



On-sky validation of image-based adaptive optics wavefront sensor referencing

Nour Skaf, Olivier Guyon, Éric Gendron, Kyohoon Ahn, Arielle
Bertrou-Cantou, Anthony Boccaletti, Jesse Cranney, Thayne Currie, Vincent
Deo, Billy Edwards, et al.

► To cite this version:

Nour Skaf, Olivier Guyon, Éric Gendron, Kyohoon Ahn, Arielle Bertrou-Cantou, et al.. On-sky validation of image-based adaptive optics wavefront sensor referencing. Astronomy and Astrophysics - A&A, 2022, 659, pp.A170. 10.1051/0004-6361/202141514 . hal-03616801

HAL Id: hal-03616801



<https://hal.science/hal-03616801>

Submitted on 22 Mar 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

On-sky validation of image-based adaptive optics wavefront sensor referencing

Nour Skaf^{1,2,3} , Olivier Guyon^{2,4,5,6}, Éric Gendron¹, Kyohoon Ahn², Arielle Bertrou-Cantou¹, Anthony Boccaletti¹ , Jesse Cranney⁷, Thayne Currie^{2,8}, Vincent Deo², Billy Edwards³, Florian Ferreira¹, Damien Gratadour^{1,7}, Julien Lozi², Barnaby Norris⁹, Arnaud Sevin¹, Fabrice Vidal¹, and Sébastien Vievard^{2,4}

¹ LESIA, Observatoire de Paris, Univ. PSL, CNRS, Sorbonne Univ., Univ. de Paris, 5 Pl. Jules Janssen, 92195 Meudon, France
e-mail: nour.skaf@obspm.fr

² National Astronomical Observatory of Japan, Subaru Telescope, 650 North A'ohōkū Place, Hilo, HI 96720, USA

³ Department of Physics and Astronomy, University College London, London, UK

⁴ Astrobiology Center of NINS, 2-21-1 Osawa, Mitaka, Tokyo 181-8588, Japan

⁵ Steward Observatory, University of Arizona, Tucson, AZ 85721, USA

⁶ College of Optical Sciences, University of Arizona, Tucson, AZ 85721, USA

⁷ Research School of Astronomy and Astrophysics, Australian National University, Canberra, ACT 2611, Australia

⁸ NASA-Ames Research Center, Moffett Field, California, USA

⁹ Sydney Institute for Astronomy, School of Physics, University of Sydney, Sydney, NSW 2006, Australia

Received 10 June 2021 / Accepted 23 October 2021

ABSTRACT

Context. Differentiating between a true exoplanet signal and residual speckle noise is a key challenge in high-contrast imaging (HCI). Speckles result from a combination of fast, slow, and static wavefront aberrations introduced by atmospheric turbulence and instrument optics. While wavefront control techniques developed over the last decade have shown promise in minimizing fast atmospheric residuals, slow and static aberrations such as non-common path aberrations (NCPAs) remain a key limiting factor for exoplanet detection. NCPAs are not seen by the wavefront sensor (WFS) of the adaptive optics (AO) loop, hence the difficulty in correcting them. **Aims.** We propose to improve the identification and rejection of slow and static speckles in AO-corrected images. The algorithm known as the Direct Reinforcement Wavefront Heuristic Optimisation (DrWHO) performs a frequent compensation operation on static and quasi-static aberrations (including NCPAs) to boost image contrast. It is applicable to general-purpose AO systems as well as HCI systems.

Methods. By changing the WFS reference at every iteration of the algorithm (a few tens of seconds), DrWHO changes the AO system point of convergence to lead it towards a compensation mechanism for the static and slow aberrations. References are calculated using an iterative lucky-imaging approach, where each iteration updates the WFS reference, ultimately favoring high-quality focal plane images.

Results. We validated this concept through both numerical simulations and on-sky testing on the SCExAO instrument at the 8.2-m Subaru telescope. Simulations show a rapid convergence towards the correction of 82% of the NCPAs. On-sky tests were performed over a 10 min run in the visible (750 nm). We introduced a flux concentration (FC) metric to quantify the point spread function (PSF) quality and measure a 15.7% improvement compared to the pre-DrWHO image.

Conclusions. The DrWHO algorithm is a robust focal-plane wavefront sensing calibration method that has been successfully demonstrated on-sky. It does not rely on a model and does not require wavefront sensor calibration or linearity. It is compatible with different wavefront control methods, and can be further optimized for speed and efficiency. The algorithm is ready to be incorporated in scientific observations, enabling better PSF quality and stability during observations.

