

## Exoplanet detection techniques for direct imaging dark zone maintenance datasets

Susan F. Redmond<sup>a,b,\*</sup>, Laurent A. Pueyo<sup>c</sup>, Emiel H. Por<sup>b,c</sup>,  
Raphaël Pourcelot<sup>c</sup>, Iva Laginja<sup>b,d</sup>, Bryony Nickson<sup>c</sup>, Ananya Sahoo<sup>c</sup>,  
Meiji M. Nguyen<sup>c</sup>, N. Jeremy Kasdin<sup>e</sup>, Marshall D. Perrin<sup>b,c</sup>,  
Rémi Soummer<sup>b,c</sup> and Leonid Pogorelyuk<sup>f</sup>

<sup>a</sup>California Institute of Technology, Pasadena, California, United States

<sup>b</sup>Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California, United States

<sup>c</sup>Space Telescope Science Institute, Baltimore, Maryland, United States

<sup>d</sup>Université PSL, Sorbonne Université, Université Paris Cité, CNRS, LESIA, Observatoire de Paris, Meudon, France

<sup>e</sup>Princeton University, Princeton, New Jersey, United States

<sup>f</sup>Rensselaer Polytechnic Institute, Troy, New York, United States

**ABSTRACT.** With the commencement of the development of the Habitable Worlds Observatory, it is imperative that the community has an understanding of (1) the stability requirements for the observatory to inform the design and (2) the gains expected from post-processing to inform observing scenarios and science yield estimates. We demonstrate that a previously developed, photon-efficient dark zone maintenance (DZM) algorithm that corrects quasi-static wavefront error drifts using only science images is compatible with traditional post-processing techniques. Further, we augment the DZM algorithm to estimate the coherent and incoherent light separately and introduce three novel post-processing techniques that leverage the concurrent estimation of coherent and incoherent light. With the DZM algorithm implemented on the High-contrast imager for Complex Aperture Telescopes testbed at the Space Telescope Science Institute, artificial drifts are injected as a random walk on a set of deformable mirrors and are corrected with DZM. We present an injected fake planet recovered in post-processing using a variety of techniques, such as angular differential imaging (ADI), and three additional techniques: incoherent accumulated imaging, software-based coherent differential imaging, and coherent reference differential imaging. All post-processing techniques can recover an injected planet at the same contrast level as the dark zone background contrast ( $\sim 8 \times 10^{-8}$ ), and the ADI technique is shown to recover a  $4 \times 10^{-8}$  planet in a  $8 \times 10^{-8}$  dark zone. For a space-based observatory, this would mean that, if the instrument can reach a contrast level, we can maintain it and recover a planet that is undetectable in a single frame.

© The Authors. Published by SPIE under a Creative Commons Attribution 4.0 International License. Distribution or reproduction of this work in whole or in part requires full attribution of the original publication, including its DOI. [DOI: [10.1117/1.JATIS.10.4.049003](https://doi.org/10.1117/1.JATIS.10.4.049003)]

**Keywords:** exoplanets; direct imaging; coronagraph; wavefront error; drift; active optics

Paper 24048G received Apr. 19, 2024; revised Aug. 23, 2024; accepted Sep. 11, 2024; published Oct. 25, 2024.

\*Address all correspondence to Susan F. Redmond, [sredmond@caltech.edu](mailto:sredmond@caltech.edu)

## 1 Introduction

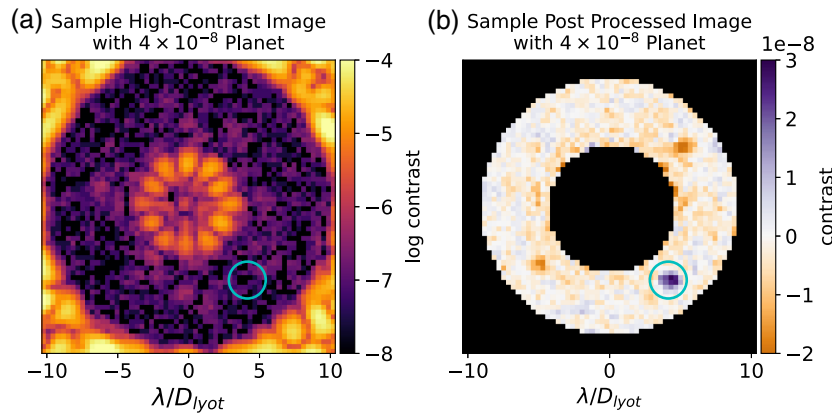
### 1.1 Motivation for Drift Control in Space-Based High-Contrast Imaging Systems

With the upcoming development of the Habitable Worlds Observatory (HWO) recommended by the Astro2020 decadal survey,<sup>1</sup> high-contrast imaging wavefront control and data analysis tools that are robust to changes in telescope wavefront are more important than ever. HWO aims to measure the spectra of  $\sim 25$  Earth-like exoplanets. This is an ambitious goal and will require the development of the hardware, wavefront correction algorithms, and post-processing algorithms. Directly imaging an exo-Earth is difficult due to the small angular separation between the host star and the planet, as well as how dim the planet is in comparison with the host star because exo-Earths are small, rocky, and not self-luminous. Coronagraphs are a critical hardware component in directly imaging exoplanets; they comprise a set of optical components that redistribute and mask the starlight in a region of the image where a planet is suspected to be [the dark zone (DZ)]. To further suppress the light in the DZ, deformable mirrors (DMs) are used to correct for manufacturing defects and misalignments of the optical surfaces. Due to the limited number of photons, directly imaging planets requires long integration times. The wavefront must be stable on the same time scale, which is often difficult in space due to time-varying wavefront errors from thermal gradients and other mechanical instabilities. The level of starlight suppression provided by the coronagraph masks and DMs needs to be maintained over the long science exposure times necessary for exo-earth detection and characterization. This stability can be obtained by designing an exquisitely stable telescope<sup>2,3</sup> and/or by stabilizing the contrast using the DMs.<sup>4,5</sup> Alternatively, post-processing methods, such as the ones used for James Webb Space Telescope (JWST) coronagraphs,<sup>6</sup> can be leveraged to calibrate contrast variations a posteriori. The Nancy Grace Roman Space Telescope (RST) Coronagraph Instrument is designed to be a space-based high-contrast imaging technology demonstration and consists of two DMs and a number of coronagraph options. RST will provide invaluable information on coronagraph and wavefront control performance in space as well as the opportunity to explore post-processing methods. This paper presents the first demonstration of algorithms that combine DZ maintenance and coronagraph post-processing, a combination that would benefit RST and likely be necessary for HWO.

For the purpose of this paper, the level of suppression achieved by the combination of the coronagraph and the DMs is referred to as the contrast and is calculated via

$$\text{contrast} = \frac{I_{\text{coron}}}{\max(I_{\text{direct}})}, \quad (1)$$

where  $I_{\text{coron}}$  is the intensity of the image with the coronagraph in place and  $I_{\text{direct}}$  is the intensity of the unocculted stellar point spread function (PSF). To image an exo-Earth, a high-contrast imaging system must achieve a contrast of at least  $10^{-10}$  (Ref. 7) and a contrast stability at least an order of magnitude better.<sup>8</sup> In this paper, we aim to address the technology gap of maintaining the wavefront in the presence of quasi-static wavefront error drifts. Addressing the stability problem using the coronagraph instrument is fundamentally limited by the finite number of photons from the surveyed host star. The fundamental limit for DZ maintenance can be quantified using the methods described in Pogorelyuk et al.<sup>9</sup> under the assumption of an unbiased drift estimator and perfect wavefront control (no lag and no fitting errors for instance). Under more realistic scenarios, a posterior calibration of the wavefront drift (e.g., post-processing) can still mitigate the impact of contrast fluctuations on our ability to detect a planet. Real-time drift correction and a posterior wavefront calibration are very complementary but historically have been developed separately: in this paper, we show how they can be combined. In particular, we demonstrate the entire dark zone maintenance (DZM)–planet detection pipeline for a variety of “observing scenarios” and post-processing techniques that will enable direct imaging of exo-Earths by future space-based-high-contrast instruments such as HWO. Note that the contrast achieved by DZM is currently limited by the in-air nature of the testbeds used to perform experiments. In general, if a testbed can generate a dark zone of a certain contrast, DZM can maintain it. Experiments on the Decadal Survey Testbed at the Jet Propulsion Laboratory are scheduled for late 2024 to demonstrate DZM at HWO-like contrast levels.



**Fig. 1** Sample high-contrast image (a) with an injected planet located in the cyan circle. The planet is not visible to the eye but can be extracted using post-processing (ADI in this example, see Sec. 2.5.2) as shown in the right panel. Note that, for this paper, we use the plasma colormap for images in units of log contrast (a) and the PuOr colormap for post-processed images in units of contrast (b).

## 1.2 Planet Signal Extraction from Direct Imaging Data

Even with a coronagraph and DMs, direct imaging of exoplanets benefits from, and often requires, data post-processing. The goal when post-processing direct imaging data is to find and remove the residual starlight in the dark zone. This does two things: (1) it removes bright speckles that could be falsely identified as planets and (2) it allows for imaging of planets below the raw dark zone contrast. This is shown by the sample image shown in the left panel of Fig. 1. A planet is located in the cyan circle, but the dark zone is still speckle-limited; hence, the planet is invisible and undetected. In this case, the planet contrast is  $4 \times 10^{-8}$ , and the mean contrast in the dark zone is  $7.8 \times 10^{-8}$ . We are, however, able to extract this planet using post-processing, as shown in Fig. 1(b). For this example, angular differential imaging (ADI) (discussed in detail in Sec. 2.5.2) is used as the post-processing technique.

In the absence of polarization optics, there are four main techniques to extract an exoplanet from direct imaging data:

1. ADI<sup>10-12</sup>
2. reference-star differential imaging (RDI)<sup>13,14</sup>
3. coherent differential imaging (CDI)<sup>15,16</sup>
4. spectral differential imaging (SDI).<sup>17,18</sup>

ADI is the main technique planned to be implemented on RST, with the telescope being rolled about the boresight every  $N$  hours while pointing at a target star (it is also used in conjunction with RDI on JWST<sup>6</sup>). Note that rolling the telescope induces a change in the thermal environment, and the angle is limited by the instrument sunshield; for RST, the roll angle is planned to be 26 deg,<sup>19</sup> and for JWST, the limit is nominally 10.4 deg.<sup>20</sup> This roll moves the planet in the dark zone during an observation sequence whereas, in theory, the residual starlight speckle pattern should stay the same as it depends on the optics. The movement of the planet with respect to the residual starlight speckles allows for the extraction of the planet signal. RDI uses a reference library of the speckle patterns in the dark zone from reference star data that does not contain a planet. It is the main method employed for Hubble Space Telescope (HST) high-contrast imaging<sup>21</sup> and is also used for JWST.<sup>6</sup> RDI does not require rolling the telescope, and often instruments point at a bright reference star to generate the dark zone before slewing to the target star and thus have a library of reference images already. CDI is a post-processing technique that generally requires additional hardware components, such as a self-coherent camera and a pupil-plane mask with a pinhole. It leverages the fact that the planet light is incoherent with the star light, whereas the residual starlight speckles are coherent with each other. This means the starlight speckles can be modulated by the DMs, but only the wings of the planet PSF are modulated and the core of the planet PSF is largely unaffected.<sup>15,16</sup> The CDI concept inspired one of the

approaches that is discussed later in Sec. 2.4.2. SDI is mainly used in conjunction with an integral field spectrometer, which provides the spectra for Nyquist-sampled points on the focal plane. Planets can be differentiated from star-induced speckles based on the wavelength-dependent behavior of the spectra. In the absence of astrophysical priors on the planet spectrum, SDI unfortunately does not work well for Earth-like planets as they are dim, small, and too close to the main lobe of the stellar PSF. Indeed, speckles near the main stellar PSF move sub-pixel distances as the wavelength varies, making it very difficult to distinguish them from planets. Planet detectability can be greatly enhanced with priors about the atmospheric composition in the case of self-luminous objects,<sup>22</sup> and ongoing work is quantifying realistic gains for reflected-light planets.<sup>23</sup> SDI will not be used in this paper as all experiments are conducted with monochromatic light. We refer the reader to the recent review paper by Follette et al.<sup>24</sup> for a more detailed description of post-processing techniques. In this paper, we demonstrate the use of DZM data with ADI as well as a software-based version of CDI and two other novel methods. Though only a few select post-processing methods are demonstrated with DZM data in this paper, there is nothing about DZM datasets that inhibits the use other post-processing methods such as SDI and RDI.

