Novel approaches for direct exoplanet imaging: Theory, simulations and experiments

Por, E.H.

Author: Por, E.H.
Title: Novel approaches for direct exoplanet imaging: Theory, simulations and experiments
Issue date: 2020-12-11

Chapter 8

First laboratory demonstration of the phase-apodized-pupil Lyot coronagraph with integrated high-order wavefront sensor

Adapted from
E. H. Por, A. Potier, P. Baudoz, R. Galicher,
M. A. Kenworthy & C. U. Keller
To be submitted to A&A

Context The next generation of high-contrast imaging instruments on space-based observatories requires sophisticated wavefront sensing and control in addition to a high-performance coronagraph.
Aims We provide a first laboratory demonstration of the phase-apodized-pupil Lyot coronagraph (PAPLC). We show that a single deformable mirror (DM) can serve as the phase-apodizer in monochromatic light. Additionally, we present the integration of a phase-retrieval wavefront sensor to measure high-order wavefront errors simultaneously with coronagraphic images.
Method We installed both a non-reflective and a reflective knife-edge in the focal plane of the *Très Haute Dynamique 2* (THD2) testbed at the Observatoire de Paris. We used electric field conjugation using pairwise DM diversity to minimize light in the dark zone. The light reflected by the focal-plane mask is reimaged with a slight defocus onto the phase-retrieval camera. The resulting image allows us to reconstruct the wavefront using weighted least squares and an empirical interaction matrix.
Results We demonstrate a mean raw contrast of 1.9×10^{-8} in monochromatic light, and 6.7×10^{-8} in 7.5% broadband light for a dark zone between $2\lambda/D$ and $9\lambda/D$ using a coronagraph with an inner working angle of $1.2\lambda/D$. Furthermore, we demonstrated open-loop reconstruction of the wavefront with an integrated phase-retrieval wavefront sensor. The reconstruction error was 30pm per mode for the first 32 Zernike modes for small wavefront aberrations, demonstrating a performance within 3× the fundamental photon-noise limit.
Conclusions These laboratory tests confirm our earlier simulated results in Por (2020) and pave the way for an optically simple approach to broadband high-contrast imaging from space that also features unprecedented wavefront sensing capabilities.

8.1 Introduction

In the last few decades, we have started to unravel the mystery that has captivated humanity since antiquity: is there extraterrestrial life? We have indirectly detected many rocky exoplanets (Borucki et al., 2011), and have started to directly detect Jupiter-sized planets with the latest generation of extreme adaptive optics (ExAO) systems, such as VLT/SPHERE (Beuzit et al., 2019), Gemini/GPI (Macintosh et al., 2008), Clay/MagAO-X (Close et al., 2012; Males et al., 2014), and Subaru/SCExAO (Jovanovic et al., 2015). Technology developments for future space-based observatories with dedicated high-contrast imaging (HCI) instruments such as Roman/CGI (Spergel et al., 2013) and LUVOIR/ECLIPS (Pueyo et al., 2019) are underway.

A typical HCI instrument consists of a coronagraph, which minimizes the stellar light, and a wavefront control system, which corrects for static and dynamic disturbances in the optical system. These disturbances can be categorized as either low-order or high-order aberrations, based on the number of cycles across the pupil. Low-order aberrations result from large-scale flexure and vibrations across the telescope mirrors, caused by thermal and mechanical load, and the relative movement of the optics in the telescope and the instrument. Low-order aberrations typically evolve rapidly, and therefore are controlled by a wavefront control system using the telemetry from a low-order wavefront sensor measuring the light rejected by the coronagraph (Shi et al., 2017, 2018). Additionally, the coronagraph can be designed to be robust to low-order aberrations, which relaxes the constraints of the wavefront control system (eg. N'Diaye et al., 2015; Ruane et al., 2017).

High-order aberrations are caused by polishing errors of the primary mirror, and, for a segmented telescope, by segment misalignments. For a monolithic mirror in space, high-order aberrations tend to evolve slowly, and are therefore easily controlled by the speckle control system, which uses telemetry from the science camera. However, for segmented telescopes vibrations are expected to dynamically misalign segments (Coyle et al., 2019), requiring active control of high-order aberrations in a similar way to low-order aberrations. However, coronagraphs cannot reject high-order aberrations as it would reduce the throughput for the planet. Additionally, most coronagraphs do not permit high-order wavefront sensing without impacting the science image.

The phase-apodized-pupil Lyot coronagraph (PAPLC; Por 2020) is an exception. It uses a standard Lyot-style optical layout, shown schematically

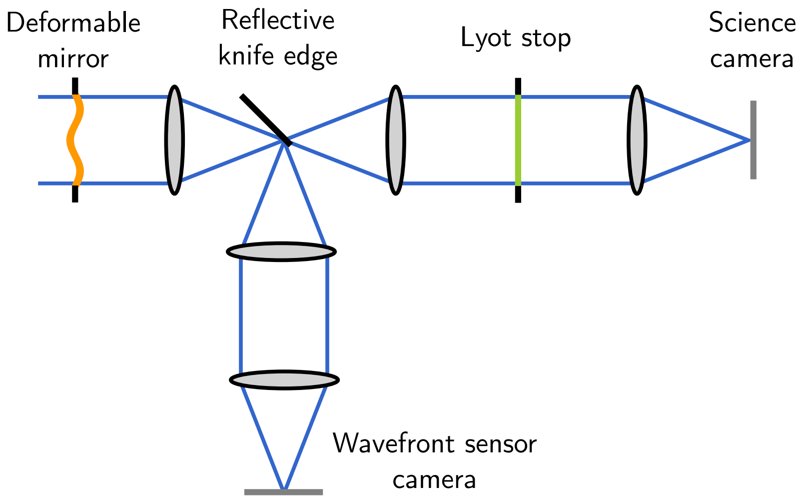

Figure 8.1: The schematic layout for the PAPLC with a deformable mirror and a phase-retrieval wavefront sensor. The PAPLC as presented by Por (2020) uses an achromatic phase-only mask instead of the deformable mirror.

in Figure 8.1. It consists of a pupil-plane phase-only apodizer that creates a one-sided dark zone. A knife-edge focal-plane mask blocks the bright side of the resulting PSF, and the dark side is further filtered using a Lyot-stop mask. Por (2020) demonstrates that PAPLC designs for the segmented LUVOIR-A telescope can achieve inner working angles (IWAs) as small as $2.4\lambda/D$ at a 10^{-10} contrast and high coronagraphic throughput. Por (2020) also eluded to the possibility of adding high-order wavefront sensing by using the light reflected by the focal-plane mask.

In this paper, we present the first laboratory demonstration of the PAPLC and its integrated high-order wavefront sensing capabilities on the *Très Haute Dynamique 2* (THD2) testbed at the Observatoire de Paris. Sect. 8.2 investigates the theoretical limitations of our implementation using optical simulations. Sect. 8.3 presents simulations of the integrated high-order wavefront sensor and its noise characteristics. Sect. 8.4 presents our experimental results for both the coronagraph and wavefront sensor. We summarize our conclusions in Sect. 8.5.

8.2 PAPLC with deformable mirror

The PAPLC, as presented in Por (2020), uses an achromatic phase-only apodizer to modify the PSF falling onto a knife-edge focal-plane mask. The PSF itself is offset by a grating superimposed on the phase-only apodizer, leading to an offset of the PSF relative to the knife-edge that grows linearly with wavelength. This results in a coronagraph that is inherently achromatic, as the changes in size of the PSF as a function of wavelength are compensated by the chromatic offset of the PSF relative to the knife-edge focal-plane mask.

In this paper, we instead use a DM as the apodizer, which introduces a chromatic phase pattern, and we physically offset the focal-plane mask instead of adding tip-tilt on the DM. Additionally, the DM has a much smaller number of degrees of freedom compared to the freeform phase mask of the PAPLC. We investigate two consequences of these changes in this section. In Sect. 8.2.1 we present the monochromatic performance of this version of the PAPLC; Sect. 8.2.2 presents the broadband performance.

8.2.1 Monochromatic performance

A globally-optimal design algorithm is presented in Por (2020) to obtain the phase pattern for the phase apodizer in a PAPLC. We now want to apply this phase pattern on the DM. A simple projection of the optimal phase pattern onto the surface of the DM is insufficient to reach the required contrast due to fitting errors on the DM surface. Additionally, directly enforcing the phase pattern to be a linear combination of DM modes in the globally-optimal design algorithm presented in Por (2020) is not feasible because the optimization is performed on complex electric field amplitudes rather than phase only. Hence, the constraints for enforcing the phase pattern to be separable into DM modes is non-linear, which increases the optimization time by many orders of magnitude.

Instead, we use electric field conjugation (EFC) as presented by Give'On (2009) to iteratively optimize the DM voltages to obtain the phase pattern for the PAPLC, starting from a flat DM. This may not produce the DM pattern with the highest-possible throughput as it does not explicitly maximize throughput, but rather minimizes stellar intensity in the region of interest. In practice, however, the EFC routine yields DM solutions with a sufficiently high throughput, as long as non-aggressive PAPLC design parameters are used. In these cases the DM solution will be driven by the dark zone contrast rather than the throughput of the coronagraph, and the EFC routine and a global optimizer will find approximately the same solution. A further advantage is that this approach can readily be used on a high-contrast imaging testbed, as most testbeds already include an implementation for performing EFC in combination with some electric field estimation method.

Similar to Mazoyer et al. (2018), the stroke on DM solutions for the PAPLC can be quite large compared to conventional stroke minimization procedures (Pueyo et al., 2009). This means that the local Jacobian can be quite different from the initial Jacobian. We follow Mazoyer et al. (2018) and recompute the Jacobian every ten iterations to ensure that the correction is computed using the local Jacobian of the non-linear system. In practice, simulations without periodic recalculations of the Jacobian tend to converge slower and to a worse contrast level compared to DM solutions that are obtained with periodic recalculations. Again, these differences are small when non-aggressive PAPLC solutions are used.

In Fig. 8.2 we show the result of an EFC simulation with an optical system that mimics the laboratory tests presented in Sect. 8.4. We used an offset of $0.9\lambda_0/D$ between the PSF center and the knife edge of the focal-plane mask; the region of interest starts at $1.8\lambda_0/D$. The other important parameters are shown in Table 8.1. Even with the reduced number of degrees of freedom compared to the apodizer, the DM is capable of producing a pattern that suppresses the stellar light in the dark zone. The total peak-to-valley stroke of the DM surface is $\sim$ 250nm, the vast majority of which is used on actuators that are blocked by the Lyot stop. For actuators inside the back-projected Lyot stop, the peak-to-valley stroke is $\sim$ 40nm, giving rise to the high throughput of the PAPLC. The coronagraphic core throughput is 82% of the core throughput without coronagraph. The achieved contrast is 4.0×10^{-10} averaged over the dark zone from $1.8\lambda_0/D$ until $9\lambda_0/D$.