Key words. instrumentation: adaptive optics – instrumentation: high angular resolution – methods: numerical

1. Introduction

Over the past 30 years, Adaptive Optics (AO) instrumentation has undergone extensive development in terms of its sophistication and scientific capabilities. Compared to the first astronomical AO experiments with the COME-ON system consisting of a 19-actuator deformable mirror (DM) driven at several hundred Hz (Kern et al. 1989; Rousset et al. 1990), current leading extreme AO (ExAO) systems consists of ≈ 2000 -actuator DMs driven at over 1 kHz speeds (e.g., Macintosh et al. 2014; Beuzit et al. 2019). The high-contrast imaging enabled by these new AO capabilities has led to key scientific breakthroughs, including the first direct image of a planet-mass companion (Chauvin et al. 2004)

and the first imaged system of jovian exoplanets (Marois et al. 2008).

Some of the current generation of ExAO systems, such as the Subaru Coronagraphic Extreme Adaptive Optics (SCExAO) at the Subaru Telescope (Jovanovic et al. 2015) and MagAO-X on the Magellan Clay Telescope (Males et al. 2018), are spurring further achievable levels of performance; for instance, in terms of Strehl ratio (SR), contrast, and sensitivity (e.g., Vigan et al. 2015; Currie et al. 2020). These also serve as technology prototypes for ExAO instruments on 25–40 m extremely large telescopes (Kasper et al. 2021; Fitzgerald et al. 2019; Close et al. 2020).

The ExAO loop corrects wavefront aberrations measured by the WFS, ideally approaching an aberration-free image.

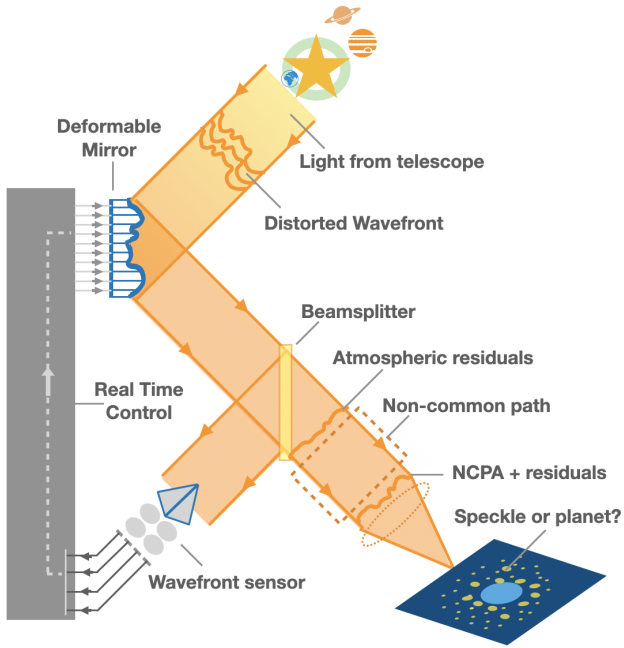


Fig. 1. Schematic view of an AO system, presenting the problem of NCPAs, especially in the case of exoplanet imaging with a coronagraphic image.

However, the ExAO loop convergence point, defined by the WFS reference, may not correspond to the optimal image quality. Reasons for such a discrepancy include several aspects. First, there are the non-common path aberrations (NCPAs) due to optics located after the beam splitting between WFS and science paths, inducing differential aberrations between the science camera and the WFS, as seen on Fig. 1. In a high-contrast imaging system, these aberrations may include coronagraph optics defects. These aberrations can also vary with temperature and mechanical deformation at timescales from minutes to hours, or with the positioning errors of moving optics, making them particularly challenging to calibrate. They are typically of the order of tens of nanometers, enough to lead to static and quasi-static speckles in coronagraphic images (Sauvage et al. 2007; Vigan et al. 2019). These speckles are also present in non-coronagraphic images but these are less concerning because they do not contribute significantly to the SR, as they are usually beneath the camera contrast limit for unsaturated images. In addition, discrepancies may come from WFS calibration errors due, for example, to a non-flat wavefront used to acquire a calibration. Then, the chromaticity for systems where the wavelength of the WFS and the scientific camera are different. Finally, there must be a choice of optimal wavefront state: the goal of the AO system may not always be to produce a flat wavefront. For example, in a HCI, cancellation of static diffraction features (airy diffraction rings, telescope spiders), such as dark hole techniques (Potier et al. 2020), may be aimed at driving the AO loop to a non-flat wavefront state.