### 1.3 Paper Overview

This paper demonstrates DZM<sup>4,25</sup> in the presence of both bright and dim planet signals, as well as how that data behave when paired with different post-processing algorithms. So far, all experimental demonstrations of this technique were carried out without planet photons in the DZ. Here, we demonstrate that the presence of a bright astrophysical signal in the DZ does not preclude exquisite contrast maintenance. For our demonstrations, high-order wavefront error drifts are artificially injected into the system and corrected for using a single image from the science detector. The pixel sampling and photon rate of the resulting images are adjusted to be similar to future space missions in the same manner as discussed in Redmond et al.<sup>26</sup> and summarized in Sec. 6.2. The experiments are performed on the High-contrast Imager for Complex Aperture Telescopes (HiCAT)<sup>27–29</sup> at the Space Telescope Science Institute. The HiCAT testbed currently contains an apodized-pupil Lyot coronagraph, an IrisAO PTT111L segmented aperture DM, and a pair of Boston Micromachine (BMC) kilo DMs. The light source is an NKT Photonics SuperK EVO4 with a tunable filter; the tunable filter allows for the central wavelength and bandpass to be set. For this paper, a wavelength of 638 nm and a bandwidth of 1.6% are used. The simulator used to emulate the testbed, catkit2,<sup>30</sup> is based on its predecessor (catkit), which has been discussed in Noss et al. 2021,<sup>31</sup> Fowler et al.,<sup>32</sup> and Moriarty et al.<sup>33</sup> Drifts are injected using the two BMC DMs for the DZM experiments, and the planet signal is added numerically to the testbed images prior to handing them to the estimator. This work is a continuation of Pogorelyuk et al.<sup>25,34</sup> and Redmond et al.<sup>4,26</sup> As previously mentioned, DZM experiments are performed with planet contrasts on par with, above, and below the mean DZ contrast. An extension of the DZM estimator is presented to ensure the stability of the control loop in the presence of bright incoherent light sources. Four post-processing techniques are considered: ADI, incoherent accumulated imaging (IAI), CDI, and a novel approach that combines RDI and CDI that we call coherent reference differential imaging (CoRDI). RDI, CDI, and IAI all work with the new DZM extension.

Section 2 covers the new features implemented with the DZM algorithm as well as how the planet signal is injected and handled in the software. It also describes how each of the post-processing techniques is implemented and how the signal-to-noise ratio (SNR) of the planet is determined in this paper. Section 3 summarizes a subset of the experiments performed and demonstrations of each of the post-processing techniques covered in Sec. 2. Section 4 compares the performance of the post-processing techniques via their ability to detect a planet, and Sec. 5 summarizes the results of this paper and outlines future work. Content from this paper is also included in Redmond et al.<sup>35</sup>

## 2 DZM and Post-Processing Pipeline Overview

### 2.1 Dark Zone Maintenance

A detailed overview of DZM is provided in Redmond et al.,<sup>4</sup> and a brief summary is provided here for context. The first step before commencing DZM is to generate the dark zone; this is done

using traditional approaches such as pair-wise probing<sup>36</sup> (PWP) to estimate the electric field and electric field conjugation<sup>37</sup> (EFC) to determine a DM command to suppress the electric field. Once the dark zone is generated (i.e., we hit the contrast limit of the testbed), the PWP-EFC algorithm is stopped, and we begin injecting drifts into the system using the two BMC DMs. For this paper, each actuator on the DMs performs a random walk with a standard deviation of  $\sigma_{\text{drift}}$ . The DZM algorithm is activated when the DMs begin injecting drifts. The open-loop electric field is then estimated using a single science image. It is passed to an extended Kalman filter (EKF), which requires input parameters for the process noise, determined by an estimate of the drift statistics, and the measurement noise, set by the detector properties. Once an estimate of the electric field is obtained by the EKF, it is passed to the EFC controller, which determines the DM command that will suppress the injected drift in the open-loop electric field. On top of the EFC DM command, a small random command is applied to the DMs, which we refer to as the “dither.” This command is described by the standard deviation  $\sigma_{\text{dither}}$ . The purpose of the dither command is to help the EKF converge as it introduces phase diversity from one image to the next.<sup>25,38</sup> Definitions for the nominal DZM parameters are provided in Sec. 6.1.

The DZM algorithm can be run using the raw images from the testbed, which nominally have a high photon count, or in low-photon mode. The low-photon mode allows us to run experiments representative of a real space-based observing scenario in which there is a finite amount of flux on a much shorter timescale. The low-photon conversion procedure as well as results from HiCAT experiments using it are described in Redmond et al.,<sup>26</sup> and a brief description is provided in Sec. 6.2. In low-SNR mode, a high-fidelity image is taken using the testbed and then degraded to have the realistic photon rate and noise properties. The low-SNR image is then provided to the estimator, and the measurement noise parameter in the EKF is updated to represent the modeled detector. All experiments in this paper use the low-SNR mode with representative flux values for a real space-based observing scenario.

## 2.2 DZM Enhancement with Real-Time Incoherent Light Estimation

In this paper, we expand on the DZM algorithm capabilities to make it more robust and enable different post-processing methods. One potential limitation of the original DZM implementation is the estimation bias caused by bright planets or bright exozodiacal dust emission.<sup>39</sup> If these incoherent light sources are very bright, the estimator will capture it, and subsequently, the DMs will try to suppress it. However, as the light is incoherent, the DMs will not succeed in correcting for it, which can cause the control loop to diverge. This concept is demonstrated using a HiCAT simulation in Sec. 6.3. A way to ensure stability of the control loop in the case of bright planets, incoherent background, and exozodiacal dust emissions is to augment the original DZM EKF, described in Redmond et al.,<sup>4</sup> to have a third state, which is the incoherent light in the image. This augmented EKF (AEKF) is described in detail for digging a dark zone in Riggs;<sup>40</sup> we adapt it in this paper for maintaining the dark zone. The nominal DZM algorithm described in Redmond et al.<sup>4</sup> can be referred to for the nominal EKF setup; definitions are also provided in Sec. 6.1. The response matrix for the AEKF (equivalent to the Jacobian  $G$  for the nominal EKF) and state vector become

$$\Gamma = \sqrt{t_{\text{exp}}} \begin{bmatrix} \Re \left( \frac{\partial E_{OL_0}}{\partial u} \right) \\ \Im \left( \frac{\partial E_{OL_0}}{\partial u} \right) \\ 0 \\ \vdots \\ \Re \left( \frac{\partial E_{OL_{n-1}}}{\partial u} \right) \\ \Im \left( \frac{\partial E_{OL_{n-1}}}{\partial u} \right) \\ 0 \end{bmatrix}, \quad (2)$$

$$x = \begin{bmatrix} \Re(E_{OL_0}) \\ \Im(E_{OL_0}) \\ I_{\text{inco}_0} \\ \vdots \\ \Re(E_{OL_{n-1}}) \\ \Im(E_{OL_{n-1}}) \\ I_{\text{inco}_{n-1}} \end{bmatrix}, \quad (3)$$

where  $I_{\text{inco}}$  is the incoherent intensity in the image,  $E_{OL}$  is the open loop electric field, and  $n$  is the number of pixels in the dark zone. The closed-loop version of the state vector is expressed as

$$x_{CL}^k = x^k + \Gamma u^k, \quad (4)$$

and  $x_{CL}$  comprises

$$x_{CL} = \begin{bmatrix} \Re(E_{CL_0}) \\ \Im(E_{CL_0}) \\ I_{\text{inco}_0} \\ \vdots \\ \Re(E_{CL_{n-1}}) \\ \Im(E_{CL_{n-1}}) \\ I_{\text{inco}_{n-1}} \end{bmatrix}. \quad (5)$$

This leads to the following equation for the model of the measurement:

$$z^{k+1} = E_{CL_{\text{Re}}}^k \circ E_{CL_{\text{Re}}}^k + E_{CL_{\text{Im}}}^k \circ E_{CL_{\text{Im}}}^k + I_{\text{inco}} + n^k, \quad (6)$$

where  $E_{CL_{\text{Re}}}$  and  $E_{CL_{\text{Im}}}$  are the real and imaginary parts of the closed-loop electric field, respectively, as calculated in Eq. (4);  $\circ$  indicates the Hadamard product; and  $n^k$  is the observation noise. Assuming drift on the BMC DMs, the process noise is updated to be

$$Q_{\text{BMC}}^k = \Gamma \Gamma^T \sigma_{\text{drift}}^2. \quad (7)$$

Note that, as  $\Gamma$  is padded with zeros,  $Q$  nominally has a zero row associated with the incoherent part of the state vector for each pixel. The 0 rows of  $Q$  are thus set to a small number to allow the EKF to converge and be numerically stable. The last major change for the EKF is the observation matrix ( $H^k = \frac{\partial h}{\partial x^k}$ ). For each pixel, the  $H$  matrix is augmented to be

$$H_i^k = \begin{bmatrix} 2\Re(\hat{E}_{CL_i}^k) & 0 & 0 \\ 0 & 2\Im(\hat{E}_{CL_i}^k) & 0 \\ 0 & 0 & 1 \end{bmatrix}, \quad (8)$$

where  $i$  indicates the pixel index in the dark zone. Note that we still use EFC for the controller, and, as with the nominal DZM algorithm, it is only handed the open loop electric field estimate, not the augmented state vector.

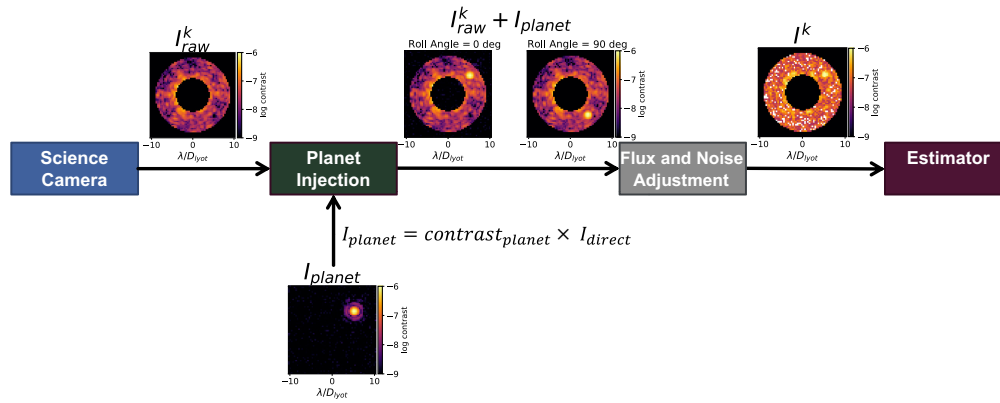
A very important aspect of any Kalman filter is how the state vector is initialized. For the original EKF, the open-loop electric field estimate can be initialized either as zero, where it will converge in  $\sim 30$  iterations, or as the final estimate of the initial dark zone algorithm that generated the dark zone. For the incoherent light, three options were considered:

$$\hat{I}_{\text{inco}}^0 = \overline{|z^0 - \hat{z}^0|}, \quad (9)$$

$$\hat{I}_{\text{inco}}^0 = |z^0 - \hat{z}^0|, \quad (10)$$

$$\hat{I}_{\text{inco}}^0 = |z^0 - z_{\text{ref}}|, \quad (11)$$

where  $z^0$  is the initial science image on the target,  $\hat{z}^0$  is the initial estimate of the closed-loop intensity (generated from the reference star), and  $z_{\text{ref}}$  is the final image from generating the dark



**Fig. 2** Flowchart for injecting a fake planet into the data acquired from the testbed during a DZM run. If operating in ADI mode, the planet can be injected at either roll angle. The low-SNR conversion occurs after the planet has been injected into the high-SNR image.

zone on the reference star. Each of these options is physical in its own way. Equation (9) assumes that the primary contribution to the incoherent light is a uniform background that should be the same order of magnitude as the error in the intensity estimate of the coherent light and the image. Equation (10) assumes that the incoherent light has a structure that matches the estimate error. Also note that, as  $\hat{z}^0$  is from the reference star, it will not contain the planet but  $z^0$  will; taking the difference and keeping the structure potentially aid the EKF in locating the planet. The last option [Eq. (11)] is based on a similar idea in that we can take the difference between the final reference star image and the first target star image to start the estimator with a potential planet location. Based on preliminary results, all approaches work well, and Eq. (9) is nominally used throughout this paper. In future work, the other methods may become more attractive as we strive to improve the tuning of the  $Q$  matrix to reduce the small amount of coherent light leaking into the incoherent estimate (further discussed in Sec. 3.2).

### 2.3 Planet Injection During Laboratory Experiments

The HiCAT testbed only has a laser source to mimic the target star around which an exoplanet could be orbiting. To test post-processing techniques to determine how DZM affects our ability to detect a planet, we inject an artificial planet signal in software. We do this while running the experiment to ensure that the effect of the planet signal on the control loop is captured. To generate a fake planet, we take the direct image and scale its intensity down so that its peak matches the desired planet contrast. We then shift the center of this intensity-scaled PSF into the dark zone and add the result to the image obtained from the camera. If we are running an ADI experiment, the roll angle of the planet varies throughout the experiment; this is further discussed in Sec. 2.5.2. The process of obtaining the science images and injecting the planet is shown in Fig. 2. If we are operating in low-SNR mode (Sec. 6.2), the planet injection happens prior to the low-SNR conversion, as described in Fig. 1 of Redmond et al.<sup>26</sup>

### 2.4 Co-add Post-Processing Methods

As outlined in Sec. 1.2, there are a number of different post-processing techniques that have been developed to extract an exoplanet from direct imaging data. Here, we discuss how to implement a set of post-processing techniques with DZM datasets in particular.