8.2.2 Broadband performance

We show the wavelength dependence of the PAPLC with a DM by calculating the EFC solution for monochromatic light and then calculating the coronagraphic image for a number of wavelengths across a 7.5% band centered on this wavelength. Fig. 8.3 shows these simulated coronagraphic images. Contrast curves for simulated images with different bandwidths are shown in Fig. 8.4.

As expected, the raw contrast decreases with increasing distance form the center wavelength, which implies that this version of the PAPLC is fairly

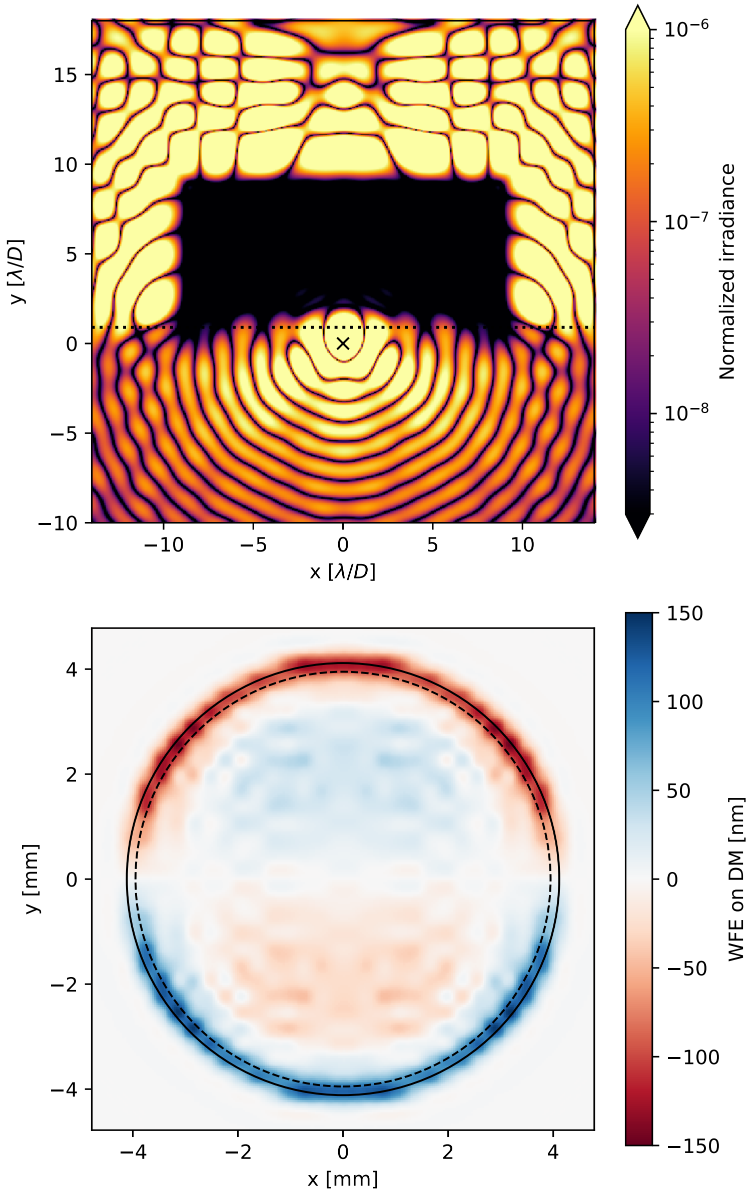

Figure 8.2: The simulated performance of a PAPLC with a DM for monochromatic light, along with the DM surface map. The knife-edge position is indicated by the dotted line. In the DM surface map, the solid and dashed lines indicate the pupil and Lyot stops, respectively.

Parameter	Value
Knife-edge offset	$0.9\lambda_0/D$
Inner working angle	$1.8\lambda_0/D$
Outer working angle	$9\lambda_0/D$
Central wavelength	785nm
Pupil diameter	8.23mm
Lyot mask diameter	7.9mm
DM actuator pitch	300μm
Size of the control region	$\sim 27 \times 27\lambda_0/D$
Number of actuators inside the pupil	591
EFC gain factor	0.9
Jacobian recalculation	Every ten iterations
Pixels across pupil	220pix
Oversampling of the PSF at the knife edge	32pix/λ/D

Table 8.1: The simulation parameters listed here were chosen to be analogous to, but not a complete end-to-end simulation of, the THD2 bench.

chromatic. However, the DM solution is not optimized for this bandwidth, which makes this the worst-case scenario. Broadband EFC can, in principle, be used to find a better broadband DM solution, but this was not attempted in our lab experiments on the THD2, and therefore is not performed in this simulation.

8.3 Simultaneous high-order wavefront sensing

8.3.1 Principle

Phase retrieval, in the field of astronomy, is a procedure in which the phase in the exit pupil is determined using one or more intensity images (Gonsalves, 1982). It relies heavily on the mathematical relationship between the phase pattern and the intensity in the measured images, and inverts this non-linear model to recover the unknown phase pattern (eg. Fienup, 1982; Gerchberg, 1972). Here, we will only use the light rejected by the PAPLC by making the knife edge in the focal plane reflective and reimaging the reflected light. Furthermore we will restrict ourselves to a linearized approach, which is appropriate given the small wavefront errors expected in a space-based system, and will only use a single measurement plane (Gonsalves, 2001; Meimon et al., 2010; Polo et al., 2013; Smith et al., 2013).

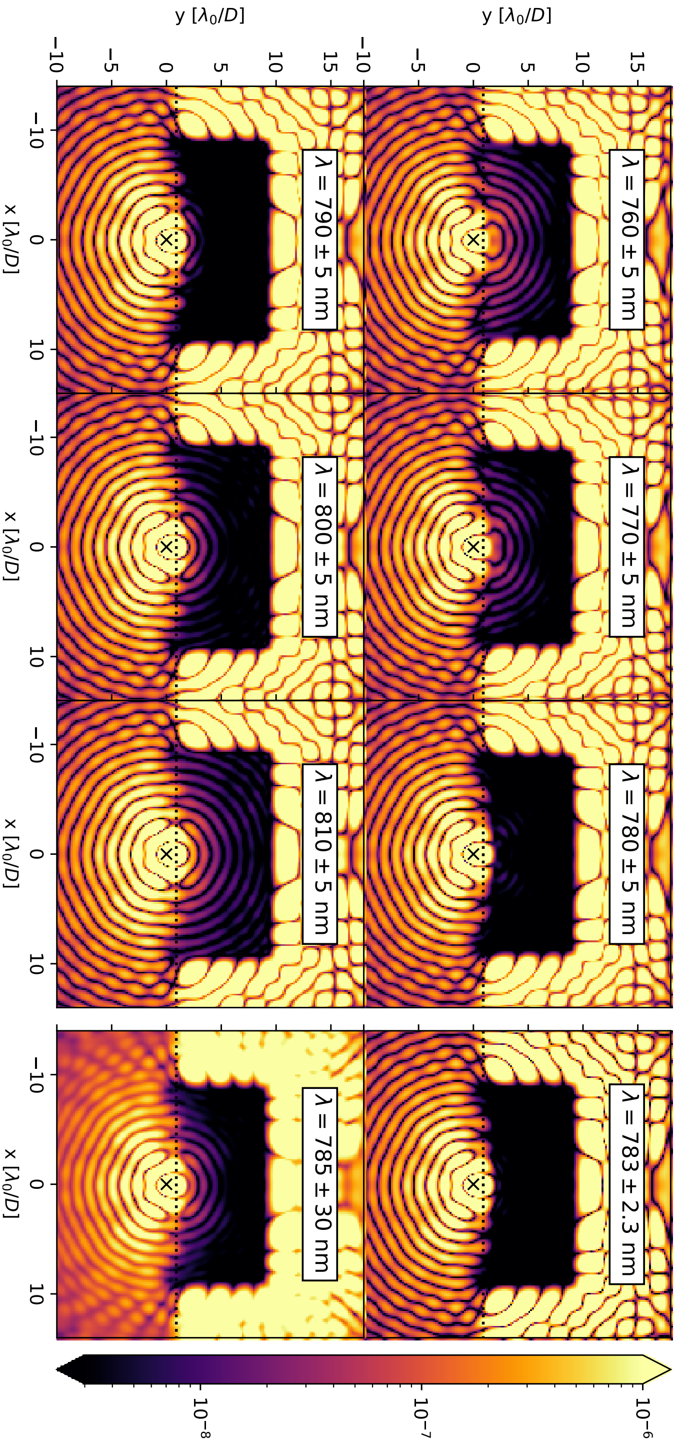

Figure 8.3: The simulated performance of a PAPLC with a single DM in broadband light.

8

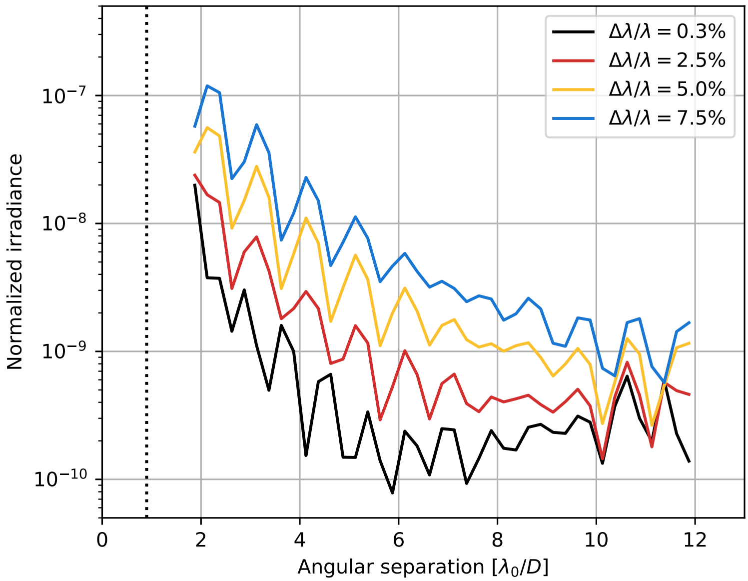

Figure 8.4: Normalized, radial irradiance profiles of simulated images with varying spectral bandwidths.

One well-known problem with single-image phase retrieval is the sign-ambiguity for even modes (Gonsalves, 2001). For odd modes, such as tip, tilt and coma, the focal-plane image will look different depending on the sign of the aberration. Therefore, even from a single image it is immediately clear what sign the aberration has, so we can uniquely reconstruct the phase. However, for even aberrations things are not that easy. When an even aberration, such as defocus, is present, we cannot determine the sign of this defocus since the focal-plane image will look the same regardless of the sign of the aberration. There are two solutions to this problem.

8

1. *Use an asymmetric telescope pupil*, either by blocking part of the pupil with an opaque mask, or by using the natural asymmetry of the telescope pupil itself (Bos et al., 2019; Martinache, 2013). While this is an extremely simple modification, the partial blocking of the telescope pupil, which increases the sensitivity to even modes, simultaneously reduces the number of photons at the camera, reducing the overall efficiency of the wavefront sensor.

