correction iteration for (a) the batch process incoherent estimate and (b) the recursive incoherent estimate. Each line corresponds to a different injected planet brightness. Except in the case of the faintest planet, the correlation settles near its final value by about the fifth correction iteration.

shown as dotted lines to reveal biases in the estimates. The values settled in about five correction iterations, and the RIE contrast values showed much less variability among correction iterations compared to the BPIE values. At each planet brightness, the BPIE contrast values started near the measured value but settled with a slightly negative bias. The RIE contrast values started positively biased then settled at the measured contrast. The initial positive bias in the RIE estimates likely arose from starting the IEKF at poor contrast levels, where nonlinearities in the model and observation are large. We previously observed this early-iteration estimate error in Fig. 5(b). It may therefore be beneficial to start running the IEKF at moderate contrast levels to avoid the large initial bias in the incoherent estimate.

We have demonstrated that the incoherent light estimate can be utilized for planet detection during wavefront correction. Because the RIE utilizes the whole history of images to average out noise, it gives the best planet contrast estimate and best

PSF correlation compared to PS and the BPIE. These results hold when the other background light can be fully subtracted or is nonexistent, neither of which is safe to assume for a space mission. Nonuniform background light makes planet detection harder for all three of these methods, but these results indicate that the RIE is best at separating starlight from incoherent light. PS may still be the best option if the dark hole speckles are stable long enough to image two different stars. However, if the dark hole does change significantly from slewing the telescope, then coherence diversity via the RIE should be the best option for detecting a companion and estimating its contrast.

5.6 Limitations in the High Contrast Imaging Laboratory

While we have demonstrated the use of an IEKF for generating a dark hole and simultaneously detecting a planet, there are

Fig. 14 Estimated planet contrast after each correction iteration for (a) the batch process incoherent estimate and (b) the recursive incoherent estimate. The measured contrast of the planet is shown for comparison as a dotted line. The BPIE planet contrast values had more variability than the RIE contrasts. The BPIE values settled below the measured contrast levels, while the RIE values started above the measured values before converging to them.

several error sources limiting our attainable contrast and dark hole size. Scattered light, particularly from the DM mounts, is contributing to a larger than desired background floor. Our final achievable contrast and the speed at which we reach it are also heavily dependent on our model accuracy. This is currently limited by our knowledge of the unactuated DM surface, the influence function shapes, and the nominal DM gains. Current work is directed at improving the lab characterization. Finally, the ultimate contrast we can measure is determined by the read noise in the camera, which is high at 4.9 ADU rms for a 40,000 ADU linear range. Future experiments will be performed with a near photon-counting detector. With these changes we expect to reach contrasts close to 10^{-8} , significantly improving our ability to characterize these algorithms for space missions.

6 Future Work

As we demonstrated by augmenting the state with the incoherent signal, the IEKF also allows the estimation of other system parameters by adding them to the state vector. This is known as parameter adaptive filtering, and it could improve the performance of both the estimator and the controller since both rely on the model. Because the IEKF is suboptimal from using a linearization of the nonlinear observation, we would like to implement other nonlinear filters such as the unscented KF to obtain less biased estimates.

In our HCIL wavefront correction trials, we showed that the recursive incoherent estimate provides a higher detection probability and better planet contrast estimate than PSF subtraction. We plan to investigate scenarios that are more representative of a space mission and to quantify the conditions for which either PSF subtraction or the recursive incoherent estimate is better suited. In particular, we would like to simulate trials at higher contrast, with expected levels of speckle dynamics, and at different SNRs. Comparisons assuming a dynamic optical system should also include advanced PSF subtraction techniques such as the locally optimized combination of images (LOCI)²³ and Karhunen–Loève image projection (KLIP);²⁴ Ygouf et al.²⁵ have already begun this analysis of reference differential imaging for WFIRST-AFTA.