The phenomena responsible for the above effects exhibit timescales smaller than the millisecond-level atmospheric turbulence timescale and can be (nearly) static. In this paper, we refer to all discrepancies between WFS reference and optimal image wavefront as NCPAs, noting that in previous publications it usually only refers to differential optical defects.

NCPAs affect all AO systems, but they are particularly acute in high-contrast imaging where small wavefront errors can easily mask faint exoplanet images (cf. Fig. 1). Furthermore, quasi-

static aberrations do not average out to a spatially smooth halo as fast as the atmospheric speckles, leaving structures in the coronagraphic image that resemble planets – unlike the speckles that are generated by atmospheric turbulence alone.

The most reliable approach to measure and correct optical misalignments is to use focal plane images (Gonsalves 1982). In the particular case of NCPAs, these images are ideally generated by the science detector. Several solutions have been suggested for static NCPAs (Paxman & Fienup 1988; Frazin 2018; Vigan et al. 2018; Vievard et al. 2020; Bos et al. 2020; Potier et al. 2020; Codona & Kenworthy 2013), the quasi-static part is shown to be particularly challenging to measure and correct. The evolutionary timescale of the aberrations with respect to the frequency of the correction is a main difficulty in this regard.

In a closed-loop AO system, NCPAs are usually compensated by subtracting a biasing reference signal in the WFS reference corresponding to the aberration to correct. Doing so requires good WFS response knowledge and stability, so that the adequate offset can be applied and maintained. Ideally, this offset should be updated in nearly real-time due to optical and mechanical variations in the instrument, however, in practice it is typically updated on a daily or weekly basis. Finding the adequate offset requires focal plane images to update the reference and this can furthermore be particularly challenging if the WFS exhibits a non-linear response that must be accounted for, as is the case in some high-performance WFSs such as the Pyramid WFS (PyWFS) (Esposito et al. 2020; Deo et al. 2019; Chambouleyron et al. 2020, 2021).

In this paper we present the DrWHO algorithm, a model-free focal plane wavefront sensing approach that is aimed at finding the offset mentioned above to correct slow and static wavefront aberrations, including the NCPAs, on a timescale of a few seconds. We present our discussion in Sect. 2 and we describe the algorithm in Sect. 3. Numerical simulations with the COMPASS software (Ferreira et al. 2018) are presented in Sect. 4. In Sect. 5, we show the results of the algorithm as obtained on-sky with SCEXAO. Section 6 presents the main results of the on-sky validation, in terms of PSF and WFS reference evolution. In Sect. 8, we review the characteristics of the algorithm and describe further implications for use in HCI.

2. Pursuing a good WFS reference

In a closed-loop AO system, each WFS measurement is compared to a WFS reference (usually by subtraction after flux normalization) to isolate residual wavefront errors that must be corrected by the DM. The real-time controller performs the wavefront reconstruction, thus bringing the WFS signal to a point where it gives a flat wavefront, such that the WFS reference defines the convergence point of the AO control loop.

The WFS reference can be measured with an internal light source inside the instrument, before on-sky observations. This can be done as a standalone step or by averaging the WFS signal when taking a response matrix, corresponding to the WFS signal when there is no aberration induced. However, in reality, the “ideal” WFS reference may be different from the internal source WFS reference due to optical illumination discrepancies. The WFS reference is also not static, constantly evolving due to quasi-static aberrations and wavefront chromaticity. The accurate calibration of the WFS is therefore a significant challenge for any AO system. A continuous way to measure the “ideal” reference is therefore required for high accuracy wavefront correction. Finding this “ideal” reference is particularly critical to

ExAO systems, where a slight deviation can introduce slow or static speckles that appear similar to a planet image.

3. Algorithm description

Since the goal is to ensure the AO loop converges to the “ideal” reference, the “actual” reference should be updated continuously. To do so, DrWHO regularly re-estimates this reference, chosen by selecting the WFS images corresponding to the best PSF image quality. The algorithm thus comprises a lucky imaging selection of the PSF leading to the selection of the best WFS images.

The algorithm first requires the focal plane PSFs and WFS frames streams to be synchronized in time, as described in Appendix B. The algorithm furthermore requires that there be a one-to-one mapping provided from PSF to WFS images, which is explored in Sect. 7.