2. *Adding a defocus on the phase-retrieval camera.* This adds the necessary diversity for reconstructing the even modes without any loss in

photon flux (eg. Tokovinin & Heathcote, 2006). However, its application in a coronagraphic instrument is limited due to the simultaneous constraint of a high contrast in the dark zone. Therefore, the defocus must be applied sequentially to coronagraphic observations (Paul et al., 2013), or require reflection off a fold mirror close to the focal-plane mask of the coronagraph (Brady et al., 2018). Both methods reduce the duty cycle of the whole system, but can serve as a calibration method for other wavefront sensors, or initial calibration of the coronagraphic system.

Here we use the light reflected by the focal-plane mask of the coronagraph itself. This has the advantage of relaxing the constraint between contrast and wavefront sensitivity to even modes, by allowing an arbitrary amount of defocus on the reimaging arm for the reflected light while leaving the coronagraphic arm unaffected. This is similar to the current approach for low-order wavefront sensing for Roman/CGI (Shi et al., 2017, 2018), where a Zernike phase-dimple is imprinted onto the light reflected off the coronagraphic focal-plane mask. As this dimple only affects the reflected beam, it has no influence on the transmitted, coronagraphic beam, allowing for simultaneous low-order wavefront sensing and coronagraphic observations. In the next section, we will perform numerical simulations showing the trade-off between sensitivity and the amount of defocus on the phase-retrieval camera.

There is an important difference between our PAPLC and the low-order Zernike wavefront sensor mentioned above: the focal-plane of the PAPLC has an infinite extent, albeit in a single direction, instead of a finite extent as is the case for the focal-plane mask for the hybrid Lyot and Shaped Pupil Coronagraphs in the Roman Space Telescope. This means that the reflected light of the PAPLC also contains information on high-order wavefront aberrations. A Zernike wavefront sensor implementation, similar to the one for the Roman Space Telescope, can also be used for the PAPLC and would most likely perform better than a phase-retrieval-based solution. However, it requires a more sophisticated focal-plane mask implementation, making it more complicated, and such an implementation will be left for future work. For the proof-of-concept high-order simultaneous coronagraphic wavefront sensing presented in this paper, we opted to use the simpler approach of phase retrieval at the cost of a slightly-worse theoretical photon-limited sensitivity.

8.3.2 Empirical modal response and reconstruction

Phase retrieval requires an accurate optical model of the system, which is inverted for wavefront reconstruction. In the small-aberration regime, the optical system can be well approximated using a linearized version of the original model (Gonsalves, 2001). This linearized model can easily be inverted to retrieve the phase aberration. Essentially, the image is modelled as the linear deviation from a reference image a function of of the aberration coefficients. The reference image can be taken either from the original model or from the testbed, at a point when a high contrast has been achieved. The former case is generally limited by the accuracy of the model, while the latter requires another wavefront control system to achieve this high contrast by itself before acquisition of the reference image. In this paper, we used the latter approach.

The response of the system to a specific aberration mode is called a *response function*. The response function for each mode can be obtained in two ways:

1. *From a model.* We can create an accurate model of the optical system, and calibrate all relevant parameters of this model with measurements from the physical optical system. In this case, the accuracy of the reconstruction is generally limited by the accuracy of the model.

2. *From the testbed.* We can simply poke each of the modes on the deformable mirror in the optical setup and empirically measure the response on the phase-retrieval camera. In this case, the reconstruction is typically limited either by the noise floor of the phase-retrieval camera, or by small changes in the optical setup between calibration time and run time.

Here we will use the second method due to its simplicity and accuracy. The reconstruction is performed using a weighted least-squares fit on the phase-retrieval images. Each pixel in the image is weighted using a pixel-by-pixel noise model, which is a combination of the expected read noise and photon noise.

Figure 8.5 shows an example of a phase retrieval using the PAPLC described in Sect. 8.2. Again, simulation parameters are chosen to be comparable to the laboratory demonstration described in Sect. 8.4. A defocus of 1 radian RMS was applied on the reimaged PSF reflected off the focal-plane mask. We used a photon count of 1.8×10^8 per image in the input light beam. The PSF sampling was matched to that of the camera in the

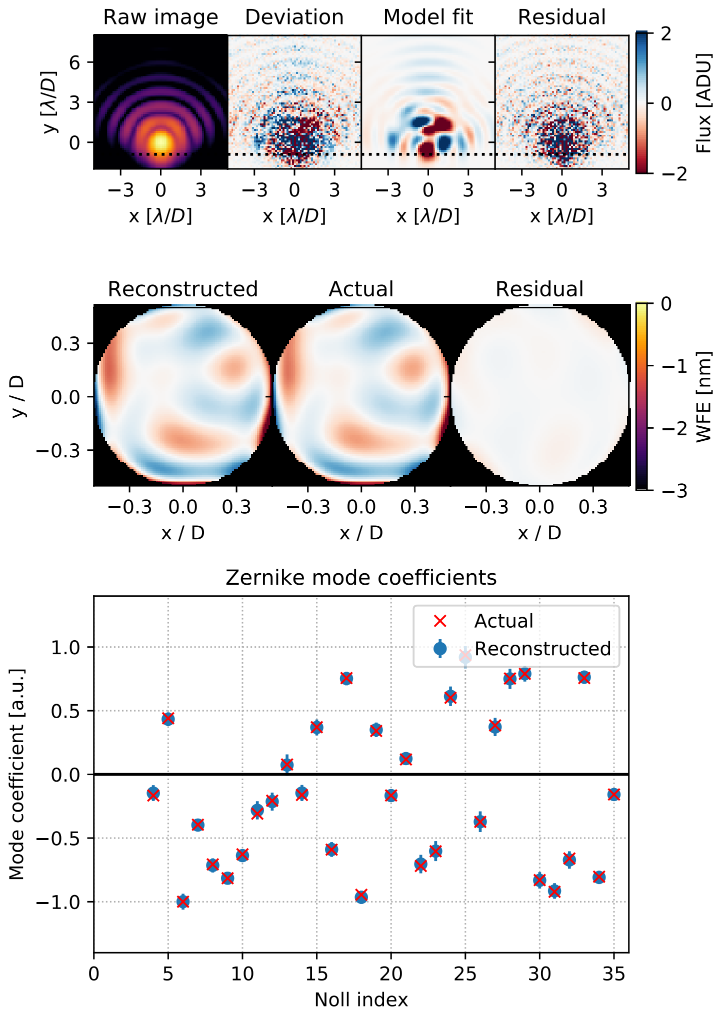

Figure 8.5: An example simulated reconstruction for a phase-retrieval wavefront sensor for a PAPLC coronagraph. The top row shows the reduction of a single frame in the simulated data set with, from left to right, the raw image on a logarithmic scale, the deviation from a reference image, the model fit, and the fit residuals. The middle row depicts the reconstructed wavefront, the actual wavefront, and the residuals of the reconstruction. The bottom plot shows the average reconstructed and actual Zernike mode coefficients. The error bar indicates the 1σ reconstruction error for each frame. The added aberration is 543pm RMS, which is reconstructed with 56pm RMS.

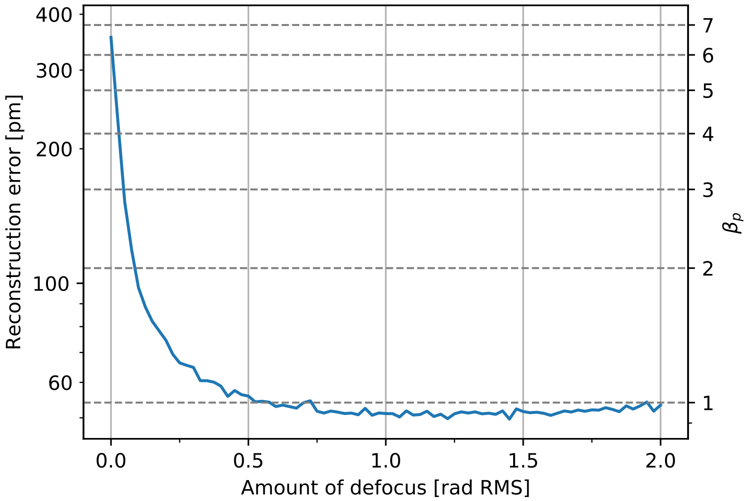

Figure 8.6: The simulated photon noise sensitivity for the phase retrieval. A random aberration of 543pm was applied. When no defocus is applied, even modes do not have a linear response and can therefore not be reconstructed. This yields a reconstruction error of $\sim \frac{1}{\sqrt{2}}$ times the original aberration.

THD2 at 12 px/(λ_0/D). We obtain an empirical interaction matrix with noiseless images. The model fits the simulated image well with the residual containing only photon noise. The reconstructed phase aberration matches the actual aberration perfectly, only deviating by stochastic noise with no observable bias. The same can be seen in the Zernike mode decomposition.

8.3.3 Sensitivity to photon noise

To choose the optimal amount of defocus to apply in our laboratory experiments, we varied the defocus in our simulations and retrieved the reconstruction error due to photon noise. For each amount of defocus, the same random aberration of 543pm RMS was applied. No other noise was present in the simulation. The results are shown in Fig. 8.6. We can see that without any defocus on the reimaged PSF, we retrieve a reconstruction accuracy of 384pm RMS. In this case, even modes cannot be estimated as they have no linear response. The tiny regularization applied during the inversion of the model yields a zero coefficient for these modes. Odd modes do have a linear response and are estimated well, which yields a reconstruction error of $\sim 1/\sqrt{2}$ of the applied aberration. When more defocus is applied, the reconstruction error gradually drops to the fundamental photon noise limit

$\beta_p = 1$ at around 1 to 1.5 radians RMS of added defocus.

8.4 Laboratory demonstration

8.4.1 The THD2 bench

We performed our laboratory demonstration on the THD2 bench that has already been described extensively in the literature (Baudoz et al., 2018a,b; Galicher et al., 2020; Patru et al., 2018; Potier et al., 2018; Singh et al., 2019). A schematic layout of the bench including our additions for the phase-retrieval wavefront sensor is shown in Fig. 8.7. The THD2 contains several light sources; in this work, we used a laser diode (783 ± 2.3nm) and the supercontinuum source with spectral filters with a ~ 10nm spectral bandwidth and center wavelengths of $760, 770, 780, 790, 800$ and 810nm. The flux and spectrum that enters the THD2 bench are continuously monitored for calibration purposes.

A single-mode fiber transports the light to the main THD2 bench. The light is collimated, reflects off several deformable mirrors (DMs) and ends up at the knife edge in the focal plane. Light transmitted by this mask is collimated, filtered by a Lyot-stop mask, and focused onto the science camera. The phase apodization is implemented using DM3, which is conjugated to the pupil. DM1, which is 269mm away from the pupil, was left in its flat state. DM2 was replaced with a flat mirror.