#### 2.4.1 Incoherent accumulated imaging

The IAI method requires the AEKF version of DZM. Here, we take the mean of the incoherent estimates because the incoherent estimate has captured the planet:

$$F_{IAI} = \frac{1}{n_{\text{iter}}} \sum_{k=1}^{n_{\text{iter}}} \hat{I}_{\text{inco}}^k. \quad (12)$$

where  $F_{\text{IAI}}$  is the post-processing image for the incoherent estimate approach. This assumes that the DZM residuals are stationary and ergodic, resulting in an improvement of the estimate of the incoherent intensity as more frames are added.

### 2.4.2 Coherent differential imaging

Here, we implement a simplified software-based version of CDI, which leverages the closed-loop coherent intensity estimate produced by the AEKF. We take the closed-loop coherent intensity estimate for each iteration ( $|\hat{E}_{\text{CL}}^k|^2$ ) and subtract it from the science image ( $z^k$ ):

$$I_{pf}^k = z^k - |\hat{E}_{\text{CL}}^k|^2, \quad (13)$$

$$F_{\text{CDI}} = \frac{1}{n_{\text{iter}}} \sum_{k=1}^{n_{\text{iter}}} I_{pf}^k, \quad (14)$$

where  $I_{pf}^k$  is the “post-processing” image for a single iteration. As the planet is not in the coherent electric field estimate,  $I_{pf}^k$  should reveal the planet in the science image. When all of the estimate differences are obtained, we obtain the final post-processed image  $F_{\text{CDI}}$  by taking an average over all iterations to reduce the noise [Eq. (14)]. As with the IAI approach, this assumes that the DZM residuals are stationary and ergodic.

## 2.5 Karhunen–Loève Image Projection Post-Processing Methods

### 2.5.1 Coherent reference differential imaging

Here, we introduce the CoRDI method, which is based on using a combination of CDI and Karhunen–Loève Image Projection<sup>41,42</sup> (KLIP) in RDI mode. The basis of the KLIP algorithm is to perform a principal component analysis (PCA) of a reference image library and project the PCA results onto the set of science images to extract the planet signal. KLIP can be used with most differential imaging techniques (described in Sec. 1.2) and is implemented using the python package `pyklip`;<sup>43</sup> `pyklip` is an open-source implementation of the PCA approach described in Soummer et al.<sup>41</sup> For CoRDI, we use the RDI PSF library philosophy, except that reference images are composed of only coherent estimates and thus, in theory, are devoid of the planetary signal. The optimal synthetic PSF based on these coherent estimates is then calculated using a PCA. To implement CoRDI with the HiCAT DZM data, we leverage the closed-loop coherent intensity estimate produced by the AEKF. The idea is to use the closed-loop coherent intensity estimates as the “reference images” as they do not contain the planet and should capture the PSF structure. The closed-loop intensity estimate is calculated via

$$\hat{I}_{\text{CL}} = |\hat{E}_{\text{CL}}|^2, \quad (15)$$

where  $\hat{E}_{\text{CL}}$  is drawn from the appropriate rows of  $\hat{x}_{\text{CL}}$ , as shown in Eqs. (4) and (5). First, we generate a PSF library that contains the closed-loop coherent intensity estimates and the science images. Using this PSF library, `pyklip` generates a correlation matrix for the entire dataset. Then, the science dataset, which only contains the science images, is assembled. Both the PSF library and the science dataset are passed to `pyklip`, which performs the following PCA procedure:

1. perform PCA of all reference images and keep the first  $p$  modes
2. project  $p$  PCA modes onto each science image and sum the  $p$  projections for each image
3. subtract each science image from the sum of the projections and sum the result for each image.

Note that the images are indexed and `pyklip` automatically removes any science images from the PSF library when it assembles the reference PSF dataset. Because a coherent PSF estimate ( $\hat{I}_{\text{CL}}$ ) exists for each science frame, the PSF library contains information on how the PSF evolves over time and changes for each of the dither commands. A reference PSF dataset is generated in parallel of the science observations for each target star, which means that the reference PSF dataset contains any target-star-specific characteristics such as stellar diameter and temperature. We generate results for  $p = [1, 5, 10, 20, \text{and } 50]$ .



This process can be largely explained by Eqs. (5) and (8) in Soummer et al.<sup>41</sup> Here, we adjust those equations to match the nomenclature in this paper. We start by defining the reference images as

$$\{R_k(i) = \hat{I}_{\text{CL}_k}(i)\}_{k=1 \dots n_{\text{iter}}}, \quad (16)$$

where  $k$  indicates an iteration in the experiment and  $i$  indicates a pixel in the image. From the reference images, we generate a covariance matrix of the references over the search area  $\mathcal{S}$ :

$$E_{\text{RR}}[l, j] = \sum_{i=1}^{N_{\mathcal{S}}} R_l(i)R_j(i), \quad (17)$$

where  $l, j \in 1 \dots n_{\text{iter}}$ . From the covariance matrix, we take the eigenvectors ( $c_p$ ) and eigenvalues ( $\Lambda_p$ ) to calculate the Karhunen–Loève transform ( $Z$ ) of the set of reference PSFs:

$$Z_p^{\text{KL}}(i) = \frac{1}{\sqrt{\Lambda_p}} \sum_{k=1}^{n_{\text{iter}}} c_p(k)R_k(i), \quad (18)$$

where  $p$  indicates the KLIP mode,  $k$  indicates the iteration, and  $R_k$  is the reference image for iteration  $k$ . We then generate the optimal synthetic PSF by projecting a science target image  $z_k$  (previously  $z^k$ ) onto the Karhunen–Loève transform of the reference PSFs:

$$\hat{I}_{\psi_k}(i) = \sum_{p=1}^{K_{\text{klip}}} \langle z_k, Z_p^{\text{KL}} \rangle_{\mathcal{S}} Z_p^{\text{KL}}(i), \quad (19)$$

and the final post-processed image  $F_{\text{CoRDI}}$  is obtained via

$$F_{\text{CoRDI}}(i) = \sum_{k=1}^{n_{\text{iter}}} z_k(i) - \hat{I}_{\psi_k}(i). \quad (20)$$

## 2.5.2 Angular differential imaging

ADI is a well-developed post-processing technique as described in Sec. 1.2 for directly imaging exoplanets, and here we pair it with DZM. ADI rotates the telescope around the bore sight to move the planet on the detector. The speckles induced by the optics should rotate with the telescope, which allows for extraction of the planet signal. On HiCAT, instead of rotating the optical bench, we move where we inject the planet as shown in Fig. 2 using a roll angle of 90 deg. When running ADI experiments, we also cycle the dither commands used by the EKF [Eq. (21) in Redmond et al.<sup>41</sup>]. We do this so that the speckles introduced by the dither commands are more consistent, which aides with any type of differential imaging post-processing algorithm. For the ADI experiments, the roll angle is changed either halfway through, so  $\delta u_{\text{dither}}^k = \delta u_{\text{dither}}^{k+n_{\text{iter}}/2}$ , or every  $n_{\text{iter}}/4$  iterations.

To post-process the ADI data, we again use KLIP. When used with ADI, the main steps consist of the following:

1. perform PCA of all images from roll 1 and keep the first  $p$  modes
2. project  $p$  PCA modes onto each image from roll 2 and sum the  $p$  projections for each image
3. subtract each roll 2 image from the sum of the projections and sum the result for each image
4. repeat steps 1 to 3 but switch the roles of rolls 1 and 2
5. rotate the roll 2 result to align with roll 1 and add.

Note that, for this analysis, we generate results for  $p = [1, 5, 10, 20, \text{ and } 50]$ .

Using KLIP with ADI data is explained in detail in Sec. 6 in Pueyo et al.<sup>42</sup> Here, we adjust those equations to match the nomenclature in this paper. The procedure is very similar to the RDI mode of pyklip with a slight change to how the reference library is built. We split the dataset into

two such that  $z^+$  contains all images for  $+\theta_{\text{roll}}/2$  and  $z^-$  contains all images for  $-\theta_{\text{roll}}/2$ . We build two reference libraries as

$$\begin{cases} R_k^+(i) = z_k^+(i) & \text{for } k = 1 \dots n_{\text{iter}}/2 \\ R_k^-(i) = z_k^-(i) & \text{for } k = 1 \dots n_{\text{iter}}/2 \end{cases} \quad (21)$$

Starting with the minus roll case, we use  $z^-$  as our science dataset and  $R^+$  as the reference dataset. We take the Karhuene–Loève transform of the “reference” set  $R^+$ :

$$Z_p^{\text{KL}^+}(i) = \frac{1}{\sqrt{\Lambda_p^+}} \sum_{k=1}^{n_{\text{iter}}/2} c_p^+(k) R_k^+(i), \quad (22)$$

where the covariance matrix is assembled using Eq. (17). The optimal synthetic PSF follows as

$$\hat{I}_{\psi_k}^+(i) = \sum_{p=1}^{K_{\text{klip}}} \langle z_k^-, Z_p^{\text{KL}^+} \rangle_{\mathcal{S}} Z_p^{\text{KL}^+}(i), \quad (23)$$

from which we get the post-processed image for the minus roll case

$$F^-(i) = \sum_{k=1}^{n_{\text{iter}}/2} z_k^-(i) - \hat{I}_{\psi_k}^+(i). \quad (24)$$

We repeat Eqs. (22) and (23) for the  $+\theta_{\text{roll}}/2$  case to get the other post-processed image

$$F^+(i) = \sum_{k=1}^{n_{\text{iter}}/2} z_k^+(i) - \hat{I}_{\psi_k}^-(i), \quad (25)$$

at which point we rotate  $F^+$  and  $F^-$  by  $\theta_{\text{roll}}/2$  and co-add them to obtain the final post-processed image  $F_{\text{ADI}}$ .

## 2.6 SNR Calculations

To calculate the SNR of the post-processed images, we use the “planet detection” procedure developed by pyklip. First, the post-processed image is loaded and cross-correlated with the planet PSF, which we represent as a 2D Gaussian kernel; we refer to the result as the cross-correlated image. Using the cross-correlated image, we can generate an SNR map using the `get_image_stat_map_perPixMasking` function. This function goes pixel-by-pixel in the image that has been cross-correlated with the PSF and takes the signal to be the value of the current pixel. To calculate the noise, it masks a region around the current pixel and computes the standard deviation of an annulus outside of the masked region. For the purposes of this paper, the mask region has a radius of five pixels, and the annulus has a width of two pixels. Note that we mask irrelevant pixels in the cross-correlated image prior to handing it to the SNR map generator. All pixels outside of the dark zone, as well as inner and outer radii of the dark zone where leakage is most likely to occur, are masked; the resulting zone spans from  $4.2 - 8.8\lambda/D_{\text{Lyot}}$ .

## 3 Experimental Results

Here, we present the results from experiments run on the HiCAT testbed. All experiments are monochromatic at 640 nm and use a  $4 - 9\lambda/D_{\text{Lyot}}$  annular dark zone. The parameters common to all experiments are provided in Table 1. For both the ADI and AEKF DZM configurations, experiments are run with a variety of planet contrasts ranging from  $0.5 - 12.5\times$  the dark zone contrast. Note that the PCA-based post-processing methods (Sec. 2.5) remove the mean speckle pattern but not the standard deviation over the course of the experiment. The lower limit on the planet contrast is chosen such that it is larger than the standard deviation of the mean dark zone contrast to ensure that we are operating in a range at which the PCA-based methods are expected to work. The experiments in which the single-frame-planet SNR is low highlight the post-processing gains, and the experiments in which the single-frame-planet SNR is high demonstrate the DZM capabilities. Providing results for this range of planet contrasts shows the flexibility and

**Table 1** Experiment parameters used. The low-photon method and parameters are described in Redmond et al.<sup>26</sup>

Parameter	Value	Units
$\sigma_{\text{dither}}$	250	Pm
$\sigma_{\text{drift}}^2$	400	$\text{pm}^2/\text{iter}$
$t_{\text{HiCAT}}$	0.1	S
$t$	39	S
$F$	$1.5 \times 10^7$	photons/s
$\mu_I$	0.005	e-/s
$\sigma_{rn}$	2	e-
$\theta_{\text{planet}}$	6.62	$\lambda/D_{\text{Lyot}}$

robustness of both the DZM algorithm and the post-processing techniques. Note that only a subset of the experiments performed are included; these cases were chosen to be the most extreme ends of the range at which the post-processing results were promising. These results should not be used to place hard limits on predicted post-processing gains.