The quality metric (hereafter, “the score”) for the PSF is flexible as far as it optimizes the PSF quality (cf. Sect. 7): this can be the SR, contrast, minimum of intensity, etc. The algorithm proceeds as follows: on a defined timescale (set arbitrarily at 30 s), the algorithm first selects the best PSF frames according to a score metric and a predefined selection fraction (arbitrarily set to 10%). Then, out of those 10% frames, the corresponding WFS raw frames are extracted from the AO telemetry and averaged; the resulting WFS frame replaces the WFS “actual” reference and the algorithm is iterated to continuously optimize the WFS reference.

As the algorithm proceeds, the AO loop point of convergence removes NCPAs and other slow aberrations. Thus, the PSF quality improves, and converges towards a better score. The temporal frequency of the aberrations that DrWHO can correct depends on the timescale of one iteration of the algorithm (for a timescale of 30 s, the algorithm will be able to correct dynamic aberrations up to 0.03 Hz).

We present the procedure step by step of DrWHO with the description in Algorithm 1.

4. Validation via numerical simulations

This section presents the AO simulation set up with the COMPASS simulator (Ferreira et al. 2018). COMPASS is a versatile AO simulator that was designed to meet the need of high-performance for the simulation of AO systems, modeling several kinds of AO features, including several types of WFS, atmospheric simulations, DMs, telescopes, and RTCs. It offers a suitable environment for simulating Single Conjugate Adaptive Optics (SCAO), for 8-m class telescopes and 30-m class telescopes as described in Vidal et al. (2018). Simulations of DrWHO on the COMPASS software were first implemented to explore the algorithm feasibility and potential.

4.1. Extreme AO simulation

The first step was to simulate an ExAO on COMPASS. We looked for simulating an AO close to the SCEXAO instrument, which is a high-order AO system mounted behind Subaru’s facility adaptive optics AO188 (Minowa et al. 2010) that provides a first level of correction (woofer) using a 188-actuator DM and a curvature WFS.

SCEXAO performs high-contrast imaging thanks to a second level of correction (tweeter), using a 2000-actuator DM and a visible PyWFS. We simulated a simplified ExAO tweeter system

Table 1. ExAO numerical simulations parameters.

Numerical simulation configuration	
Telescope	$D = 8$. m diameter No support spiders Central obstruction = 0.12
Turbulence layer	von Kármán, ground layer only $r_0 = 0.16$ at 500 nm $L_0 = 20$ m (to simulate post-woofer residuals) $ v = 10 \text{ m s}^{-1}$
PyWFS	
Subapertures	64×64
Wavelength	Monochromatic, 750 nm
Guide star magnitude	7
Modulation	Circular, $3 \frac{\lambda}{D}$ radius, 24 points
Photon noise only	
Deformable mirror	48 actuators across the diameter
Focal Plane Camera	
Wavelength	Monochromatic, 1.65 μm
Photon noise only	
RTC controller	
Loop rate	2 kHz
Method	Linear modal integrator
Basis	DM Karhunen-Loève (KL) basis ^(a)
Loop gain	0.3
Controlled modes	1505
Modes filtered	300

References. ^(a)Ferreira et al. (2018).

inspired by SCEXAO, without trying to be an exact copy. Table 1 synthesizes the parameters used in the COMPASS simulations.

The simulation is idealized, with no source of noise other than the photon noise on both the WFS and the science detector (observing $\text{mR} = 7$ with 50% efficiency), and does not take into consideration the telescope spiders. Furthermore, we considered a PyWFS with a higher sampling than on the SCEXAO instrument (64 pixels over the PyWFS pupil diameter, versus 50 pixels), hence higher performance than achieved on-sky. The PyWFS is working in the visible at 750 nm, and the detector in the infrared at 1650 nm (cf. Table 1). Finally, the COMPASS simulation allows us to have synchronized exposure times, which is not the case on SCEXAO (cf. Appendix B).

The quality criteria used for quantifying the PSF on COMPASS is the SR.

4.2. DrWHO on COMPASS

We implemented the DrWHO algorithm as described in Algorithm 1. Results are presented for a case with 15 DrWHO iterations, each consisting of 10 000 AO loop iterations (corresponding to 5 s at 2 kHz). The number of AO loop per DrWHO iteration was constrained by simulation time. Henceforth, the focus remains on the low order modes.