We performed experiments with two different knife edges, one opaque and one reflective. Photos of both focal-plane masks are shown in Fig. 8.8. The opaque knife edge was based on a razor blade, mounted in an optical mount. The reflective knife edge was a standard D-shaped flat mirror with a broadband dielectric coating from Thorlabs (BBD1-E02). We obtained slightly better results with the opaque knife edge, but also spent more time optimizing the speckle control algorithm and alignment with this focal-plane mask. Therefore, the difference is presumably not caused by the quality of the knife edge itself. Another design, based on a coated right-angle knife-edge prism was also considered.

Light reflected off the knife edge is reimaged by a single lens onto the phase-retrieval camera (PR-Cam). The lens, with a focal length of 200mm is 600mm from the knife edge. This yields an unused pupil plane image at 200mm and a focal plane image at 300mm from the lens. The camera (Allied Vision Manta G-319B) was offset by a conservative ~ 4mm from the optimal focus, corresponding to a $\sim 20\%$ reduction in peak flux at the

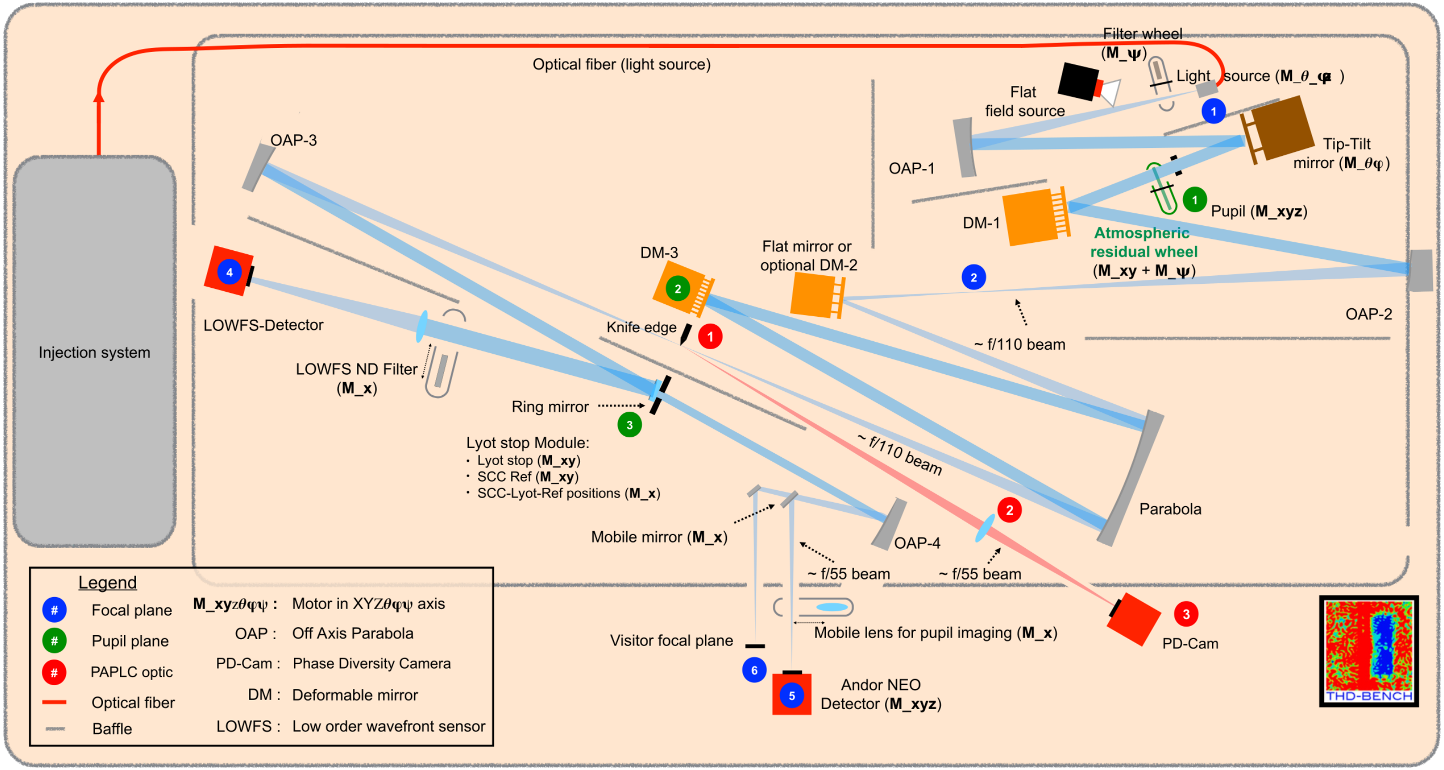

Figure 8.7: Schematic layout of the THD2 bench. Active optics are shown in orange. Numbered green and blue points are pupil and focal planes, respectively. The location of PAPLC optics, including those used by the phase-retrieval wavefront sensor are indicated by numbered red points. The light reflected off the knife edge focal-plane mask is colored red for clarity.

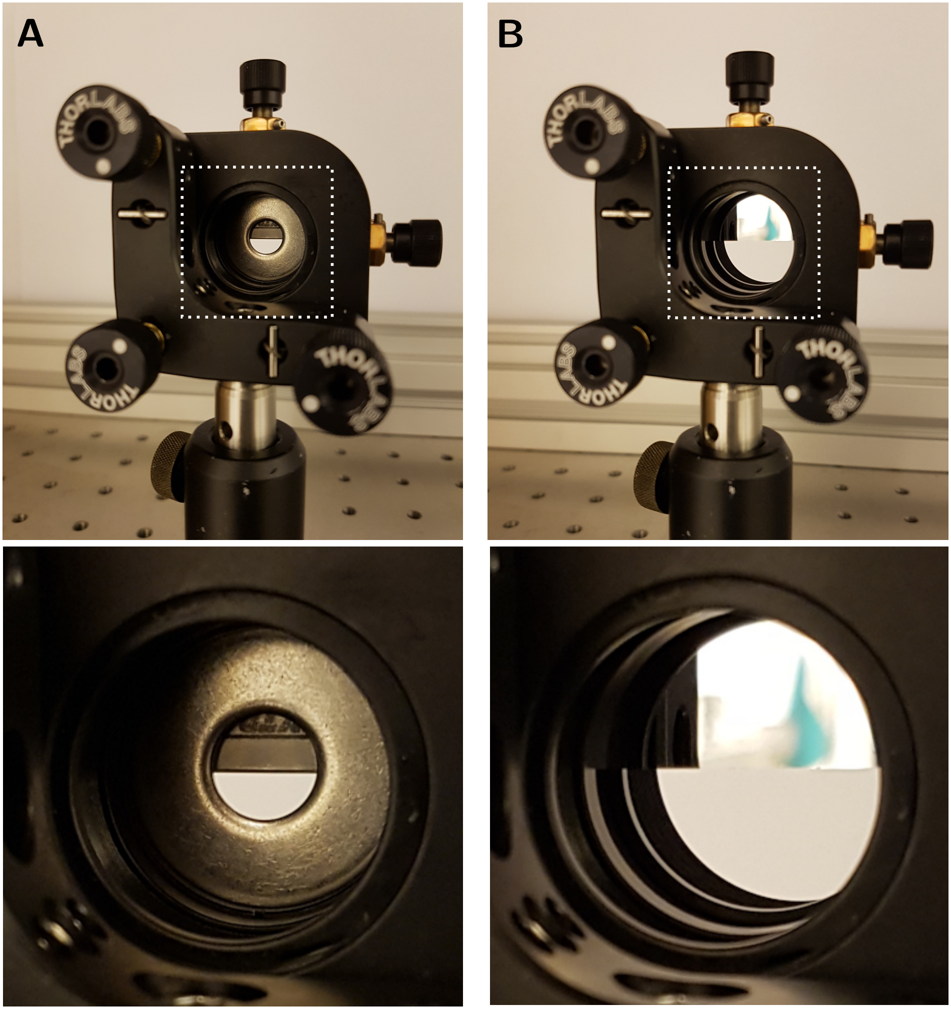

Figure 8.8: Photos of the knife-edge focal-plane masks used for the laboratory demonstration: a) an opaque knife edge based on a razor blade, and b) a reflective knife edge based on a D-shaped flat mirror.

laser diode wavelength and a $\sim$ 60nm rms defocus aberration. Furthermore, several neutral-density filters were placed in front of the camera to reduce the flux and allow the PR-Cam to take unsaturated images at reasonable integration times while simultaneously having sufficient flux on the science camera for speckle control. This, however, had the unintended effect of lengthening the focus, decreasing the amount of defocus to an estimated $\sim$ 0.8mm corresponding to $\sim$ 10nm rms defocus aberration. This decreased the expected wavefront-sensing performance due to photon noise based on simulations. Future experiments will increase the amount of defocus to improve the wavefront-sensing efficiency.

8.4.2 Coronagraphic performance

Monochromatic performance

We performed EFC as described by Potier et al. (2020). For electric field estimation, we used two DM probes, each consisting of a single poked actuator on DM3. Only DM3 was controlled; DM1 was left in its flat position for the whole run. Control voltages were based on a linearized mathematical model of the instrument with a knife-edge focal-plane mask with all DMs in their flat position.

No recalculation of the EFC matrix was performed during the EFC runs. This is not optimal, especially given the large strokes of the PAPLC solution, yielding slow convergence at the end of the speckle control loop as crosstalk between modes is handled by increasing the number of wavefront control iterations, and a higher chance of breaking of the loop. Simulations proved that while better results can be obtained with recomputation of the EFC matrix, it is not strictly necessary to achieve a sufficiently good result at the contrast levels achieved by the THD2 in previous experiments.

The EFC matrix was inverted with Tikhonov regularization. Its regularization parameter was gradually relaxed after convergence to gradually improve the raw contrast. Periodically, the run was stopped and the stellar PSF was recentered by hand using reference spots created by superimposing sinusoidal patterns on DM3. This periodic recentering of the PSF compensates for the tip-tilt induced by the EFC routine. The position of the focal-plane mask was calibrated at the start of the experiment using the flat-field light source, which uniformly illuminates the focal plane, clearly showing the shadow of the knife edge. The focal-plane mask is assumed to be static during the experiment.

Figure 8.9 shows the final monochromatic image in the science focal plane, the corresponding DM surface deviation from the flat position and the normalized radial irradiance curves in the region of interest.

We partially correct for speckles outside of the final dark zone by first performing speckle control on a larger dark zone with an outer working angle of $13\lambda_0/D$. Halfway through the run, the outer working angle was reduced to $9\lambda_0/D$ to improve the contrast in the smaller dark zone from a better starting point. This is evident in the final image, seen as residual lines of speckles just outside of $13\lambda_0/D$.

Additionally, an incoherent ghost is seen, outlined by the dashed ellipse in Fig. 8.9. This ghost is likely an incoherent copy of the coronagraphic PSF positioned towards the bottom left at $(-2,-3)\lambda_0/D$, and we are seeing a copy of the extended edge of speckles at the top of the dark zone at $(0,10)\lambda_0/D$. This ghost is not seen by the electric field sensing using DM probes, confirming its incoherent nature. We performed multiple experiments with different orientations of the focal-plane mask to avoid known ghosts in the THD2.