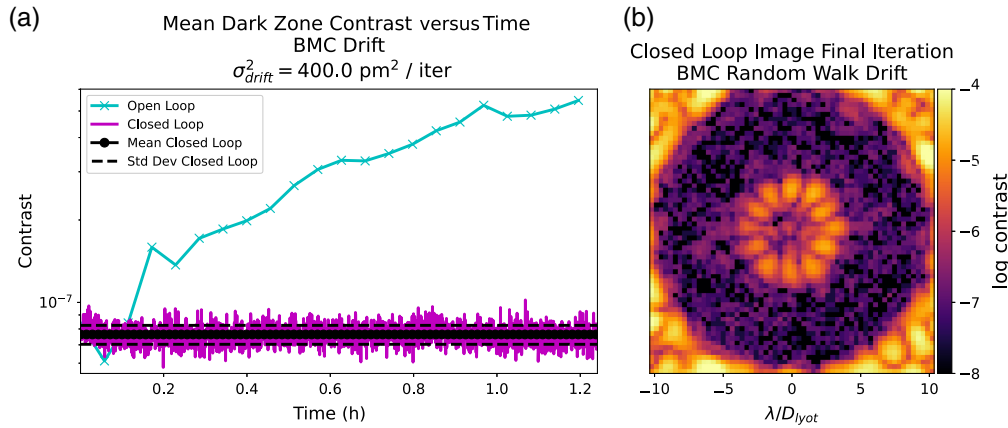
The number of iterations chosen for each experiment is tied to the Observing Scenario 11 (OS11) for RST. Krist<sup>19</sup> outlines the plan for target 47 UMa in which the instrument spends a total integration time of 26.7 h on the target. Every 6.7 h (400 min), RST slews to the reference star to touch up the dark zone. The exposure time used for a single image in the OS11 dataset is 39 s, which indicates  $\sim 2461$  total exposures on the target and  $\sim 615$  exposures in between slews to the reference star. In the interest of time, full 2400 iteration experiments were not run. Instead, for the ADI experiments, we use 2200 iterations, and for the AEKF experiments, we use 1800 iterations. It was determined that  $\geq 1800$  iterations are sufficient to demonstrate that DZM can remove the requirement of slewing to a reference star and that the post-processing methods tested can work with a representatively sized dataset. Note that DZM can maintain the dark zone for longer, as shown by the 4000 iteration experiment in Redmond et al.,<sup>26</sup> but for this paper, the priority was to produce a number of datasets with a range of input values; thus to be able to acquire the data in a reasonable amount of time, the number of iterations was reduced.

### 3.1 Angular Differential Imaging

For this experiment, we inject a  $4 \times 10^{-8}$  planet into a dark zone with a mean DZ contrast of  $7.8 \times 10^{-8}$ . The BMC DMs perform a random walk drift over 2200 iterations. The parameters for this experiment are summarized in Table 2, and the low-photon procedure is described in Redmond et al.<sup>26</sup> To match RST, the exposure time simulated for the low-photon mode (Sec. 6.2) is 39 s, which results in a total simulated integration time of 23.8 h. For this experiment, we “roll” the telescope four times (every 550 iterations) to match the RST Observing Scenario 11 procedure.<sup>19</sup> Note that the magnitude of the roll angle (90 deg) is 3.5 times larger than that planned for RST. This large roll angle was chosen as these are proof-of-concept

**Table 2** ADI DZM experiment parameters used.

Parameter	Value	Units
$\text{contrast}_{\text{planet}}$	$4 \times 10^{-8}$	—
$\theta_{\text{roll}}$	90	deg
$n_{\text{rolls}}$	4	—
$n_{\text{iter}}$	2200	—

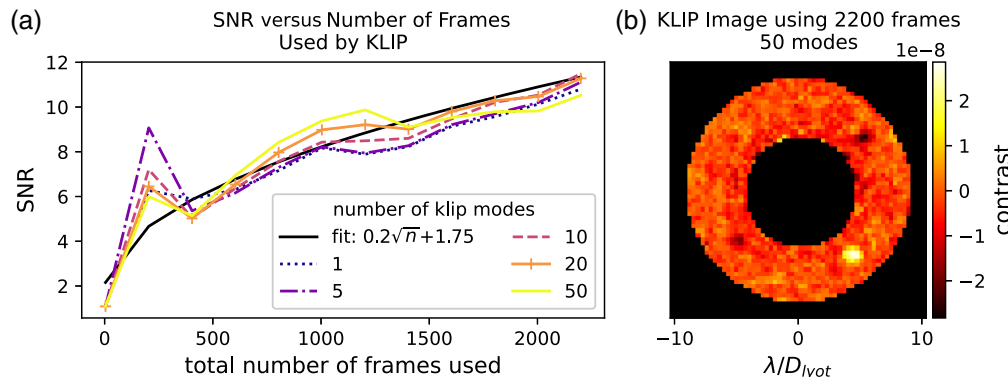


**Fig. 3** Low-photon ADI DZM experiment results. (a) Mean DZ contrast versus time for the low-photon case with a  $4 \times 10^{-8}$  planet. The cyan is the open-loop contrast, the magenta is the closed-loop contrast, and the black dots and dashes are the mean and standard deviation of the magenta curve, respectively. Note that the time axis is the actual duration of the experiment on the testbed and does not use the 39 s of simulated exposure time that RST would require. (b) Sample science image from the final iteration. The planet is not visible to the eye in the dark zone.

experiments, but it does mean that the ADI results are optimistic with respect to a real observing scenario.

The experiment summary plot is provided in Fig. 3 with the mean dark zone contrast versus time in the left panel and the closed loop image in the right panel. The time axis for this plot is HiCAT time. As shown in the right panel of Fig. 3, the planet is not visible, even when using a log scale for the image.

As shown in Fig. 4(a), when all 2200 frames are used, the planet achieves an SNR of 10.52. The behavior of all modes is similar, so we know that we are not fitting the noise. Also, as shown by the black line, the SNR increases with  $\sqrt{n}$ , which is expected from Nemati et al.<sup>44</sup> An SNR of 10.52 is well above the commonly used detection threshold of 5, and one can also clearly see the two negative lobes created by the pyklip algorithm in the right panel of Fig. 4. The low signal lobes at  $\pm\theta_{\text{roll}}/2$  on either side of the planet are a result of taking the autocorrelation of the planet signal and is described in detail in Cantalloube et al.<sup>45</sup> This experiment demonstrates how DZM can be used in conjunction with ADI in the case of dim planets below the contrast floor.



**Fig. 4** Post-processed KLIP results for an ADI DZM experiment. (a) SNR as a function of number of frames provided to KLIP. This can be interpreted as SNR versus integration time. The SNR trends up as the number of frames increase, and all modes have similar behavior. For 50 modes, the planet has an SNR of 10.52 when all 2200 frames are used. The black line shows a fit as the SNR should go as  $\sqrt{n}$ . (b) Image produced by pyklip for the 50-mode case when all 2200 frames are used. The planet, as well as the low noise lobes on either side of it, is clearly visible in the lower right quadrant.

## 3.2 Augmented DZM Experiments

### 3.2.1 Bright planet

We first investigate the bright planet case with the AEKF to easily visualize the extraction of the planet light by the AEKF and to show that the augmented DZM algorithm is stable in the presence of a bright incoherent source. For this experiment, we inject a  $10^{-6}$  planet into a dark zone with a mean contrast of  $9.8 \times 10^{-8}$ . The images passed to the estimator are flux-adjusted as per Table 1. The experiment runs for 1500 iterations, all of which are used in the post-processing. The contrast versus time plot is provided in Fig. 12 in Sec. 6.4.

The results of the CDI, IAI, and CoRDI post-processing techniques are shown in Fig. 5. The left column shows the post-processed image using four frames, the middle panel shows the SNR evolution as more frames are provided, and the right panel shows the final post-processed image. Note that the SNR calculations use the pyklip methodology described in Sec. 2.6, and the post-processed images are in units of contrast. For all cases, the post-processed image using four frames is quite granular, but the planet is still clearly visible. The post-processed image in the right column using all 1500 frames is very smooth, and the SNR has increased. An interesting thing to note is that IAI/CDI and CoRDI seem to suffer from different systematics. IAI and CDI have a residual sixfold spoke pattern with brighter speckles at  $\sim 6\lambda/D_{\text{Lyot}}$  (further discussed in Sec. 3.2.2), whereas CoRDI has negative speckles near the IWA.

A more quantitative method of comparing the results is to look at the SNR evolution provided in the middle column of Fig. 5. For each row, the dotted blue line is associated with the post-processed images on either side, and the solid black line is a fit of the data with the form  $A\sqrt{n} + B$ , where  $n$  is the number of frames used. Based on Nemati et al.,<sup>44</sup> we know that, in the ideal case, the SNR should be proportional to the square root of the integration time. Determining the appropriate scaling factor is outside the scope of this paper, but we can see that, for all approaches,  $A\sqrt{n} + B$  provides a good fit to the data. As this is a bright planet, we start with an SNR of  $\sim 40$  in all curves. After 1500 frames, IAI and CDI perform slightly better with a final SNR of  $\sim 68$  versus 58.6 for CoRDI.

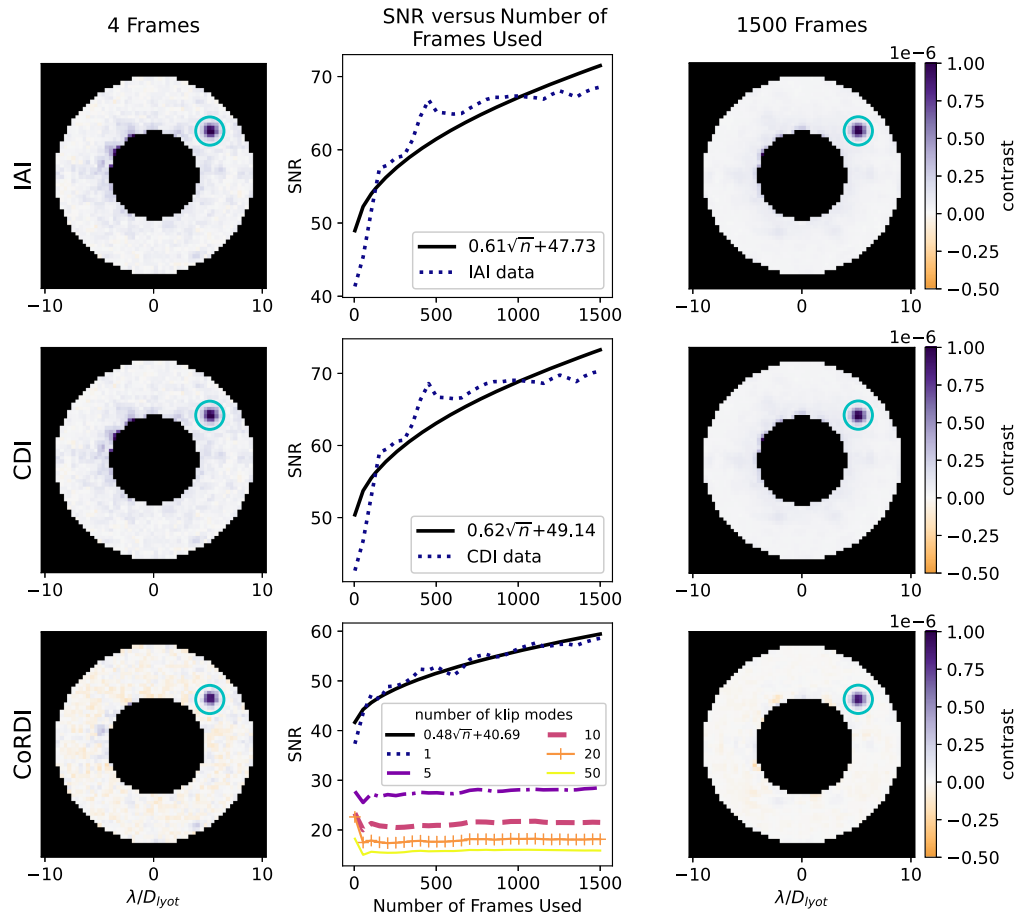
It should be noted that the top two rows in Fig. 5 look very similar. The residual sixfold speckles are erroneously captured by the incoherent intensity estimate instead of the coherent electric field estimate, which is why the top two rows look the same. This may be a fundamental limit of the AEKF as the incoherent estimate will tend to capture any light that is not being modulated by the DMs. If there are other sources of incoherent light in the dark zone, they will be captured by  $\hat{I}_{\text{inco}}$ . If other incoherent light sources form a uniform background, it is much less of an issue than if it has structure.

Another interesting note is that the CoRDI results are significantly impacted by how many KLIP modes are used when a bright planet is present, as shown in the bottom middle panel of Fig. 5. In Fig. 6, we show the final post-processed CoRDI image for all KLIP modes calculated. If one KLIP mode is used, the planet has an SNR of 58.6, but if 50 modes are used, it has an SNR of 15.8. As the number of modes increases, KLIP begins to fit the residual structure as shown in the increase in speckles between the left and right panel of Fig. 6. This KLIP effect is not present when using the ADI mode with the DZM data (Fig. 4) or when using CoRDI with a dim planet (further discussed in Sec. 3.2.2).


### 3.2.2 Dim planet

For this experiment, we inject an  $8 \times 10^{-8}$  planet into a dark zone with a mean contrast of  $8.9 \times 10^{-8}$ ; note that here the planet cannot be detected in an individual frame and extraction of the planet signal will rely on post-processing techniques. The images passed to the estimator are flux-adjusted as per Table 1. The experiment runs for 1500 iterations, all of which are used in the post-processing. The contrast versus time plot is provided in Fig. 13 in Sec. 6.4.