The NCPAs, corresponding to a linear combination of the first 12 Karhunen-Loève modes (which is the linear and orthogonal modal basis used for computing the response matrix), with randomized amplitude, were applied on the science path (not visible to the WFS), with a total amplitude of 30 nm rms (cf. Fig. 2). Such amplitude for the NCPAs is slightly greater than what is

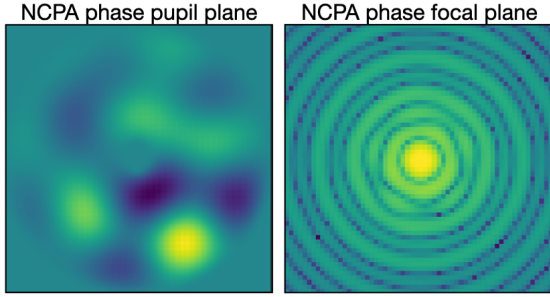


Fig. 2. Image simulation of the NCPAs in the pupil plane (*left*), with an amplitude of 30 nm rms, with the corresponding simulated PSF of size 60×60 pixels, or $12 \lambda/D$, at 750 nm (*right*).

Table 2. Comparison of the long-exposure (LE) Strehl ratio between the following cases of the AO loop: no NCPAs, NCPAs alone, and post-DrWHO with the NCPAs.

	No NCPAs	With NCPAs	With DrWHO
SR LE (%)	87.9	86.8	87.7
σ_{OPD} (nm rms)	94.3	98.8	95.1

Notes. The second line corresponds to the rms of the wavefront for each of those cases, calculated from the SR of the line above with the Maréchal approximation.

expected in reality. In fact, on SCExAO, NCPAs are expected to have a total amplitude of approximately 20 nm. Furthermore, in those simulations, the NCPAs only include differential optical aberrations between the science camera and WFS.

On SCExAO, and on most ExAO systems, the impact of NCPAs over the SR is relatively small, therefore, it is not the best metric to use for quantifying the efficiency of the algorithm, as the AO is dominated by dynamical wavefront residuals. In fact, NCPAs are usually small enough that they do not affect the SR significantly relative to dynamical wavefront residuals, but they add quasi-static speckles in the focal plane that can mimic planets, especially when a coronagraph is used. However, the SR is a good metric to make sure the algorithm does not diverge, in particular through the high-order modes as they could adopt a random walk behavior. Furthermore, the simulation presents a simplified system where a coronagraph is not simulated, hence the SR remains useful.

Table 2 presents the SR for different cases: (1) when the loop is running without NCPAs, the SR reaches 87.9% at $1.65 \mu\text{m}$; (2) when the NCPAs are added and left uncorrected, the SR is reduced to 86.8%. Figure 2 shows the impact of the NCPAs on a PSF free of other aberrations.

4.3. Simulation results

4.3.1. Results in terms of PSF quality

Figure 3 and Table 2 present the result of the SR after the DrWHO run: the final SR after the 15 iterations. DrWHO run reaches 87.7%, with a maximum of 87.9% when there are no NCPAs. As shown in Fig. 3, nearly half of the SR improvement is obtained within the first iteration of DrWHO. Then, the SR keeps slowly improving, until it reaches a plateau. After 15 iterations, the SR reached is almost the same as the SR measured without NCPAs.

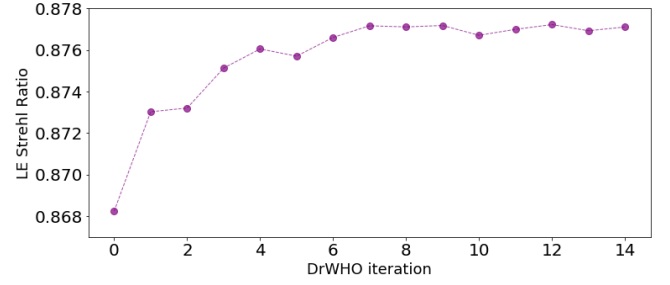


Fig. 3. Evolution of the long exposure (LE) Strehl ratio during the DrWHO run on COMPASS, corresponding to 15 DrWHO iterations of 10 000 AO loop each.

In order to better understand the impact on the optical path difference (OPD), we used the Maréchal approximation

$$\text{SR} = \exp(-\sigma_{\phi}^2) = \exp\left(-\frac{4\pi^2 \sigma_{\text{OPD}}^2}{\lambda^2}\right), \quad (1)$$

with σ_{ϕ} being the variance of the phase and σ_{OPD} being the variance of the OPD.