Additionally we see a line of speckles from the center towards the top, outlined by the straight, dashed lines. While these were coherent, as they were also seen by the DM probe measurements, they were not well controlled by the EFC routine. While similar features were seen in simulations, these were not as persistent, and usually disappeared after convergence. We do not have a clear explanation for their existence.

8

The DM surface map shows the general phase pattern of the apodizer in a PAPLC and the simulations performed in Sect. 8.2.1. However, there are a few peculiarities. The first are the visible vertical lines of actuators that deviate from their immediate surrounding. This can be explained by the assumed "flat" position of the DM. The starting DM flats were based on the last wavefront command of a wavefront control experiment with an FQPM coronagraph (Rouan et al., 2000). As the FQPM coronagraph inherently cancels all speckles on a horizontal and vertical line from the center, these flats still contain any aberrations that produce those speckles. Therefore, we still see an imprint of those aberrations on the DM surface map, as it shows the deviation from the FQPM flat.

Additionally, we see a pronounced structure at the edge of the pupil, the part that is blocked by the Lyot stop. The asymmetry of this edge points towards a small transverse translation of the Lyot stop on the THD2 bench as compared with the EFC model. Further experiments with a better centration of the Lyot stop should confirm this.

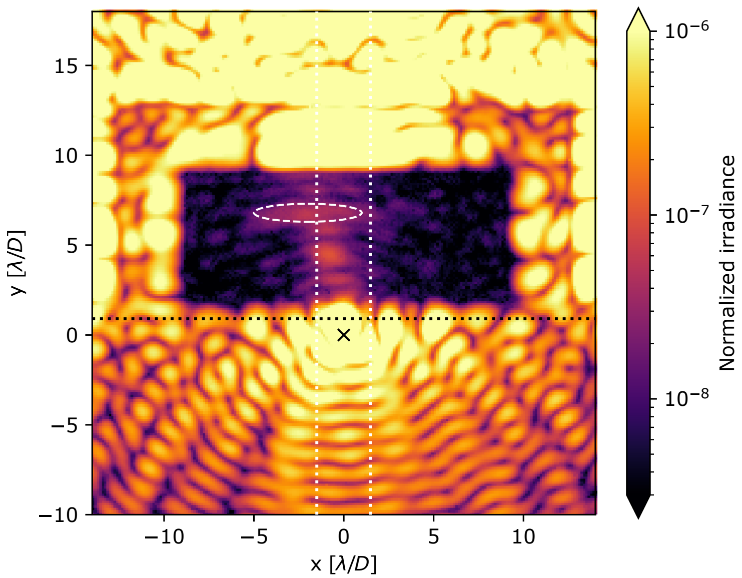

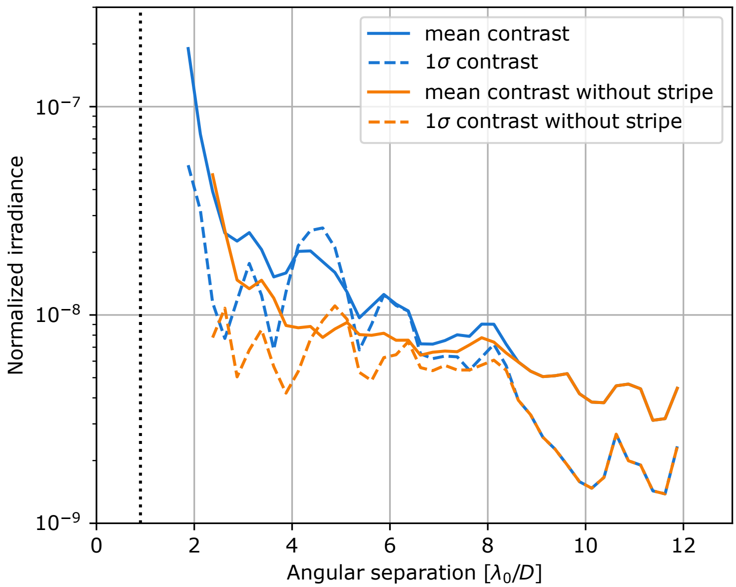

Figure 8.9: The coronagraphic normalized irradiance image and the normalized radial irradiance. The known incoherent ghost, outlined with the dotted white ellipse, has been excluded from the image before taking the radial profile. This curve has been corrected for the coronagraphic throughput of an off-axis source.

8

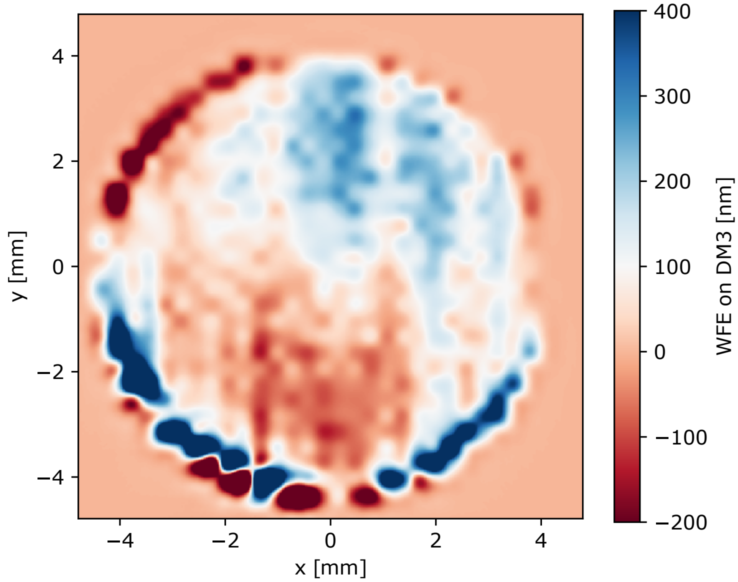

Figure 8.10: The DM3-induced wavefront corresponding to the image in Fig. 8.9 and the normalized radial irradiance profile. The large spikes on the bottom of the DM surface are likely due to a transverse translation of the Lyot-stop mask and pupil with respect to the used model.

8

Broadband performance

Due to the limited photon count in broadband light on the THD2, we cannot perform electric field sensing using a number of wavelengths in the spectral band, as is usually done. Instead, we perform wavefront control using a laser diode in the center of the band, and assess its performance using several broadband filters with the supercontinuum light source, without changing the DM shape.

Figure 8.11 shows the captured images for all filters. Clearly, the further away the wavelength is from the center wavelength of 783nm, the brighter the speckles in the region of interest become. Furthermore, speckles on opposite sides of the center wavelength have a striking similarity, except close to the center of the PSF. This indicates that the electric field is scaling linearly with wavelength, passing through zero at the center wavelength. Therefore, these are likely caused by amplitude aberrations that are corrected by a phase aberration on DM3, which is conjugated to the pupil (Pueyo & Kasdin, 2007). Furthermore, the more speckled and brighter

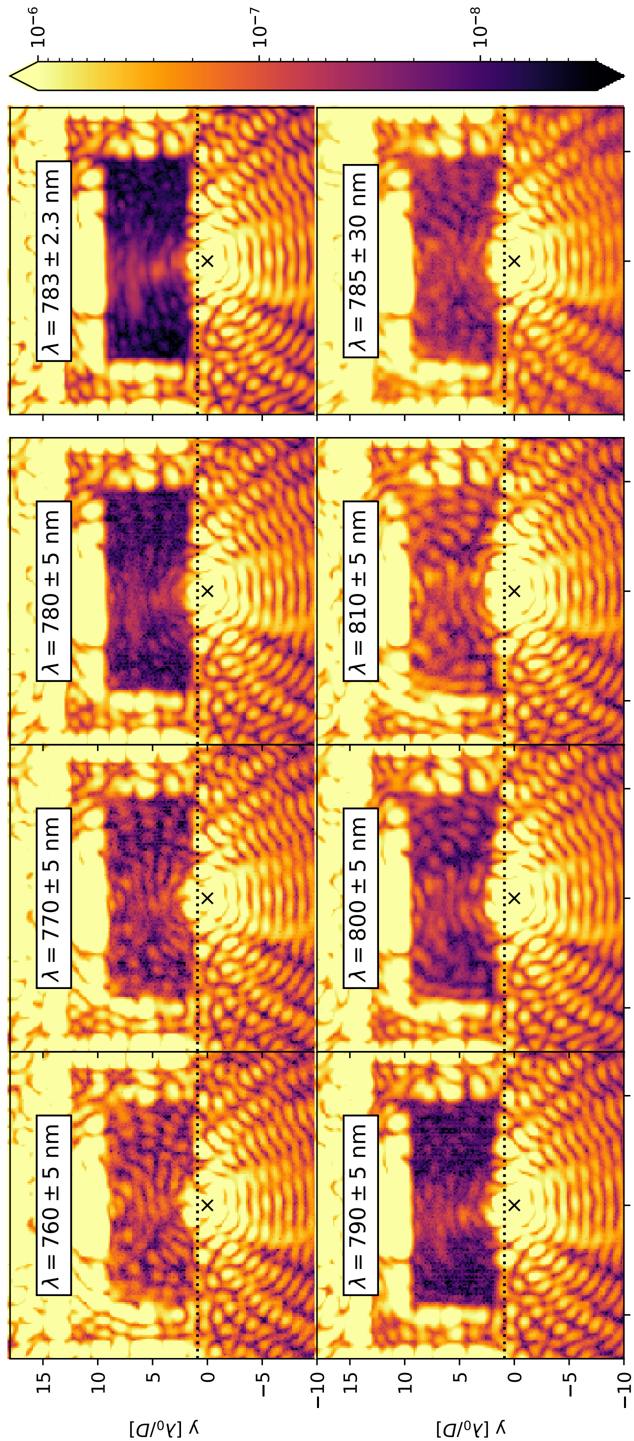

Figure 8.11: Coronagraphic images obtained at several wavelengths and spectral bandwidths. The DM command is optimized for the laser diode (783 ± 2.3nm, upper right). The images for the six narrowband filters between 760 and 810 ± 5nm (first, second and third columns) are combined into a single broadband image (785 ± 30nm, lower right). All images were taken with the same DM shape on January 16th, 2020.