The results of the CDI, IAI, and CoRDI post-processing techniques are shown in Fig. 7. The left column shows the post-processed image using four frames, the middle panel shows the SNR evolution as more frames are provided, and the right panel shows the final post-processed image. Note that the SNR calculations use the histogram overlap approach discussed in Sec. 7, and

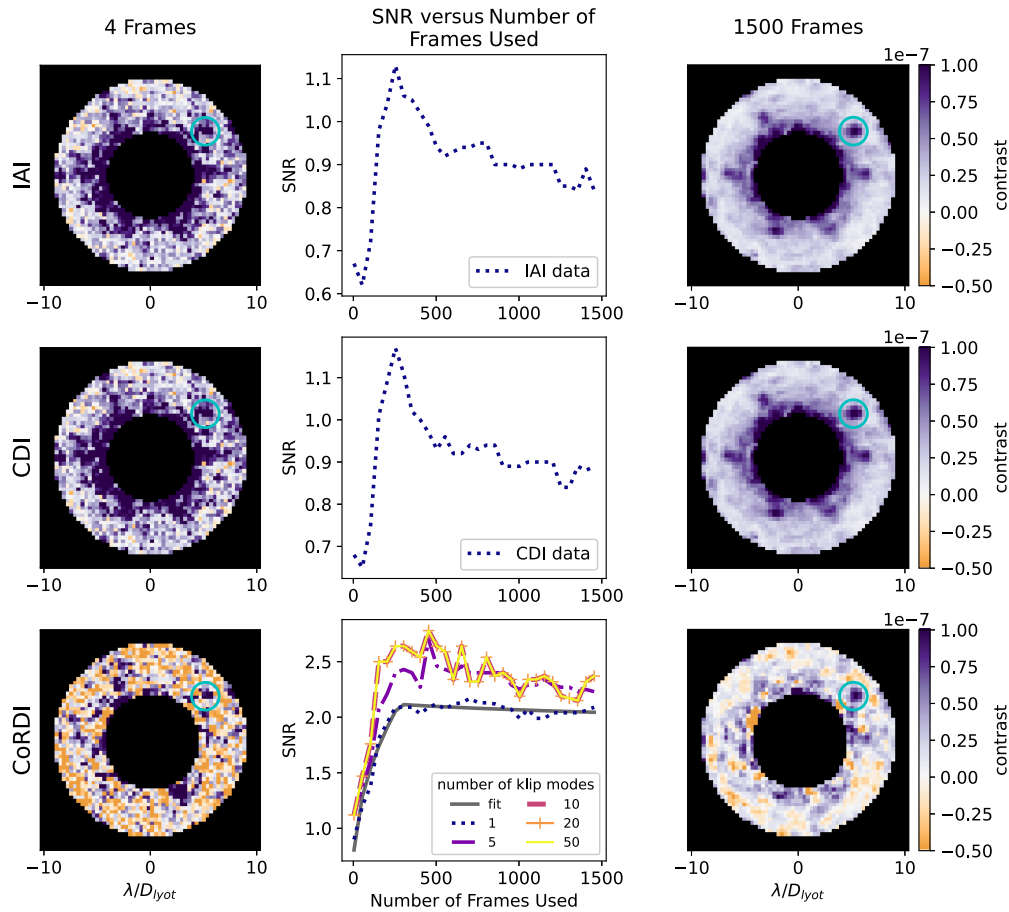


**Fig. 5** AEKF DZM post-processing results for a  $10^{-6}$  planet injected in the dark zone. The left and right columns show the post-processed images (in units of contrast) for the IAI, CDI, and CoRDI approaches with the planet circled in cyan. The left column only provides the first four frames to the associated algorithm, and the right column uses all 1500 frames. The middle column shows the SNR evolution as more frames are provided to each algorithm. The dotted blue curve is associated with the images in the left and right columns, and the solid black curve is a fit of the dotted blue curve of the form  $A\sqrt{n} + B$ . For the CoRDI results (bottom row), the SNR curve is provided for the five KLIP mode options calculated.



**Fig. 6** KLIP CoRDI result using 1500 frames for a  $10^{-6}$  planet in a low-photon dark zone; the AEKF version of DZM is used. The planet is clearly visible to the eye for all KLIP modes. The SNR decreases as the number of KLIP modes increases because KLIP begins to fit the residual structure in the image, which increases the noise.

the post-processed images are in units of contrast. The right column of Fig. 7 highlights why the alternate SNR calculation method is required. The colormap is centered such that purple is positive, white is zero, and orange is negative. For IAI and CDI, the post-processed images are clearly not zero-mean, and the residual sixfold spoke pattern creates a long tail in the histogram.




**Fig. 7** AEKF DZM post-processing results for a  $8 \times 10^{-8}$  planet injected in the dark zone. The left and right columns show the post-processed images (in units of contrast) for the IAI, CDI, and CoRDI approaches with the planet circled in cyan. The left column only provides the first four frames to the associated algorithm, and the right column uses all 1500 frames. The middle column shows the SNR evolution as more frames are provided to each algorithm. The dotted blue curve is associated with the images in the left and right columns, and the solid gray curve is a fit of the dotted blue curve of the form  $A\sqrt{n} + B$ . For the CoRDI results (bottom row), the SNR curve is provided for the five KLIP mode options calculated.

For the CoRDI approach, the noise is much closer to being zero-mean and is much less structured. Using the AEKF with DZM is still a very novel concept, and even showing that the incoherent state in the AEKF captures the dim planet (top row in Fig. 7) is a very promising result. As in Sec. 3.2.1, the more frames that are used in the stack, the more the background smooths out.

For the dim planet case, there are a significant number of residual speckles on the same order of magnitude as the planet in the IAI and CDI post-processed images. These residual speckles are due to the same issue discussed in Sec. 3.2.1, but in this case, they have a similar level of signal as the planet. One way to differentiate between a speckle and a planet for IAI or CDI would be through target revisits. By returning to a target multiple times, we image the planet at different locations along its orbit. The speckle pattern will either stay the same (if the optics are sufficiently stable) or move in a manner that is inconsistent with an orbit, which facilitates planet identification.

In the middle column of Fig. 7, we see the SNR evolution as more frames are provided to the post-processing algorithm. For all methods, the SNR increase is drastic from 0 to 300 frames but then either declines (IAI and CDI) or asymptotes (CoRDI). Fits are not provided for the IAI and CDI results due to the large effect of the systematics. For the CoRDI case, the gray line is a piecewise function with one  $A\sqrt{n} + B$  fit for 0 to 300 and another for 300 to 1500. For the dim planet case, we see that the CoRDI method outperforms IAI and CDI by a factor of  $\sim 2$ . In addition, the



**Fig. 8** Benefits of a stable wavefront and effective post-processing pipeline demonstrated using a 2200 iteration ADI DZM experiment with a  $4 \times 10^{-8}$  planet injected (see Sec. 2.5.2 for more information). (a) We have the final open-loop image for the experiment; this would be the result in the set-and-forget approach used by JWST and HST and planned for RST. (b) We show the final closed-loop image in the case in which the DZ is maintained; the planet is still not visible, but the speckles have been suppressed. (c) We have an SNR map of the post-processed image where the  $4 \times 10^{-8}$  planet is clearly visible. Note that, for the SNR maps, we use the PuBu colormap.

CoRDI SNR does not drastically decrease as a function of the number of KLIP modes (as we show in Sec. 5) and varies by  $<15\%$  for KLIP modes from 1 to 50.

#### 4 Discussion: Implications for Future Space-Based Missions

For on-sky direct imaging instruments, the limit to exoplanet detection and characterization will be contrast stability as post-processing benefits drastically decline when the electric field is drifting among exposures. The current space-based high-contrast instruments (HST, JWST) do not have DMs, and the contrast stability is directly coupled to the stability of the telescope. HST and JWST thus require specific observing scenarios (reference stars, rolling on targets, and target revisits) to obtain the necessary data to enable post-processing success in which the planet signal is extracted. These observing scenarios are inefficient and reduce the science time of the observatory. In the case of future missions such as RST or HWO, we will have DMs that facilitate alternate approaches. In theory, the DMs can lower the residual starlight speckles to arbitrarily low levels, and post-processing is not needed if the DZ can be suppressed below the level of the planet. In practice, even state-of-the-art experiments (Seo et al.<sup>46</sup>) do not obtain the desired level of starlight suppression even with DMs, so most future mission concepts propose a combination of telescope stability and post-processing requirements. In Fig. 8, we show the drastic benefits of both a stable wavefront and an effective post-processing pipeline.

In previous papers, we showed that maintaining the contrast was possible both in simulations (Pogorelyuk et al.<sup>25</sup>) and in practice (Redmond et al.<sup>4</sup>) using the DMs not to generate a dark zone but to maintain it in the presence of drifts. This maintenance scheme uses a single focal plane image per iteration (rather than the difference between a pair of images) and is based on dithering the DMs to perturb the dark zone. It is thus legitimate to investigate (1) if the presence of a bright incoherent source in the dark zone would affect the maintenance performance and (2) if this approach is compatible with post-processing techniques. These are the questions that we answer in this paper by reporting on the outcome of a series of dedicated experiments on the HiCAT testbed. We find that minor adjustments to the DZM algorithm ensure good performance in the presence of bright planets and that post-processing is possible using DZM data. We also demonstrate an approach in which DZM enables post-processing with no observing inefficiencies (reference stars, telescope rolls, etc.).

By augmenting the EKF and estimating the incoherent light separately, DZM performance is not affected by a planet two orders of magnitude brighter than the dark zone ( $10^{-6}$  planet in a  $10^{-8}$  DZ). The planet light is sufficiently captured by the incoherent state and is not passed to the controller to be suppressed; this avoids unstable controller behavior that would generate a large command to attempt (and fail) to suppress a bright “speckle.” Splitting the estimate enables IAI,



CDI, and CoRDI, removing the need to slew the telescope to a reference star or roll the telescope on the target star. This has profound implications on the operations concept of future missions as the roll or slew associated with ADI or RDI might induce thermally induced wavefront drifts of the telescope assembly. An example of such is shown for RST in Krist et al.<sup>47</sup> (available on the IPAC website). Although DZM can correct for the thermal-induced drifts within the instrument, it would be preferable not to move the telescope at all to increase the observing efficiency.

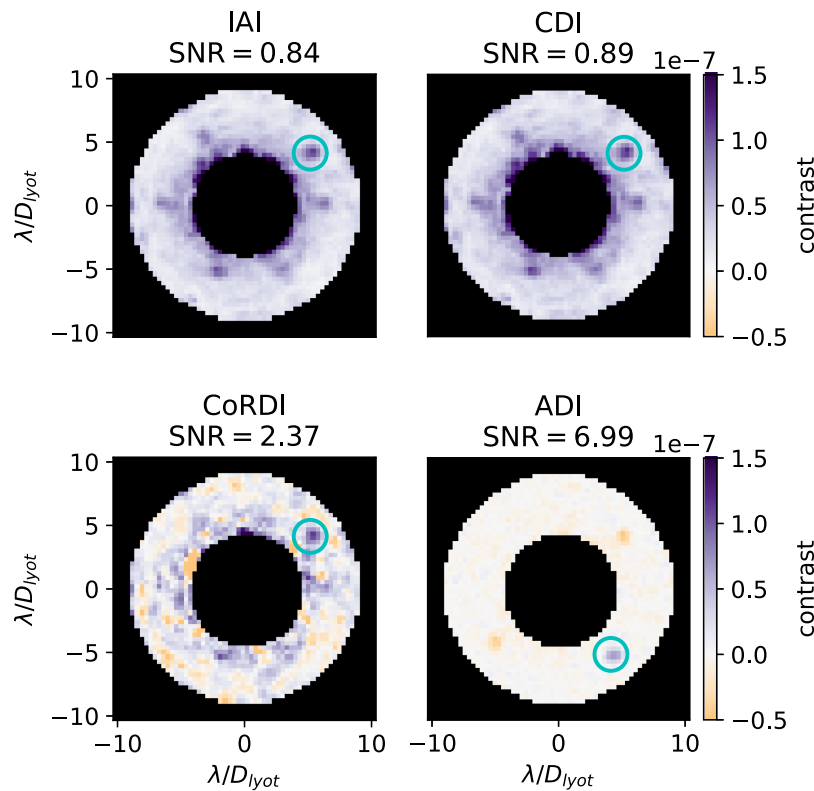
Moving on to post-processing investigations, we first confirm that dithering the DMs does not alter the ability to use PCA-based post-processing methods, such as ADI, commonly used by current missions. Our ADI results exhibit the scaling of SNR versus time to be proportional to  $\sqrt{t}$  as is expected when the analysis is limited by random noise sources; this scaling is not expected if systematics associated with residual starlight speckles are the dominating noise source in the images. It should be noted that this is a very optimistic implementation of ADI as the roll angle is 90 deg putting the planet in a very different location in the dark zone. To put this in perspective, JWST can only roll 14 deg, and RST rolls are planned to be 26 deg. For the upcoming HWO, the limitations will be similar to or more stringent than JWST, which may degrade the ADI performance.