The corresponding values for σ_{OPD} , calculated from the measured SR, can be found in Table 2.

The wavefront residuals are dominated by dynamical atmospheric residuals, with 94 nm rms in the ideal case without NCPAs. The quadratic difference with and without NCPAs is 30 nm rms of NCPAS, matching the amplitude of the modes added in the simulation. Of these 30 nm, 26.6 nm (89%) are corrected by the DrWHO algorithm.

4.3.2. Results in terms of modes corrected

The 15 references for each iteration of the DrWHO run were projected over the modal basis with which the response matrix was acquired. Figure 4 shows this modal decomposition along with the ideal NCPA correction, where the opposite of the NCPAs was added, to better compare the correction. It appears from the top plot of this figure that, as we previously concluded, the algorithm converges from the first iteration towards a compensation of the NCPAs. The average of all those references (minus the reference 0, which is the initial reference) is neighbouring the NCPAs aberrations. This proves that DrWHO measured them, with the correct amplitude, and compensated for them. The middle and bottom plots of this figure prove that the algorithm does not diverge at the higher order modes, as confirmed by the increase of SR. However, they add a noise contribution to the correction. The quadratic sum of 12 modes from the mean of all the references from the algorithm corresponds to 36 nm rms, which is roughly consistent with the amplitude of the NCPAs.

4.3.3. Discussion

DrWHO on COMPASS improved the image quality in terms of SR, and the NCPAs applied over the PSF, unseen by the WFS, were almost totally corrected. Half of the correction has occurred at the first iteration of the algorithm. The projection of the references from DrWHO is showing that the NCPAs are nearly entirely compensated for, and the higher orders, where we did not apply NCPAs, are not diverging. However close the SR got to the SR without NCPAs, it did not exactly reach the original pre-NCPAs SR. Besides the fact that SRs are defined up to a certain precision only (Roberts et al. 2004), this can be explained by

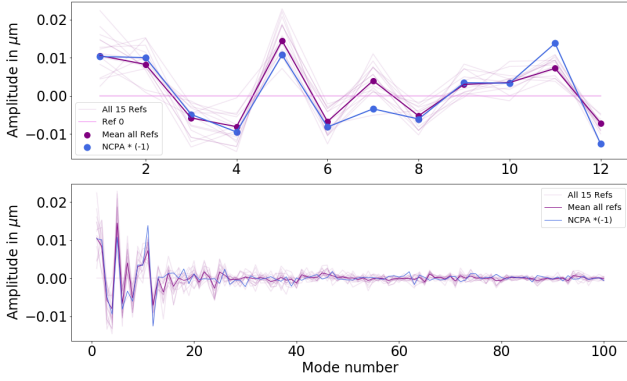


Fig. 4. We projected the references given by DrWHO on the control matrix to quantify the correction applied. *Top*: projection of the references provided by the 15-iteration DrWHO run over the modal basis, with the projection of the opposite of the NCPAs, in order to better visualize the correction. *Bottom*: projection over the first 100 modes.

two effects: first, as seen in Fig. 4, the modes are not perfectly compensated; and second, the modes of the references beyond the ones applied on the NCPAs (12 modes) did have a small – yet not fully negligible – impact on the reference update. Hence, they may have slightly impacted the SR.

5. DrWHO on SCEXAO

The DrWHO algorithm was then adapted and deployed on the SCEXAO instrument, which is located on the infrared Nasmyth platform of the Subaru Telescope, downstream of the AO188 instrument.

Several aspects are different between the COMPASS simulation and the SCEXAO instrument, however the concept remains the same.

The instrument is equipped with a Pyramid WFS operating in the 600–950 nm wavelength range (Lozi et al. 2019). The real time control (RTC) of the AO system is managed by the CACAO software – Compute And Control for Adaptive Optics – (Guyon et al. 2018). The DM has 45 actuators across the pupil, giving a control radius of $22.5 \lambda/D$. We use the CACAO software to interact with the system and implement the algorithm, to communicate between the PyWFS and the DM. There are several scientific modules downstream of SCEXAO. The one used for this paper was one of the two VAMPIRES camera, which is an instrument operating in the visible (Norris et al. 2015), with a pixel scale of 6.1 mas.