8

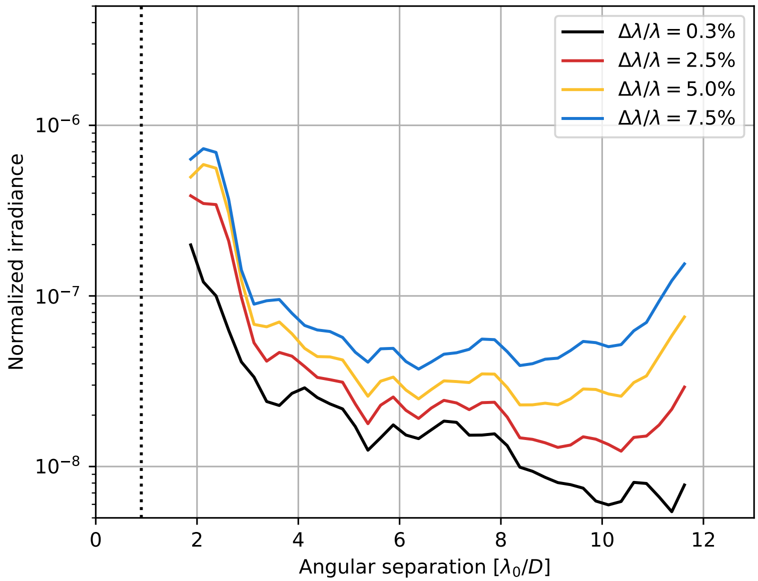

Figure 8.12: Coronagraphic normalized radial irradiance profiles for the images in Fig. 8.11, averaged to represent different spectral bandwidths. These curves are corrected for the coronagraphic transmission for an off-axis source, averaged over the annulli.

8

appearance compared to the broadband features seen in the simulations in Sect. 8.2.2 suggest that these originate from THD2 itself, rather than the coronagraph. Close to the PSF center, we see a different chromatic behaviour, which is likely caused by the chromaticity of the coronagraph itself.

Figure 8.12 shows the normalized irradiance radial profiles for synthetic broadband images based on the images shown in Fig. 8.11. Again, the radial profile excluded the known incoherent ghost, but included the stripe artefact. We can see that for a 7.5% bandpass, the contrast stays below 8×10^{-7} at $2\lambda_0/D$, and below 1×10^{-7} outside $3\lambda_0/D$.

8.4.3 Phase-retrieval wavefront sensor

Data acquisition and filtering

The PR-Cam integrated for 3ms and read out a 200x200px subregion to increase the frame rate. We performed standard dark correction and added up 80 images with a total effective integration time of 240ms. The images were corrected for drifts in the power of the light source by normalizing the images to a constant integrated flux level over the detector. All wavefront-sensing experiments were performed with the narrowband 783nm laser diode. At this wavelength, the PR-Cam produces strongly-oversampled images with around 12 pixels per λ_0/D. We manually limit the maximum spatial frequency in the images to 4 cycles per λ_0/D by applying a circular aperture in the Fourier domain. This filters out most pixel-to-pixel read and photon noise, limiting their ability to contaminate measured modal response functions later on. This filtering process was found to significantly reduce crosstalk between modes, while having little effect on the noise level in the sensed modes.

Modal response

In analogy to the simulation in Sect. 8.3, we want to acquire an empirical response matrix on the THD2. However, after subtracting the reference image from our PR-Cam images, clear tip-tilt drifts are visible, even when keeping the THD2 static. Any tip-tilt drift between the positive and negative poke for a specific mode will be attributed to that mode, leading to crosstalk between tip-tilt and all measured modes. We therefore need to remove any tip-tilt drifts from our images before determining the interaction matrix for higher-order modes.

The DM is not capable of accurately reproducing tip-tilt without introducing other modes as well. This would yield noticeable crosstalk between tip-tilt and higher-order modes. As tip-tilt drift is much stronger than the aberrations we are trying to measure, this is not a sufficiently accurate way to calibrate the tip-tilt modes. We therefore opted to reconstruct the tip and tilt modes from the principal components of the reference data set. As the tip-tilt drift represents the vast majority of dynamic instability in the THD2, accounting for $> 98\%$ of all variance in the reference data set, the first two principal components will each be some linear combination of the tip and tilt response functions. The first two components are shown in Fig. 8.14. We manually fitted the direction of the tip-tilt to these images, which are shown in the inset. By comparing the images to the x and y-derivatives of the reference image, we determined the scaling to physical movement of the PSF. The tip-tilt drift during the reference data sets is shown in Fig. 8.15.

Having removed the tip-tilt drift from these images, we can now calibrate the higher-order modes. We chose to use just the first 32 Zernike modes starting from focus. In principle, nothing limits us from using more

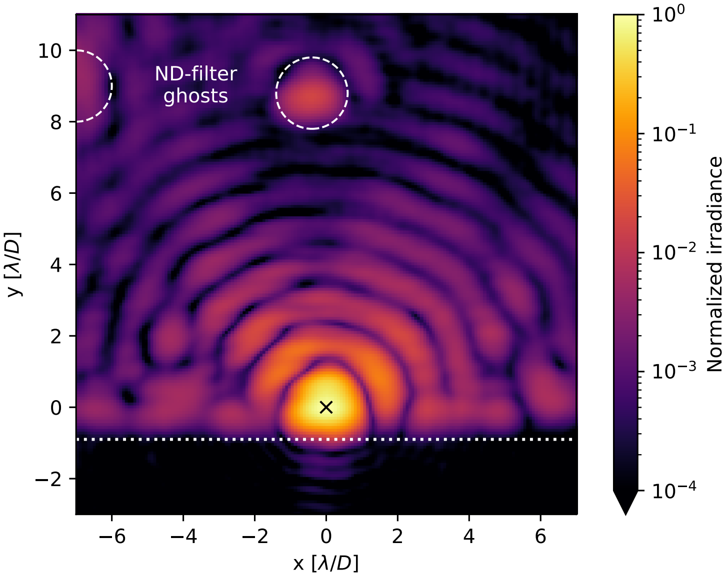

Figure 8.13: The reference image for PR-Cam. Labels identify some prominent ghosts from the ND-filters, which were added to avoid saturation. The cross indicates the location of the laboratory star.

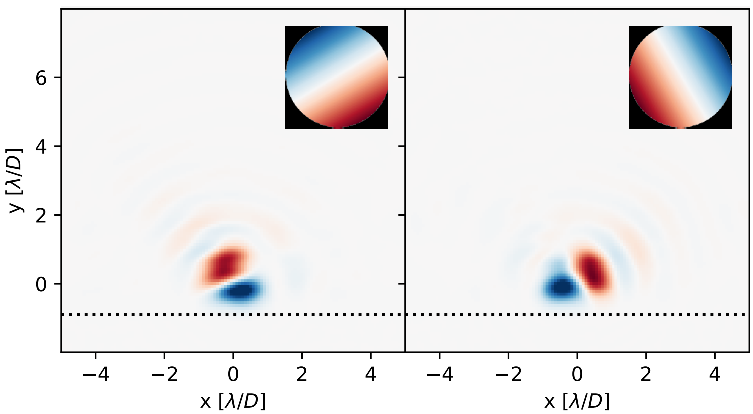

Figure 8.14: The reconstructed tip and tilt modes from the principal component analysis. The x and y tilt for each of these modes were fitted manually.

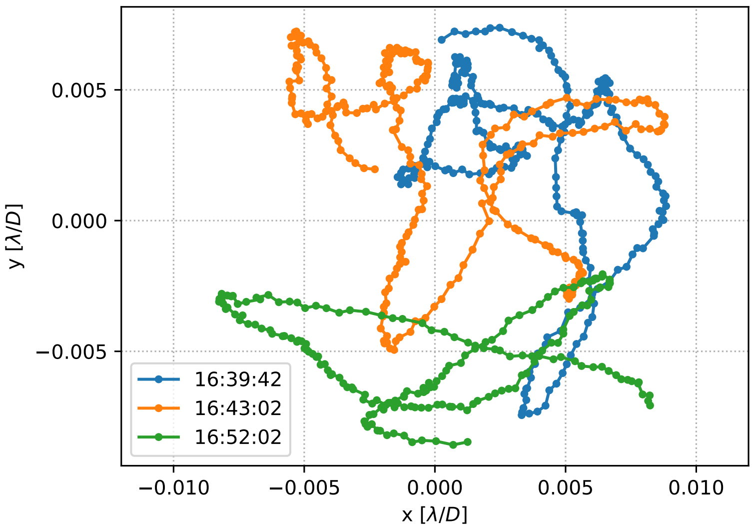

Figure 8.15: The reconstructed tip-tilt drifts during some of the data sets. The duration of each data set was about one minute, with the time between points being $\sim$ 277ms. While implemented for other coronagraphs at the THD2, we did not use any tip-tilt control loop. Future work will use this signal to drive the control loop. The data sets were taken on the 17th of January 2020; the label indicates the local time of the start of the measurement.

modes in the future. The reference diameter of the Zernike modes were slightly oversized by $\sim$ 1 actuator on each side compared to the actual pupil diameter to make sure that we do not suffer from edge effects due to the finite size of the DM influence functions. We ran PR-Cam continuously while sequentially applying our Zernike modes on the DM. Afterwards, we combined images for each mode to retrieve the response function for that mode. On average, we took $\sim$ 10 images per mode. The retrieved response functions are shown in Fig. 8.16 for all 32 modes.

We calibrated the stroke of the DM for all Zernike modes by capturing the coronagraphic science images during the acquisition of the response functions. We fitted a model PAPLC with aberrations to these images, varying the stroke of the Zernike modes, the diameter oversizing of the Zernike modes compared to the actual pupil diameter, the offset of the knife-edge focal-plane mask from the center of the PSF, and a transverse

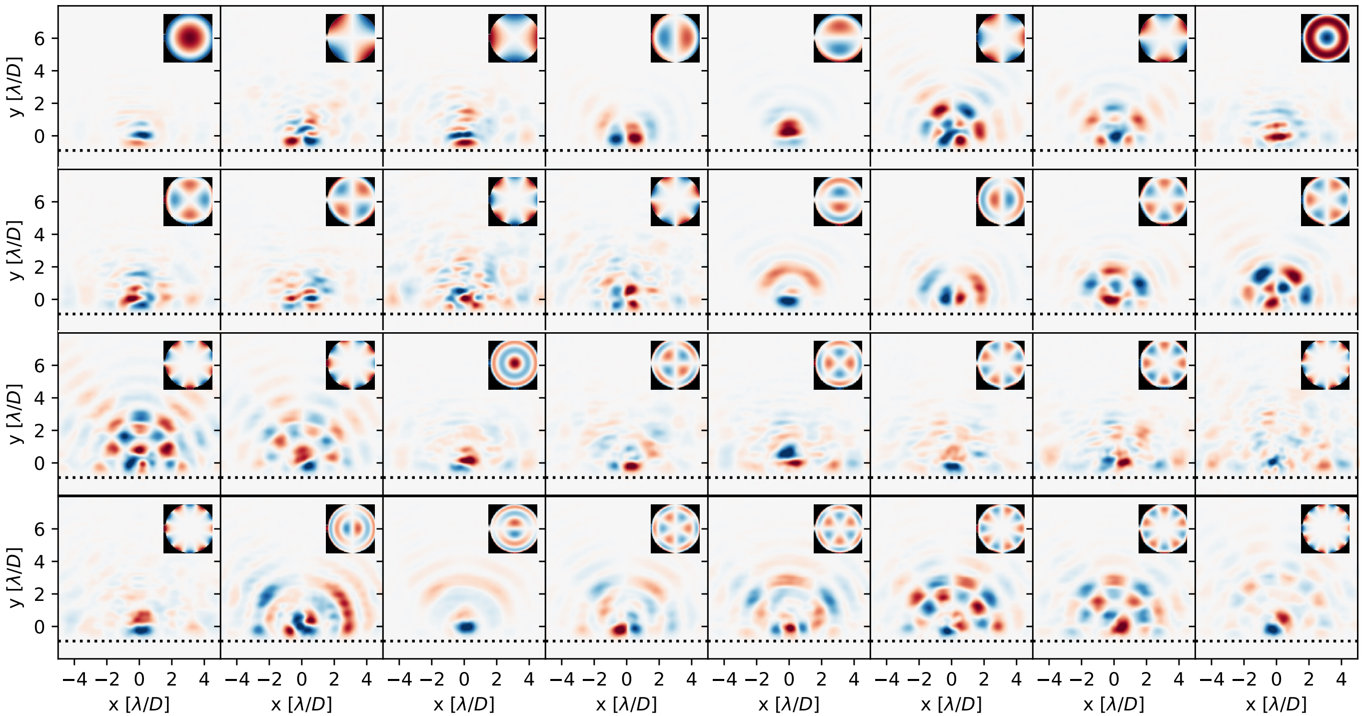

Figure 8.16: The measured PR-Cam response functions for all 32 Zernike modes. The colorscale for each mode is different in order to clearly show the structure in the response function.