We then tested three approaches for CDI (CDI, IAI, and CoRDI) that leverage the AEKF. The three approaches work equally well for the case of a bright planet, showing the  $\text{SNR} \propto \sqrt{t}$ . In the case of a faint planet, they are limited by a current model-mismatch systematic limitation of HiCAT. As a result, a small amount of the residual “starlight” speckles are captured by the incoherent estimate, and a small amount of the planet light is captured by the coherent estimate. In the dim planet case, this estimate error is sufficient such that the planet cannot be distinguished from the speckles when IAI and CDI are used. We then show that the systematic floor can be lowered if the coherent estimates are combined to create a reference PSF using a PCA framework (CoRDI). In that context, the planet gets detected, and the SNR shows the  $\sqrt{t}$  growth up until PCA is not sufficient to remove the estimate error systematics. An image combination method that captures higher order moments of the systematic in the coherent estimates or a better model of HiCAT in the future will alleviate this floor in the future. In a real space mission, target revisits would also assist in determining what is a speckle versus a planet. Regardless, our successful CoRDI results demonstrate, for the first time in a high-contrast experiment, that continuous maintenance without slewing/rolling the telescope can not only maintain raw contrast in the presence of large drifts in the observatory but also enable post-processing of faint planets under the raw contrast floor. No extra hardware is needed to execute this, just an algorithm to acquire images and control DMs making it a promising option for RST or HWO.

Figure 9 shows the sample detection results using each of the four post-processing techniques previously discussed. For all of the techniques, 1500 frames are used to generate the post-processed image, and an  $8 \times 10^{-8}$  planet is injected into the dark zone. For the ADI experiment, the mean DZ contrast is  $7.7 \times 10^{-8}$  (see Fig. 14), and for the RDI/CDI/IAI experiment (Sec. 3.2.2), the mean DZ contrast is  $8.9 \times 10^{-8}$ . The top two panels of Fig. 9 show the results for the simple stacking post-processing methods, and the bottom two show the KLIP-based post-processing methods. As we can see, ADI performs the best, but for a real space-based high-contrast imaging system it comes at the cost of observing time and potentially increasing the drifts due to the change of thermal environment. As previously discussed, the CoRDI/IAI/CDI approaches are limited by a model mismatch combined with sub-optimal tuning of AEKF parameters. ADI does not depend on the estimator output for post-processing, so these limitations are not applicable, which results in ADI currently outperforming CoRDI/IAI/CDI. It is predicted that this gap in performance can be drastically reduced with upcoming work.

## 5 Conclusions and Future Work

This paper demonstrates how we can augment the DZM EKF to estimate the incoherent light in the image and how DZM data can effectively be used with a number of different post-processing methods. Post-processing with DZM data is still in the early stages of development, and there are a number of future research avenues. Currently, when using the nominal DZM algorithm with ADI, the dither patterns applied to the DMs start to repeat whenever the planet is “rolled.” If the timescale was shorter, it is possible that pyklip would perform better as the residual starlight

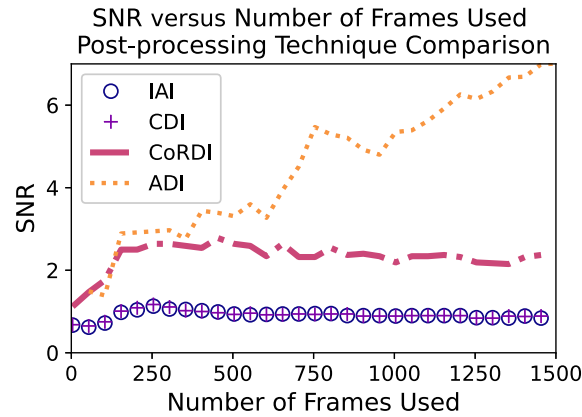


**Fig. 9** Final post-processed images for all post-processing techniques for an  $8 \times 10^{-8}$  planet. Results from an ADI experiment are included in the bottom right panel to provide a full comparison. The SNR is calculated using the histogram overlap method discussed in App. B. The two PCA-based approaches (ADI, CoRDI) perform the best, followed by the stacking methods.

speckle pattern would be more consistent. It would be worthwhile to investigate the shortest allowable dither-repeat cycle, but it should be noted that it must be longer than the memory of the EKF for the EKF to be stable.

Using the AEKF with DZM is a novel concept, and there are many exciting paths forward. Repeating the dither commands could also potentially aid with the CoRDI approach for a similar reason, so experiments should be run with the AEKF and dither-cycling. One could also run AEKF with planet-PSF-sized blocks instead of single pixels. The  $H$  matrix [Eq. (8)] would then be  $3n_p \times 3n_p$ , where  $n_p$  is the number of pixels in an expected planet PSF. As the incoherent estimate finds the planet, the blocking of the dark zone could be dynamically updated to center the block on the planet. It would also be interesting to look into the tuning of the  $Q$  matrix and have it dynamically update throughout the experiment. Often when using an EKF, the  $Q$  matrix is made to be more aggressive at the start and taper off. It is possible that, by implementing a dynamic  $Q$  matrix, we could avoid accumulating the coherent estimate error in the incoherent estimate, which would improve the CDI and IAI post-processing results. As exozodiacal dust is predicted to be a potential issue for directly imaging exo-Earths,<sup>39</sup> it would also be beneficial to test the AEKF with representative exozodiacal dust emission injected into the images. Finally, it would be good to look into using the AEKF with ADI and develop a methodology for re-initializing the incoherent estimate when a roll is performed. The AEKF is much more robust for dealing with rolls as it does not require re-initializing the coherent electric field estimate as that will not change. For this reason, it should be relatively trivial to avoid the peaks seen in Fig. 11 when the nominal DZM algorithm is used when a bright planet is present.

As shown in Fig. 10, there is room for improvement when it comes to the IAI, CDI, and CoRDI post-processing methods. All three methods encounter a systematic that causes the SNR to level off after the first 250 frames. Having an accurate estimate of the number of frames required to obtain a desired SNR on a planet is imperative for yield calculations and target planning for any high-contrast instrument. The adjustments described above to the AEKF should help



**Fig. 10** SNR versus the number of frames used for all post-processing techniques for an  $8 \times 10^{-8}$  planet. The SNR is calculated using the histogram overlap method discussed in App. B. The ADI data follow  $\sqrt{t}$  for all 1500 iterations, but the other methods hit systematics and level off. With the recent upgrades of the HiCAT software and hardware, these systematics should now be much less prominent. It is clear that IAI, CDI, and CoRDI are more sensitive to systematics in general, and work is ongoing to improve their performance.

all three post-processing methods as they will improve the incoherent estimate. We could also try testing the effect of a target re-visit on CoRDI; this would essentially be combining CoRDI with ADI. Finally, the HiCAT model and testbed have been upgraded since these experiments were performed, and these upgrades may remove a number of the systematics seen in this data. The coronagraph is now a pupil-apodized Lyot coronagraph, and the sixfold spike pattern (top right panel in Fig. 7) is no longer the dominant residual when the dark zone is generated. The model has also been better calibrated with the testbed, so model-induced errors in the AEKF resulting in mixing of the incoherent and coherent signals should be reduced.

## 6 Appendix A: Dark Zone Maintenance Parameter Definitions

### 6.1 Nominal Dark Zone Maintenance

This section is taken from Redmond et al.<sup>4</sup> and provided here as an easy reference. In Table 3, we provide the conversion between traditional WFSC notation and the EKF notation. Note that, in both notations, a  $\hat{\cdot}$  indicates that the variable is an estimate not the true value or measurement.

The Jacobian and state variable are

$$G = \sqrt{t_{\text{exp}}} \begin{bmatrix} \Re\left(\frac{\partial x_0}{\partial u}\right) \\ \Im\left(\frac{\partial x_0}{\partial u}\right) \\ \vdots \\ \Re\left(\frac{\partial x_{n-1}}{\partial u}\right) \\ \Im\left(\frac{\partial x_{n-1}}{\partial u}\right) \end{bmatrix}, \quad (26)$$

$$x = \begin{bmatrix} \Re(x_0) \\ \Im(x_0) \\ \vdots \end{bmatrix}, \quad (27)$$

where the pixel index in the dark zone ranges from 0 to  $(n - 1)$  and  $G$  is the BMC Jacobian. Note that  $G$  is the total BMC Jacobian, where the Jacobians for BMC DM1 and BMC DM2 are stacked horizontally ( $G = [G_{\text{DM1}} \ G_{\text{DM2}}]$ ). The EKF estimator operates in units of counts, so the Jacobians are in units of  $[(\text{counts/s})^{1/2}/\text{nm}]$  and must be multiplied by the square root of the exposure time ( $\sqrt{t_{\text{exp}}}$ ) each iteration. The nonlinear state space model is described as

**Table 3** Conversion table between WFSC notation and EKF notation.

Parameter	WFSC notation	EKF notation
Open-loop electric field	$E_{OL}$	$x$
Estimate of open-loop electric field	$\hat{E}_{OL}$	$\hat{x}$
Closed-loop electric field	$E_f$	$x_{CL}$
Measured closed-loop intensity	$I$	$z$
DM command	$u$	$u$

$$x^{k+1} = f(x^k) + w^k = x^k + w^k, \quad (28)$$

$$x_{CL}^k = x^k + Gu^k, \quad (29)$$

$$z^{k+1} = h(x^k, u^k) + n^k = x_{CL_{Re}}^k \circ x_{CL_{Re}}^k + x_{CL_{Im}}^k \circ x_{CL_{Im}}^k + n^k, \quad (30)$$

where  $\circ$  is the Hadamard operator,  $u^k$  is the DM command for the current iteration,  $x_{CL}^k$  is the closed-loop electric field, and  $w^k$  and  $n^k$  are the process and observation noise terms. As with the BMC Jacobians, the DM command  $u$  is a stacked vector containing the command for both BMC DMs ( $u = [u_{DM1} \ u_{DM2}]^T$ ).

## 6.2 Low-Photon Conversion

This process is described in Redmond et al.<sup>26</sup> We start by taking a direct image on HiCAT (without the Lyot stop in place) using the longest exposure time possible without saturating any pixels. This minimizes the noise from the HiCAT detector in the image. We normalize the image ( $I_{direct}$ ) to the exposure time ( $t_{direct}$ ) and sum the count rate to get

$$dI_{direct} = \frac{I_{direct}}{t_{direct}}, \quad (31)$$

$$F_{HiCAT} = \sum dI_{direct}, \quad (32)$$

where  $dI_{direct}$  is the direct HiCAT image in counts/s and  $F_{HiCAT}$  is the total count rate in the image. We then choose a photon rate ( $F$ ) to which we want to scale the image and determine the counts-to-photons conversion factor ( $\gamma$ ) for any type of image (direct or coronagraph) via

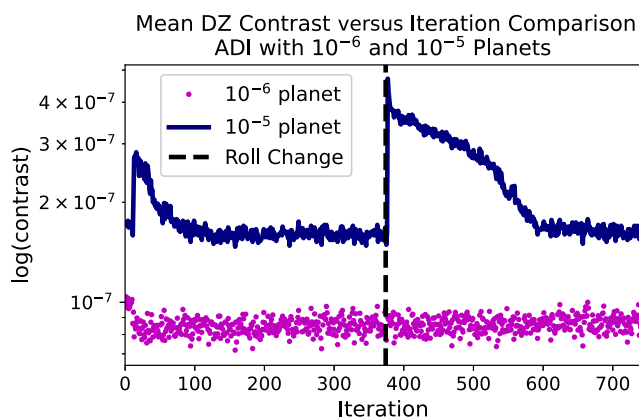
$$\gamma = \frac{F}{F_{HiCAT}}. \quad (33)$$

Note that  $\gamma$  is calculated once via the direct image and then used to convert the coronagraph images to photons/s. For this paper, we choose  $F = 15 \times 10^6$  ph/s. For reference, the direct image of the target star in the RST CGI Observing Scenario (47 UMa,  $V = 5.04$ , G1V,  $D_{star} = 0.9$  mas)<sup>48</sup> has a total photon rate (the sum of the photon rate over the entire image) of  $\sim 18.5 \times 10^6$  ph/s.  $F$  is intentionally slightly less than the RST photon rate quoted above. HiCAT operates at a lower contrast than RST and will thus have more photons in the dark zone for a given direct image photon rate.