Up to this point, we have used the same set of one unprobed image and N_{pp} probed image pairs per correction iteration because it provides sufficient diversity for wavefront correction. The original purpose of using pairs of probed images was to yield a linear relationship between the electric field and the field change from the probes, but we have just shown that the nonlinear measurement in Eq. (26) along with the IEKF works at least as well if not better. We can therefore modify our measurement equation z_k from Eq. (25) to a more general one comprised of any set of images. For example, this formulation allows the use of any unpaired probes or even probes on more than a single DM simultaneously. We plan to test other probe combinations in our recursive, nonlinear estimation scheme to further reduce the number of exposures.

7 Summary

With the prospect of the WFIRST-AFTA CGI flying in less than a decade, progress in efficient wavefront correction and exoplanet detection algorithms is critical for maximizing the science output of the mission. In our experiments in Princeton's HCIL, we demonstrated the effectiveness of the EKF for coronagraphic focal plane wavefront and bias estimation. We found that the EKF and IEKF provide faster wavefront correction by requiring

fewer images, and that the IEKF provides a better estimate of the incoherent signal by estimating it recursively along with the starlight. This provides an alternative methodology for separating exoplanet light from the stellar speckles and enables faster, more accurate planet detection. As the simplest nonlinear filters, the EKF and IEKF should be manageable for the computational limits of a space telescope's computer. By proving the viability of using a nonlinear, recursive estimator for focal plane wavefront sensing, we have enabled several possible paths for improvement such as different probe choices and parameter estimation. We found that the IEKF eliminates most of the bias error of the EKF, and we plan to implement more advanced filters for further improvement.

Acknowledgments

A. J. Eldorado Riggs is funded for this work by the NASA Space Technology Research Fellowship (NNX14AM06H). N. Jeremy Kasdin is supported by JPL subaward 1499194 under contract from NASA. Tyler D. Groff is funded by the NASA Nancy Grace Roman Technology Fellowship (NNX15AF28G).

References

1. W. A. Traub et al., "Science yield estimate with the WFIRST-AFTA coronagraph," submitted to *J. Astron. Telesc. Instrum. Syst.* (2016).
2. S. B. Shaklan and J. J. Green, "Reflectivity and optical surface height requirements in a broadband coronagraph. 1. Contrast floor due to controllable spatial frequencies," *Appl. Opt.* **45**, 5143–5153 (2006).
3. J. Mazoyer et al., "High-contrast imaging in polychromatic light with the self-coherent camera," *Astron. Astrophys.* **564**, L1 (2014).
4. B. Paul et al., "Simultaneous phase and amplitude retrieval with COFFEE: from theory to laboratory results," *Proc. SPIE* **9147**, 91479O (2014).
5. A. Give'on et al., "Broadband wavefront correction algorithm for high-contrast imaging systems," *Proc. SPIE* **6691**, 66910A (2007).
6. T. D. Groff et al., "Methods and limitations of focal plane sensing, estimation, and control in high contrast imaging," *J. Astron. Telesc. Instrum. Syst.* **2**(1), 011009 (2015).
7. P. Borde and W. Traub, "High-contrast imaging from space: speckle nulling in a low aberration regime," *Astrophys. J.* **638**, 488–498 (2006).
8. L. Pueyo et al., "Optimal dark hole generation via two deformable mirrors with stroke minimization," *Appl. Opt.* **48**(32), 6296–6312 (2009).
9. R. Belikov et al., "Demonstration of high contrast in 10% broadband light with the shaped pupil coronagraph," *Proc. SPIE* **6693**, 66930Y (2007).
10. D. Moody, B. Gordon, and J. Trauger, "Design and demonstration of hybrid Lyot coronagraph masks for improved spectral bandwidth and throughput," *Proc. SPIE* **7010**, 70103P (2008).
11. B. Kern et al., "Laboratory demonstration of phase induced amplitude apodization (PIAA) coronagraph with better than 10^{-9} contrast," *Proc. SPIE* **8864**, 88640R (2013).
12. E. Serabyn et al., "High-contrast imaging results with the vortex coronagraph," *Proc. SPIE* **8864**, 88640Y (2013).
13. A. Riggs et al., "Demonstration of symmetric dark holes using two deformable mirrors at the high-contrast imaging testbed," *Proc. SPIE* **8864**, 88640T (2013).
14. A. Give'on, B. Kern, and S. Shaklan, "Pair-wise, deformable mirror, image plane-based diversity electric field estimation for high contrast coronagraphy," *Proc. SPIE* **8151**, 815110 (2011).
15. T. D. Groff and N. J. Kasdin, "Kalman filtering techniques for focal plane electric field estimation," *J. Opt. Soc. Am. A* **30**(1), 128–139 (2013).
16. A. J. E. Riggs, N. J. Kasdin, and T. D. Groff, "Wavefront correction with Kalman filtering for the WFIRST-AFTA coronagraph instrument," *Proc. SPIE* **9605**, 960507 (2015).
17. N. J. Kasdin and I. Braems, "Linear and Bayesian planet detection algorithms for the Terrestrial Planet Finder," *Astrophys. J.* **646**, 1260–1274 (2006).