	Case A	Case B
Aberration	578 pm rms	5780 pm rms
Systematic error	87 pm rms	521 pm rms
Stochastic error	147 pm rms	147 pm rms
Total error	170 pm rms	541 pm rms

Table 8.2: The approximate error terms. The values for photon noise are independently determined for case A and B.

shift on both the pupil and Lyot-stop masks. The measured coronagraphic images for all 32 modes are shown in Fig. 8.17.

Reconstruction

To show the accuracy of the reconstruction, we applied a known random aberration, composed of the first 32 Zernike modes, on DM3 and performed reconstruction using the image from PR-Cam. We first fit and remove the tip-tilt drift, and subsequently we fit the high-order modes using the retrieved empirical response matrix. This results in 32 coefficients for the Zernike modes for each PR-Cam image. To separate systematic and stochastic errors, we took several hundred images, each composed of 80 subimages. Figure 8.18 shows the analysis of the series of images taken with an aberration with a stroke of 578pm RMS. Figure 8.19 shows the same analysis for the same aberration scaled to a stroke of 5780pm RMS, to differentiate between additive and multiplicative systematic error sources. Additive systematic errors simply add a bias to the estimation, while multiplicative systematic errors scale with the added wavefront aberration.

Stochastic reconstruction errors: photon noise

We use the gain reported by the camera manufacturer ($2.6585e^-$/ADU) to perform a calibration of the number of photons incident on the camera. On average, we record 2.25×10^6 electrons per subframe, corresponding to 1.8×10^8 electrons per frame. To focus on the performance of the wavefront-sensing algorithm rather than the quality of the camera, we assume a quantum efficiency of 100% so that the number of electrons is the same as the number of incoming photons. In this framework, the quantum efficiency is integrated in the throughput factor for the optics. This makes sure that these results can be more easily converted to cameras with better quantum efficiencies.

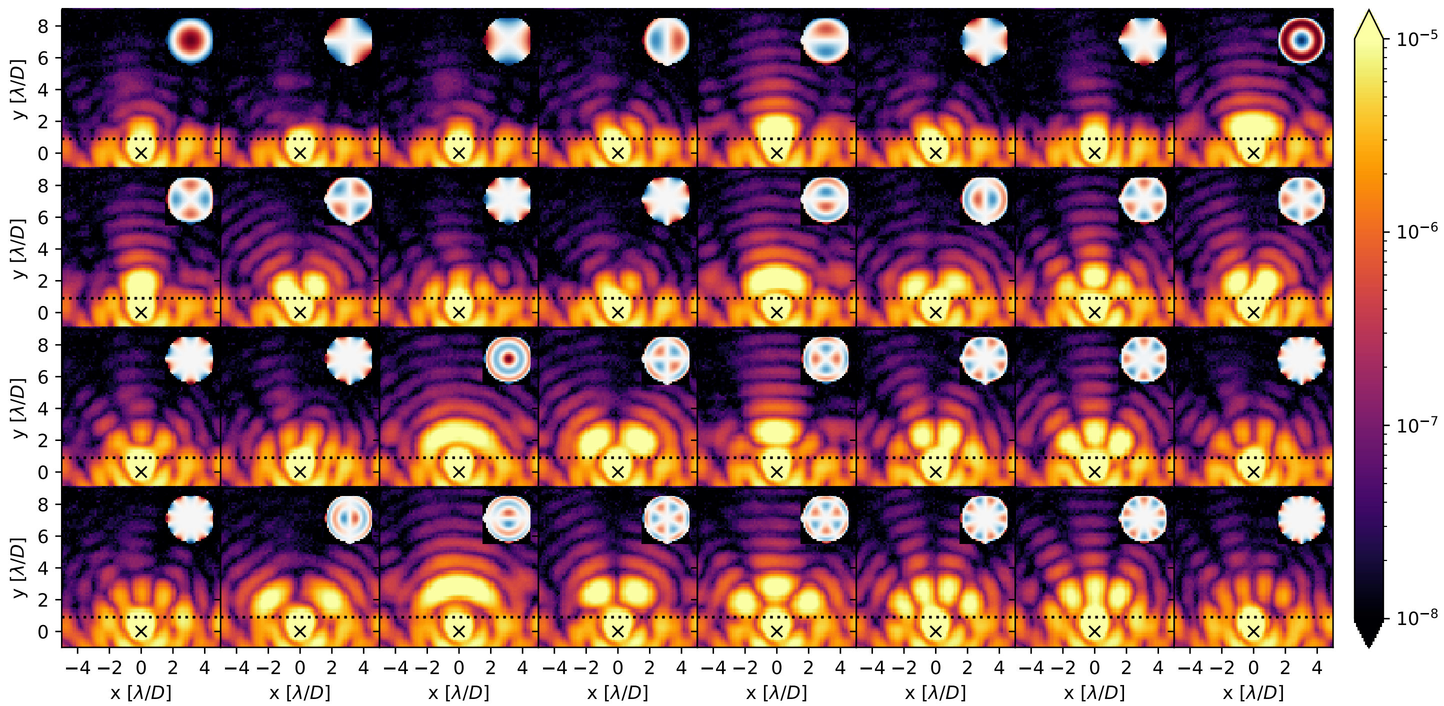

Figure 8.17: The measured coronagraphic images for all 32 Zernike modes. Each Zernike mode was applied with ~ 1.7nm rms. The Zernike modes were applied on DM3 with an $\sim 7\%$ larger diameter to avoid edge effects.

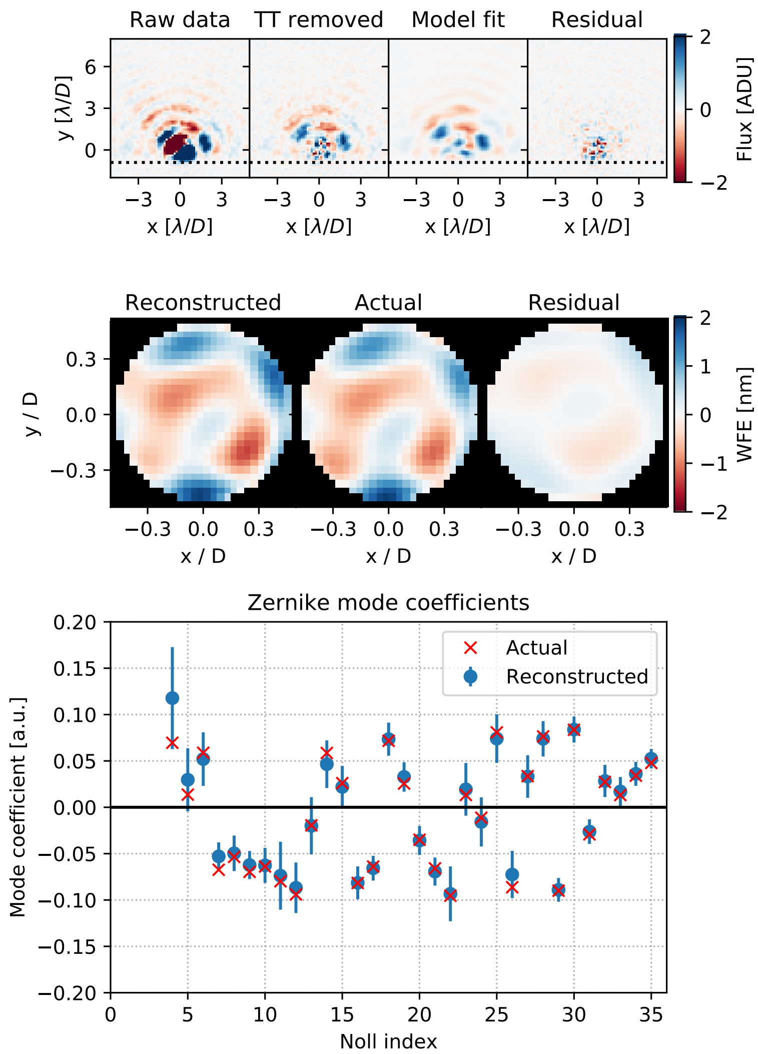

Figure 8.18: Retrieval of the wavefront for a small aberration. The top row shows the reduction of a single frame in the data set with, from left to right, the raw deviation from the reference image, having tip-tilt removed from the image, the fitted model image, and the fitting residuals. The middle row depicts the reconstructed wavefront from a single frame, the actual wavefront that was applied to the DM, and the reconstruction residuals. The bottom plot shows the average reconstructed and actual Zernike mode coefficients. The error bar indicates the 1σ reconstruction error in a single frame. Parameters for this retrieval are listed in Table 8.2 under case A.