We then scale the images and apply the desired noise characteristics. For the purpose of this paper, we use a very basic noise model that only includes dark current and read noise. For each image ( $I_{HiCAT}$ ), we first normalize by the exposure time ( $t_{HiCAT}$ )

$$dI_{HiCAT} = \frac{I_{HiCAT}}{t_{HiCAT}}, \quad (34)$$

to get  $dI_{HiCAT}$  in counts/s and then apply the scaling and noise model to obtain a low-photon image ( $I_{lsnr}$ ) to use in the estimator, given as



**Fig. 11** Simulated HiCAT contrast versus time data for a bright planet ADI experiment using the nominal DZM implementation (no AEKF). The magenta dots curve has a  $10^{-6}$  planet injected, and the mean DZ contrast creeps up over the course of the experiment. The blue curve has a  $10^{-5}$  planet injected and struggles to obtain a good estimate at the start as well as to update the estimate for the new planet location when the roll angle changes at iteration 375 (black dashes).

$$I_{\text{poisson}} = P((\gamma d I_{\text{HiCAT}} + \mu_I)t), \quad (35)$$

$$I_{\text{lsnr}} = \text{round}(I_{\text{poisson}} + \mathcal{N}(0, \sigma_{rn}^2 \mathcal{I})), \quad (36)$$

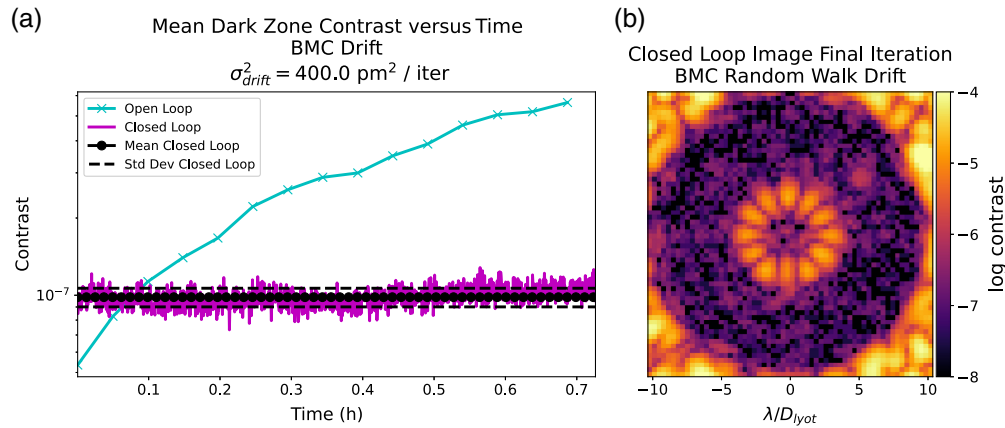
where  $P()$  is taking the Poisson distribution,  $\mu_I$  is the dark current in e-/s,  $t$  is the RST-like exposure time in seconds,  $I_{\text{poisson}}$  is an image (in units of photons) that contains Poisson noise and dark current only,  $\sigma_{rn}$  is the standard deviation of the read noise, and  $\mathcal{I}$  is the identity matrix with the same dimension as the image. The “round” is included as it is not physical to have a fraction of an electron. Note that the Jacobian and covariance matrices (for EFC and the EKF) need to be calculated with or scaled by  $t$  and  $I_{\text{lsnr}}$  as opposed to  $t_{\text{HiCAT}}$  and  $I_{\text{HiCAT}}$ .

### 6.3 Bright Incoherent Light and DZM

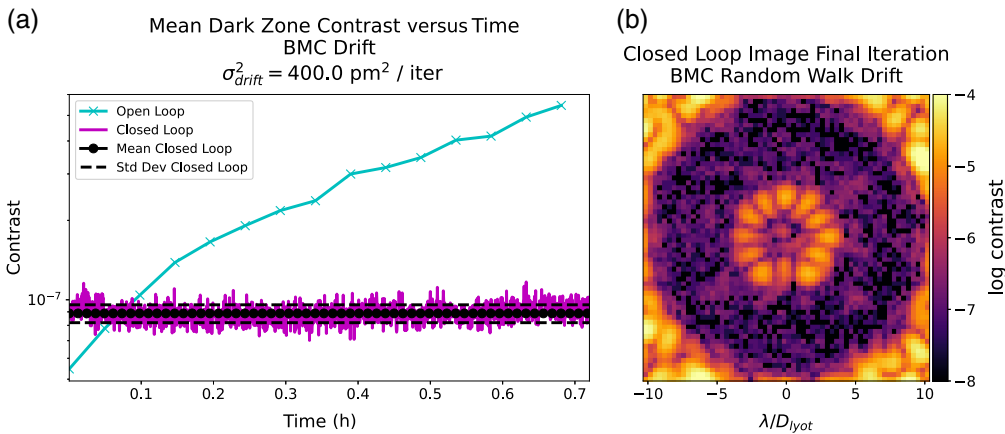
One potential limitation of the original DZM implementation is the estimation bias caused by bright planets or bright exozodiacal dust emission.<sup>39</sup> If these incoherent light sources are very bright, the estimator will capture it and subsequently, the DMs will try to suppress it. However, as the light is incoherent, the DMs will not succeed in correcting for it, which can cause the control loop to diverge or struggle. This is especially true when considering an ADI observing scenario because the EKF has memory. If a bright speckle abruptly moves in the dark zone, it takes a number of iterations for the EKF to “forget” the initial location and “learn” the new location. This is demonstrated in Fig. 11, which shows the contrast versus time for two simulated DZM runs in which a bright planet is injected and the planet location is moved at iteration 375. The initial mean DZ contrast without the planet is  $5 \times 10^{-8}$ , and the planet contrast is either  $10^{-6}$  (magenta dots) or  $10^{-5}$  (blue). The dashed black line shows where the roll angle of the planet changes. For both cases, one can see that the estimator takes longer than normal to converge; note that these peaks at the beginning could be reduced by turning on the controller later. For the  $10^{-6}$  planet, we see the mean DZ contrast degrade over the 750 iterations as the controller tries (and fails) to suppress the planet mistaken as a bright speckle. For the  $10^{-5}$  planet, the system struggles when the roll angle changes. This is something that could be mitigated by disabling the controller for a period of time after a roll change, but while the controller is off, drifts will leak in. If the controller has to be disabled for a prolonged period every roll change, which can result in contrast degradation due to changing wavefront errors.

### 6.4 DZM Experiments

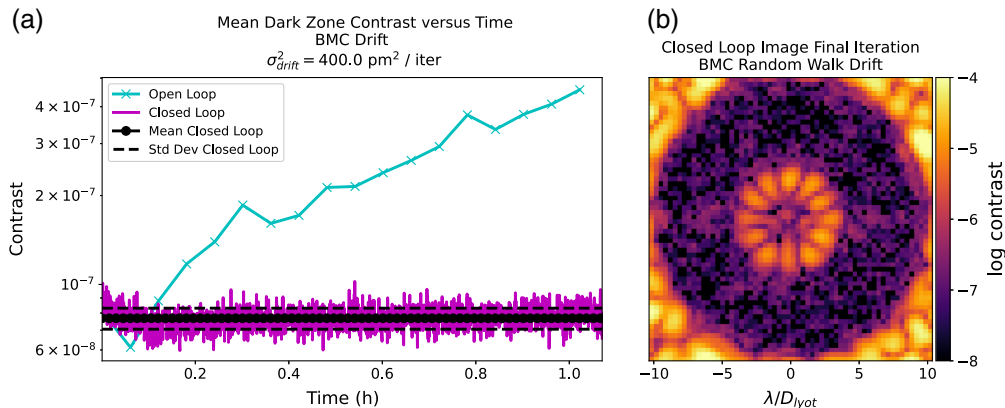
Here, we provide additional contrast versus time plots for experiments discussed in this paper. Figure 12 is discussed in Sec. 3.2.1, Fig. 13 is discussed in Sec. 3.2.2, and Fig. 14 is discussed in Sec. 4. Note that the time is the actual time it took to run the experiment. It should also be noted



**Fig. 12** AEKF DZM experiment results with a bright planet injected. (a) Mean DZ contrast versus time for the low-photon case using the AEKF version of DZM with a  $1 \times 10^{-6}$  planet. (b) Sample science image from the final iteration. The planet is visible to the eye in the dark zone. This experiment is discussed in Sec. 3.2.1.



**Fig. 13** AEKF DZM experiment results with a dim planet injected. (a) Mean DZ contrast versus time for the low-photon case using the AEKF version of DZM with an  $8 \times 10^{-8}$  planet. (b) Sample science image from the final iteration. The planet is not visible to the eye in the dark zone. This experiment is discussed in Sec. 3.2.2.



**Fig. 14** ADI DZM experiment results with dim planet injected. (a) Mean DZ contrast versus time for the low-photon case using the nominal DZM algorithm with an  $8 \times 10^{-8}$  planet. (b) Sample science image from the final iteration. The planet is not visible to the eye in the dark zone. This experiment is discussed in Sec. 4.

that the planet pixels are included in the mean DZ contrast calculation; this contributes to the higher mean DZ contrast for the experiment with the  $1 \times 10^{-6}$  planet.

## 7 Appendix B: SNR Calculations

Figure 15 shows a sample of the overlapping planet and DZ histograms for RDI, IAI, and CDI post-processed images for the experiment discussed in Secs. 3.2.2 and 6.4 (Fig. 13). For the IAI and CDI data, the histogram of the dark zone is very skewed. This means that it is not very accurate to calculate the SNR via the classic  $\text{SNR} = I_{\text{planet}}/\sigma_{\text{DZ}}$  method or using `pyklip`. Instead, we normalize the DZ and planet histograms such that the total area is one and use the percent overlap of the two histograms to determine a more representative SNR proxy. The overlap is taken to be the entire region of the DZ histogram above the lowest non-zero planet bin. The percent overlap is then used to produce a Z-score via a “one-tail test,” which we use as an SNR proxy. The Z-score is calculated via


$$p = 1 - \frac{\text{percent\_overlap}}{100}, \quad (37)$$

$$Z = \frac{p - \mu_{\text{DZ}}}{\sigma_{\text{DZ}}}, \quad (38)$$

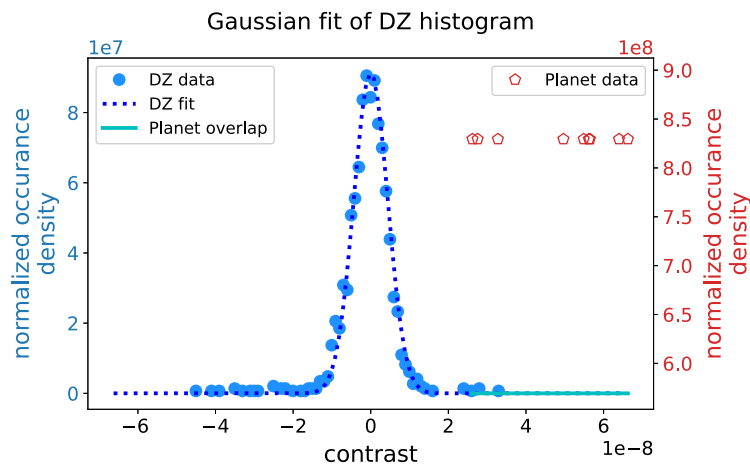
where  $\mu_{\text{DZ}}$  is the mean of the DZ and  $\sigma_{\text{DZ}}$  is the standard deviation of the dark zone. Note that, in Fig. 15, the RDI DZ histogram is much more Gaussian and symmetric than the IAI and CDI data, but there is sufficient overlap between the planet and DZ data such that this method is still valid.

We wish to compare all four post-processing techniques for the  $8 \times 10^{-8}$  planet case, which means also adjusting the method by which we calculate the SNR for the ADI data. For this, we are considering the ADI experiment shown in Fig. 14 in Sec. 6.4. In the case of the ADI data, there is not sufficient overlap of the histograms to use the numerical method described above. Luckily, the ADI DZ is sufficiently Gaussian that we can fit a Gaussian to it and extrapolate such that the DZ histogram and planet histogram overlap; this is shown in Fig. 16. The process is as follows:

1. fit Gaussian to DZ data
2. symmetrically extrapolate Gaussian to the planet data
3. normalize the extrapolated Gaussian to have an area of one (dashed blue line in Fig. 16)
4. calculate the percent overlap of the planet data and the fit (solid cyan line in Fig. 16)
5. calculate Z-score via Eq. (38).



**Fig. 15** Histogram overlap of DZ (blue) and planet (red) data. Data are drawn from the final post-processed image for the IAI, CDI, and CoRDI methods.



**Fig. 16** Demonstration of the components to calculate the planet SNR for an ADI experiment with an  $8 \times 10^{-8}$  injected planet. Right axis: normalized occurrence density of a given contrast in a  $3 \times 3$  square around the center of the injected planet (red pentagons). Note that, as there are so few pixels that make up the planet signal, we do not try to fit the planet data. Left axis: Gaussian fit of the DZ excluding the planet pixels. The light blue circles show the DZ data, and the dark blue dashed line shows the extrapolated Gaussian fit; note that the DZ data and fit are normalized such that the area under the curve is one. The cyan solid line shows the overlap between the DZ fit (dotted blue) and the planet data (red pentagons). The integral of the planet overlap line can then be used to calculate the planet Z-score.

## Disclosures

The authors declare no conflicts of interest.

## Code and Data Availability

HiCAT makes use of the NumPy,<sup>49</sup> Matplotlib,<sup>50</sup> AstroPy,<sup>51,52</sup> SciPy,<sup>53</sup> ImageIO,<sup>54</sup> HCIPy,<sup>55</sup> and catkit2<sup>30</sup> packages. The data utilized in this study were obtained from the Russell B. Makidon Optics Laboratory at the Space Telescope Science Institute. Data are available from the authors upon request and with permission from the Russell B. Makidon Optics Laboratory.