18. A. J. E. Riggs, N. J. Kasdin, and T. Groff, "Optimal wavefront estimation of incoherent sources," *Proc. SPIE* **9143**, 914324 (2014).
19. R. Stengel, *Optimal Control and Estimation*, Dover Publications, New York (1994).
20. A. Gelb et al., *Applied Optimal Estimation*, MIT Press, Cambridge, Massachusetts (1974).
21. D. Simon, *Optimal State Estimation: Kalman, H Infinity, and Nonlinear Approaches*, Wiley and Sons, Hoboken, New Jersey (2006).
22. N. J. Kasdin, R. J. Vanderbei, and R. Belikov, "Shaped pupil coronagraphy," *CR Phys.* **8**, 312–322 (2007).
23. D. Lafrenière et al., "A new algorithm for point-spread function subtraction in high-contrast imaging: a demonstration with angular differential imaging," *Astrophys. J.* **660**(1), 770 (2007).
24. R. Soummer, L. Pueyo, and J. Larkin, "Detection and characterization of exoplanets and disks using projections on Karhunen-Loève eigenimages," *Astrophys. J.* **755**(2), L28 (2012).
25. M. Ygouf et al., "Data processing and algorithm development for the WFIRST-AFTA coronagraph: reduction of noise free simulated images, analysis and spectrum extraction with reference star differential imaging," *Proc. SPIE* **9605**, 96050S (2015).

A. J. Eldorado Riggs is a PhD candidate in the Department of Mechanical and Aerospace Engineering at Princeton University. He received his BS degree in physics and mechanical engineering

from Yale University in 2011. His research interests include wavefront estimation and control, deformable mirrors, and coronagraph design for exoplanet imaging. He is a member of the American Astronomical Society and SPIE.

N. Jeremy Kasdin is a professor of mechanical and aerospace engineering and a vice dean of the School of Engineering and Applied Science at Princeton University. He is the principal investigator of Princeton's High Contrast Imaging Laboratory. He received his PhD from Stanford University in 1991. His research interests include space systems design, space optics and exoplanet imaging, orbital mechanics, guidance and control of space vehicles, optimal estimation, and stochastic process modeling. He is an associate fellow of the American Institute of Aeronautics and Astronautics and a member of the American Astronomical Society and SPIE.

Tyler D. Groff is an associate research scholar at Princeton University. He received his BS degree in mechanical engineering and astrophysics from Tufts University in 2007 and his PhD from Princeton University in 2012. His current research interests include exoplanet imaging, optomechanics, optical design, wavefront estimation and control, infrared spectroscopy, integral field spectrographs, deformable mirrors, and ferrofluidics. He is a member of the American Astronomical Society and SPIE.