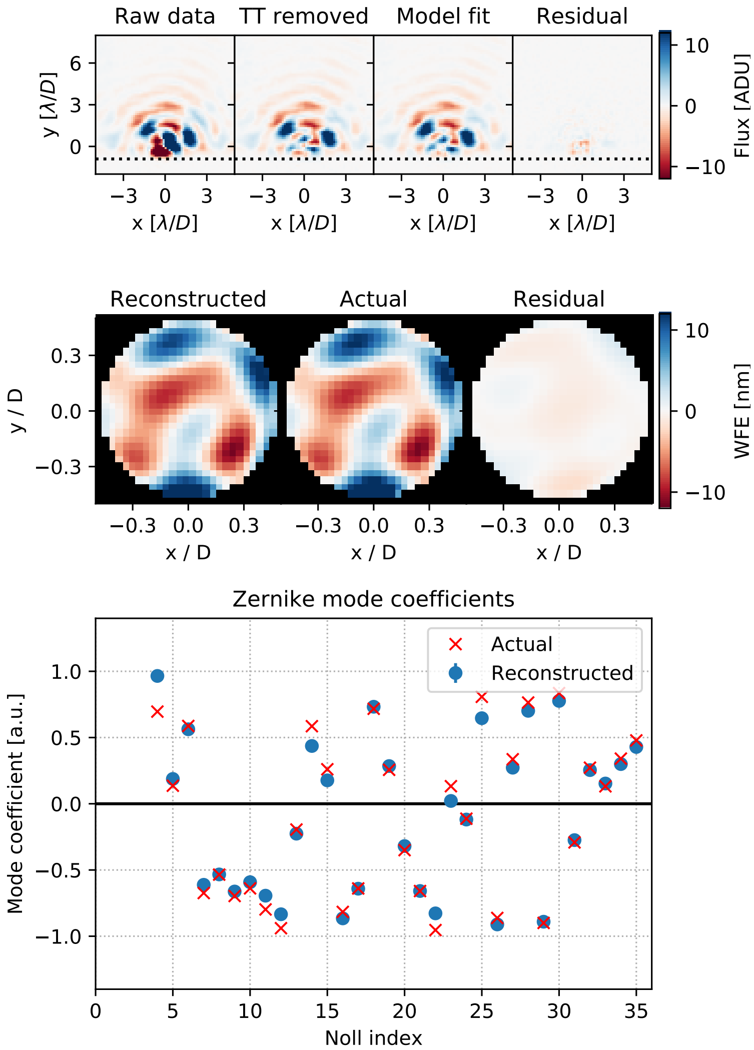

Figure 8.19: The same as Fig. 8.18, but with a larger aberration added. Parameters for this retrieval are listed in Table 8.2 under case B.

The gain yields a fundamental photon noise limit of $\sim$ 54pm for all 32 modes, and $\sim$ 9.5pm for each mode separately. This yields $\beta_p \simeq 3.1$ including systematic errors, and $\beta_p \simeq 2.7$ without systematic errors. This limited performance with respect to the photon-noise-limited performance in Sect. 8.3 is likely the result of the small defocus on the PR-Cam. Future experiments will increase the amount of defocus, yielding both increased wavefront sensor efficiency and enabling longer integration times due to the decreased flux level at the peak of the PSF. Both of these should improve the wavefront-sensing performance.

Systematic reconstruction errors: calibration error

During the acquisition of the empirical response matrix, we are still subject to photon noise. While it is tempting to ignore this noise source since multiple images are taken per mode to increase the signal-to-noise of the recorded response function, and each mode is poked with a larger amplitude to boost its signal, it can still have an effect on the reconstruction accuracy. As the response matrix is kept constant between images for the data set used for reconstruction, the photon noise during calibration will imprint a small crosstalk between modes onto the final reconstruction. In particular even modes, which have a weaker response compared to odd modes, have increased crosstalk with other even modes.

Repeating the simulations performed in Sect. 8.3 with photon noise during the acquisition of the response matrix exhibit similar levels of crosstalk as the observed levels of systematic error in our experiments. We therefore conclude that photon noise in the response matrix dominates other sources of systematic noise.

While this crosstalk impacts open-loop sensing, it should not strongly impact closed-loop performance of this wavefront sensor. After all, unless the crosstalk is extremely strong, a closed-loop experiment will still converge to zero albeit with a slight increase in the necessary number of iterations. In our case, the error caused by crosstalk is $< 10\%$, which has a negligible effect on closed-loop performance. If necessary, calibration can be improved by using stronger modes on the DM during acquisition of the response matrix, something that we opted against due to saturation of the science camera, or by integrating longer during calibration. The former is of course preferred, especially when extending this work to a larger number of modes.

Upon manual inspection of the data, we also observe a small correlation between retrieved tip-tilt and higher-order mode coefficients, indicating crosstalk between tip-tilt and other modes. The strength of this crosstalk is still visible despite the substantial tip-tilt drifts compared to the strength of high-order modes that we are trying to measure. For example, the $\sim 0.005\lambda_0/D$ tip-tilt errors that we observe correspond to $\sim$ 1000pm RMS of wavefront error. We do not have enough data to provide a quantitative estimate for this crosstalk: there is not enough diversity in the tip-tilt signal to retrieve the correlation with sufficient accuracy to separate it from the systematic error due to calibration errors. With closed-loop tip-tilt control, which was turned off during our experiments, we expect this to be negligible.

8.5 Conclusions

This paper presents the first laboratory demonstration of the phase-apodized-pupil Lyot coronagraph with the in-air THD2 testbed. We have shown mean narrowband raw contrasts of 1.9×10^{-8} in a one-sided dark zone between $2\lambda_0/D$ and $9\lambda_0/D$ with an inner working angle of $1.2\lambda_0/D$. In 7.5% broadband light we have shown a mean raw contrast of 6.7×10^{-8} in the same dark zone. This broadband performance is likely the result of testbed limitations. This demonstrates that the original idea of a simple PAPLC is extremely powerful.

Additionally, we have shown a unique capability of the PAPLC: its integrated high-order wavefront sensor that only uses light rejected by the coronagraph. We showed in simulations that this wavefront sensor achieves the fundamental photon noise limit. Furthermore, we demonstrated this wavefront sensor with the PAPLC on the THD2 at a level of 3× the fundamental photon noise limit, retrieving the first 32 Zernike modes with an accuracy of 30pm per mode in a 240ms exposure. We identified a likely cause for the calibration errors, and showed that this is unlikely to influence closed-loop operation.

References

Baudoz, P., Galicher, R., Patru, F., Dupuis, O., & Thijs, S. 2018a, arXiv e-prints, arXiv:1801.06600. https://arxiv.org/abs/1801.06600

Baudoz, P., Galicher, R., Potier, A., et al. 2018b, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 10706, Advances in Optical and Mechanical Technologies for Telescopes and Instrumentation III, 107062O

Beuzit, J. L., Vigan, A., Mouillet, D., et al. 2019, A&A, 631, A155, doi: 10.1051/0004-6361/201935251

Borucki, W. J., Koch, D. G., Basri, G., et al. 2011, ApJ, 736, 19, doi: 10.1088/0004-637X/736/1/19

Bos, S. P., Doelman, D. S., Lozi, J., et al. 2019, A&A, 632, A48, doi: 10.1051/0004-6361/201936062

Brady, G. R., Moriarty, C., Petrone, P., et al. 2018, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 10698, Proc. SPIE, 106986I

Close, L. M., Males, J. R., Kopon, D. A., et al. 2012, in Adaptive Optics Systems III, Vol. 8447, 84470X

Coyle, L. E., Knight, J. S., Pueyo, L., et al. 2019, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 11115, UV/Optical/IR Space Telescopes and Instruments: Innovative Technologies and Concepts IX, 111150R

Fienup, J. R. 1982, Appl. Opt., 21, 2758, doi: 10.1364/AO.21.002758

Galicher, R., Huby, E., Baudoz, P., & Dupuis, O. 2020, A&A, 635, A11, doi: 10.1051/0004-6361/201936971

Gerchberg, R. W. 1972, Optik, 35, 237

Give'On, A. 2009, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 7440, Proc. SPIE, 74400D

Gonsalves, R. A. 1982, Optical Engineering, 21, 829

—. 2001, Optics Letters, 26, 684, doi: 10.1364/OL.26.000684

Jovanovic, N., Martinache, F., Guyon, O., et al. 2015, PASP, 127, 890, doi: 10.1086/682989

Macintosh, B. A., Graham, J. R., Palmer, D. W., et al. 2008, in Proc. SPIE, Vol. 7015, Adaptive Optics Systems, 701518

Males, J. R., Close, L. M., Morzinski, K. M., et al. 2014, ApJ, 786, 32, doi: 10.1088/0004-637X/786/1/32

Martinache, F. 2013, PASP, 125, 422, doi: 10.1086/670670

Mazoyer, J., Pueyo, L., N'Diaye, M., et al. 2018, AJ, 155, 7, doi: 10.3847/1538-3881/aa91cf

Meimon, S., Fusco, T., & Mugnier, L. M. 2010, Optics Letters, 35, 3036, doi: 10.1364/OL.35.003036

N'Diaye, M., Pueyo, L., & Soummer, R. 2015, ApJ, 799, 225, doi: 10.1088/0004-637X/799/2/225

Patru, F., Baudoz, P., Galicher, R., et al. 2018, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 10703, Adaptive Optics Systems VI, 107032L

Paul, B., Sauvage, J. F., & Mugnier, L. M. 2013, A&A, 552, A48, doi: 10.1051/0004-6361/201220940

Polo, A., Pereira, S. F., & Urbach, P. H. 2013, Optics Letters, 38, 812, doi: 10.1364/OL.38.000812

Por, E. H. 2020, ApJ, 888, 127, doi: 10.3847/1538-4357/ab3857

Potier, A., Baudoz, P., Galicher, R., Patru, F., & Thijs, S. 2018, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 10698, Space Telescopes and Instrumentation 2018: Optical, Infrared, and Millimeter Wave, 106986G

Potier, A., Baudoz, P., Galicher, R., Singh, G., & Boccaletti, A. 2020, A&A, 635, A192, doi: 10.1051/0004-6361/201937015

Pueyo, L., & Kasdin, N. J. 2007, ApJ, 666, 609, doi: 10.1086/518884

Pueyo, L., Kay, J., Kasdin, N. J., et al. 2009, Appl. Opt., 48, 6296, doi: 10.1364/AO.48.006296

Pueyo, L., Stark, C., Juanola-Parramon, R., et al. 2019, in Techniques and Instrumentation for Detection of Exoplanets IX, ed. S. B. Shaklan, Vol. 11117, International Society for Optics and Photonics (SPIE), 37 – 65. https://doi.org/10.1117/12.2530722

Rouan, D., Riaud, P., Boccaletti, A., Clénet, Y., & Labeyrie, A. 2000, PASP, 112, 1479

Ruane, G., Mawet, D., Jewell, J., & Shaklan, S. 2017, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 10400, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 104000J

Shi, F., Cady, E., Seo, B.-J., et al. 2017, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 10400, Proc. SPIE, 104000D

Shi, F., Seo, B.-J., Cady, E., et al. 2018, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 10698, Proc. SPIE, 106982O

Singh, G., Galicher, R., Baudoz, P., et al. 2019, A&A, 631, A106, doi: 10.1051/0004-6361/201936206

Smith, C. S., Marinică, R., den Dekker, A. J., et al. 2013, Journal of the Optical Society of America A, 30, 2002, doi: 10.1364/JOSAA.30.002002

Spergel, D., Gehrels, N., Breckinridge, J., et al. 2013, arXiv e-prints, arXiv:1305.5422. https://arxiv.org/abs/1305.5422

Tokovinin, A., & Heathcote, S. 2006, PASP, 118, doi: 10.1086/506972

8