## Acknowledgments

The authors are especially thankful to the extended HiCAT team (over 50 people) who have worked over the past several years to develop this testbed. This work was supported in part by the National Aeronautics and Space Administration (Grant No. 80NSSC19K0120) issued through the Strategic Astrophysics Technology/Technology Demonstration for Exoplanet Missions Program (SAT-TDEM; PI: R. Soummer). E.H.P. was supported by the NASA Hubble Fellowship (Grant No. HST-HF2-51467.001-A) awarded by the Space Telescope Science Institute, which is operated by the Association of Universities for Research in Astronomy, Incorporated, under NASA (Grant No. NAS5-26555). I.L. acknowledges the support by a postdoctoral grant issued by the Centre National d'Études Spatiales (CNES) in France. The research was carried out in part through the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration (Grant No. 80NM0018D0004).

## References

1. “Interactive Overview: Pathways to Discovery in Astronomy and Astrophysics for the 2020s,” National Academies of Sciences Engineering (2021).
2. L. E. Coyle et al., “Achieved technology maturation of key component-level technologies for ultra-stable optical systems,” *Proc. SPIE* **12180**, 121802K (2022).
3. L. E. Coyle and J. S. Knight, “Integrated modeling of large, segmented telescopes with ultra-stable wavefronts,” *Proc. SPIE* **12180**, 121805Z (2022).
4. S. F. Redmond et al., “Implementation of a dark zone maintenance algorithm for speckle drift correction in a high contrast space coronagraph,” *J. Astron. Telesc. Instrum. Syst.* **8**, 035001 (2022).



5. P. K. Poon et al., “Experimental demonstration of spectral linear dark field control at NASA’s high contrast imaging testbeds,” *Proc. SPIE* **12680**, 492–503, (2023).
6. A. L. Carter et al., “The JWST early release science program for direct observations of exoplanetary systems I: high-contrast imaging of the exoplanet HIP 65426 b from 2 to 16 m,” *Astrophys. J. Lett.* **951**, L20 (2023).
7. J. Kasting et al., “Exoplanet characterization and the search for life,” arXiv:0911.2936 (2009).
8. B. Nemati et al., “Method for deriving optical telescope performance specifications for Earth-detecting coronagraphs,” *J. Astron. Telesc. Instrum. Syst.* **6**, 039002 (2020).
9. L. Pogorelyuk et al., “Information-theoretical limits of recursive estimation and closed-loop control in high-contrast imaging,” *Astrophys. J. Suppl. Ser.* **256**, 39 (2021).
10. G. Schneider and M. D. Silverstone, “Coronagraphy with HST/NICMOS: detectability is a sensitive issue,” *Proc. SPIE* **4860**, 1–9 (2003).
11. D. A. Fraquelli, A. B. Schultz, and H. Bushouse, et al., “NICMOS coronagraphy,” *Publ. Astron. Soc. Pa.c* **116**, 55–64 (2003).
12. C. Marois et al., “Angular differential imaging: a powerful high-contrast imaging technique,” *Astrophys. J.* **641**, 556 (2006).
13. D. Lafrenière et al., “HST/NICMOS detection of HR 8799 b IN 1998,” *Astrophys. J.* **694**, L148 (2009).
14. R. Soummer et al., “Orbital motion of HR 8799 b, c, d using Hubble space telescope data from 1998: constraints on inclination, eccentricity, and stability,” *Astrophys. J.* **741**, 55 (2011).
15. R. Galicher et al., “Self-coherent camera as a focal plane wavefront sensor: simulations,” *Astron. Astrophys.* **509**, A31 (2010).
16. M. Bottom et al., “Speckle suppression and companion detection using coherent differential imaging,” *Mon. Not. R. Astron. Soc.* **464**, 2937–2951 (2017).
17. W. H. Smith, “Spectral differential imaging detection of planets about nearby stars,” *Publ. Astron. Soc. Pac.* **99**, 1344 (1987).
18. W. B. Sparks and H. C. Ford, “Imaging spectroscopy for extrasolar planet detection,” *Astrophys. J.* **578**, 543 (2002).
19. J. Krist, “Roman CGI observing scenario 11 time series simulations,” Nancy Grace Roman Space Telescope at IPAC, [https://roman.ipac.caltech.edu/data/sims/coronagraph/os11/spc\\_spec\\_os11\\_version3.pdf](https://roman.ipac.caltech.edu/data/sims/coronagraph/os11/spc_spec_os11_version3.pdf) (2022).
20. “Aperture position angle special requirements—JWST user documentation,” (2022).
21. E. Choquet et al., “First images of debris disks around TWA 7, TWA 25, HD 35650, and HD 377,” *Astrophys. J. Lett.* **817**, L2 (2016).
22. J.-B. Ruffio et al., “JWST-TST high contrast: achieving direct spectroscopy of faint substellar companions next to bright stars with the NIRSpec IFU,” arXiv:2310.09902 [astro-ph] (2023)
23. J. Wang et al., “High resolution spectroscopy of directly imaged exoplanets with KPIC,” *Proc. SPIE* **11823**, 1182302 (2021).
24. K. B. Follette, “An introduction to high contrast differential imaging of exoplanets and disks,” *Publ. Astron. Soc. Pac.* **135**, 093001 (2023).
25. L. Pogorelyuk and N. J. Kasdin, “Dark hole maintenance and a posteriori intensity estimation in the presence of speckle drift in a high-contrast space coronagraph,” *Astrophys. J.* **873**, 95 (2019).
26. S. F. Redmond et al., “Dark zone maintenance for future coronagraphic space missions,” *Proc. SPIE* **12180**, 869–877 (2022).
27. R. Soummer et al., “High-contrast Imager for Complex Aperture Telescopes (HiCAT): 5. first results with segmented-aperture coronagraph and wavefront control,” *Proc. SPIE* **10698**, 106981O (2018).
28. M. N’Diaye et al., “High-contrast Imager for Complex Aperture Telescopes (HiCAT): 3. first lab results with wavefront control,” *Proc. SPIE* **9605**, 96050I (2015).
29. R. Soummer et al., “High-contrast imager for complex aperture telescopes (HiCAT): 8. Dark zone demonstration with simultaneous closed-loop low-order wavefront sensing and control,” *Proc. SPIE* **12180**, 1218026 (2022).
30. E. Por et al., “The control and automation for Testbeds Kit 2 (CATKit2),” Zenodo (2024).
31. J. Noss et al., “Spacetelescope/catkit: v0.36.1,” (2021).
32. J. Fowler et al., “The generalized lab architecture for restructured optical experiments (GLARE),” in *Am. Astron. Soc. Meeting Abstr.*, Vol. 235, pp. 373.09 (2020).
33. C. Moriarty et al., “High-contrast imager for complex Aperture Telescopes (HiCAT): 6. Software Control Infrastructure and Calibration,” *Proc. SPIE* **10698**, 1069853 (2018).
34. L. Pogorelyuk et al., “Reduced order estimation of the speckle electric field history for space-based coronagraphs,” *Astrophys. J.* **881**, 126 (2019).
35. S. F. Redmond et al., “Exoplanet detection techniques for direct imaging dark zone maintenance data sets,” *Proc. SPIE* **13092**, 1309224 (2024).
36. 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**, 011009 (2015).

37. A. Give'on et al., "Broadband wavefront correction algorithm for high-contrast imaging systems," *Proc. SPIE* **6691**, 66910A (2007).
38. A. J. E. Riggs et al., "Recursive starlight and bias estimation for high-contrast imaging with an extended Kalman filter," *J. Astron. Telesc. Instrum. Syst.* **2**, 011017 (2016).
39. A. Roberge et al., "The exozodiacal dust problem for direct observations of exo-earths," *Publ. Astron. Soc. Pac.* **124**, 799 (2012).
40. A. J. E. Riggs, *Integrated Wavefront Correction and Bias Estimation for the High-Contrast Imaging of Exoplanets*, Princeton University, New Jersey, United States (2016).
41. R. Soummer et al., "Detection and characterization of exoplanets and disks using projections on Karhunen-Loève eigenimages," *Astrophys. J.* **755**, L28 (2012).
42. L. Pueyo, "Detection and characterization of exoplanets using projections on Karhunen-Loeve eigenimages: forward modeling," *Astrophys. J.* **824**, 117 (2016).
43. J. J. Wang et al., "pyKLIP: PSF subtraction for exoplanets and sisks," *Astrophys. Source Code Lib. ADS Bibcode: 2015ascl.soft06001W* (2015).
44. B. Nemati et al., "HabEx Telescope WFE stability specification derived from coronagraph starlight leakage," *Proc. SPIE* **10743**, 109–120 (2018).
45. F. Cantalloube et al., "Direct exoplanet detection and characterization using the ANDROMEDA method: performance on VLT/NaCo data," *Astron. Astrophys.* **582**, A89 (2015).
46. B.-J. Seo et al., "Testbed demonstration of high-contrast coronagraph imaging in search for Earth-like exoplanets," *Proc. SPIE* **11117**, 111171V (2019).
47. J. Krist, "WFIRST CGI OS9 Time Series Simulations (Hybrid Lyot Coronagraph, Band 1)," Roman Space Telescope Coronagraph Instrument Public Simulated Images, [https://roman.ipac.caltech.edu/sims/Coronagraph\\_public\\_images.html](https://roman.ipac.caltech.edu/sims/Coronagraph_public_images.html) (2020).
48. J. Krist, "Roman CGI Observing Scenario 11 Time Series Simulations (Hybrid Lyot Coronagraph, Band 1)," Roman Space Telescope Coronagraph Instrument Public Simulated Images, [https://roman.ipac.caltech.edu/sims/Coronagraph\\_public\\_images.html#CGI\\_OS11\\_report](https://roman.ipac.caltech.edu/sims/Coronagraph_public_images.html#CGI_OS11_report) (2022).
49. C. R. Harris et al., "Array programming with NumPy," *Nature* **585**, 357–362 (2020).
50. J. D. Hunter, "Matplotlib: a 2D graphics environment," *Comput. Sci. Eng.* **9**(3), 90–95 (2007).
51. T. P. Robitaille et al., "Astropy: a community Python package for astronomy," *Astron. Astrophys.* **558**, A33 (2013).
52. T. A. Collaboration et al., "The Astropy Project: building an inclusive, open-science project and status of the v2.0 core package," *Astron. J.* **156**, 123 (2018).
53. P. Virtanen et al., "SciPy 1.0: fundamental algorithms for scientific computing in Python," *Nature Methods* **17**, 261–272 (2020).
54. S. Silvester et al., "Imageio/imageio v2.8.0," Zenodo (2020).
55. E. H. Por et al., "High contrast imaging for Python (HCIPy): an open-source adaptive optics and coronagraph simulator," *Proc. SPIE* **10703**, 1070342 (2018).

**Susan F. Redmond** is a David and Ellen Lee Distinguished Scholar fellow at the California Institute of Technology. She received her BEng from the Memorial University of Newfoundland and Labrador, her MASc from the University of Toronto, and her PhD from Princeton University. Her research interests focus on the questions of "How did we get here?" and "Are we alone?" via cluster weak lensing and directly imaging exoplanets. She focuses on the optical system characterization for the Super-pressure Balloon-borne Imaging Telescope, which includes analysis of the 39 days of data acquired during the 2023 science flight, and on developing and implementing novel focal plane wavefront estimation and control algorithms for exoplanet direct imaging adaptive optics systems as part of the Exoplanet Technology Laboratory at Caltech and the High Contrast Imaging Testbed team at Jet Propulsion Laboratory (JPL).

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

**Emiel H. Por** is an NASA Hubble Fellowship Program (NHFP) Sagan fellow at the Russell B. Makidon Optics Laboratory at the Space Telescope Science Institute in Baltimore, United States. He received his BSc degree in astronomy and physics, his MSc degree in astronomical

instrumentation, and his PhD in astronomy from Leiden University, the Netherlands. His research interests include coronagraphy, wavefront sensing and control, and high-contrast imaging.

**Iva Laginja** is a postdoctoral researcher in astronomical instrumentation at the Laboratoire d'Etudes Spatiales et d'Instrumentation en Astrophysique (LESIA) in Meudon near Paris. Previously, she was an astronomical optics scientist at the Space Telescope Science Institute. She received her doctorate degree in astronomy and instrumentation at STScI/ONERA/LAM.

**N. Jeremy Kasdin** is the Eugene Higgins professor of mechanical and aerospace engineering, emeritus, at Princeton University and is the former assistant dean for engineering at the University of San Francisco. He was the chief systems engineer for NASA's Gravity Probe B spacecraft. From 2014 to 2016, he was vice dean of the School of Engineering and Applied Science at Princeton. From 2014 to 2021, he was the adjutant scientist for the coronagraph instrument on NASA's Nancy Grace Roman Space Telescope.

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

Biographies of the other authors are not available.