On SCEXAO, the PyWFS and the short exposure focal plane camera run at different frequencies, with typical frequencies of 2 kHz for the former, and between 200 Hz to 2 kHz for the latter; hence, there is the need to synchronize the data streams to run DrWHO. This is further described in Appendix B.

5.1. PSF quality Flux criteria

We introduce the Flux Concentration (FC), which is the α -norm defined as

$$FC = \frac{\sum_i (x_i^\alpha)}{(\sum_i x_i)^\alpha}, \quad (2)$$

with x_i being the value of the pixel i , and $\alpha > 1$. If some pixels have a negative value following a dark subtraction due to camera readout noise, then the pixel value is set to zero. This quantity is

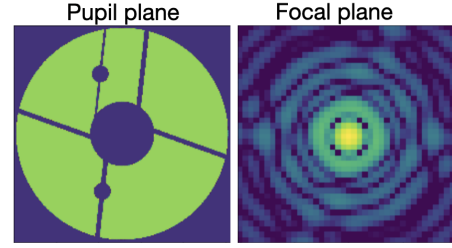


Fig. 5. Image simulation of the pupil plane, including the spiders and a mask for two dead actuators (*left*), and the corresponding simulated perfect SCEXAO PSF (*right*).

independent of absolute flux, since it is normalized by the total flux as seen in the denominator of Eq. (2).

The FC approximately tracks SR, but is less sensitive to noise and sampling effects, as it considers an ensemble of pixel values instead of the PSF peak. If the intensity is concentrated in one pixel, FC equals 1. For the intensity spread over several pixels, $FC < 1$. If the light is equally spread over many pixels, then FC tends to 0. We go on to use $\alpha = 2$, noting that higher values of α favor the brightest pixels.

We chose to compute FC for a subarray of the full image covering the central region of the PSF, in order to focus on the low order modes: wavefront aberrations at low spatial frequencies (except for Tip-Tilt) redistribute the flux in such a way that it spreads across more pixels in the central region of the PSF, and thus the higher the values of the low order aberrations, the lower the FC will get. The FC can be normalized by the FC of a simulated perfect PSF, which we refer to as FC_0 , computed for the pupil of the telescope (cf. Fig. 5), as the FC is a relative number. This new quantity is referred to as the normalized FC (nFC): $nFC = FC/FC_0$.

6. Testing the principle on-sky

6.1. Observation setup

We performed observations on the engineering night of December 8 UT, 2020, during which we ran DrWHO over a period of 10 min, corresponding to 20 iterations of 30 s. Each iteration entails synchronized data cubes resampled to a common 1 kHz frame rate for both VAMPIRES and PyWFS images, running at 500 Hz and 2 kHz respectively. 1200 modes were corrected. The observation conditions at that time were a seeing of $0.5''$, and a wind speed of 4 m s^{-1} . The observed target was the star ξ Leo star, at an airmass of 1.03. The window size of the image for the selection of the quality criteria was of 70×70 pixels, corresponding to 0.43×0.43 arcsec field of view, and to approximately 250 modes.

6.2. Evolution of the PSF quality

Figure 6 shows the PSF just before running DrWHO, followed by the PSF after the first iteration and the PSF after the last iteration of this run. After 20 iterations, the PSF visually looks better and more circular, attesting to a partial correction of the low orders. Furthermore, the outer part of the PSF looks darker compared to the first iteration. The improvement in PSF quality is stronger between the pre-DrWHO PSF and the first iteration PSF, compared to the first iteration PSF and the last one. Hence, it appears that a major part of the correction was applied by the first iteration, as in the simulations.

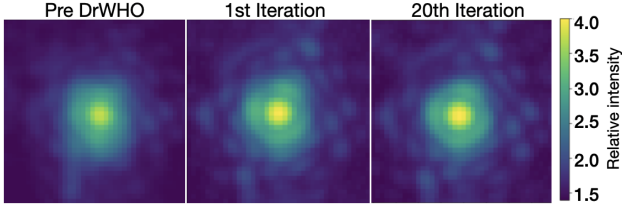


Fig. 6. Evolution of the on-sky PSF before running the algorithm, after the first iteration, and the after last iteration. Each image is 0.25 arcsec (40×40 pixels) across, acquired at $\lambda = 750$ nm, 30 s exposure time (computed by co-addition of 15 000 frames acquired at 500 Hz). Those PSFs should be compared to the ideal PSF in Fig. 5, with matching wavelength, field of view and orientation of Fig. 6 shown in logarithmic brightness scale.

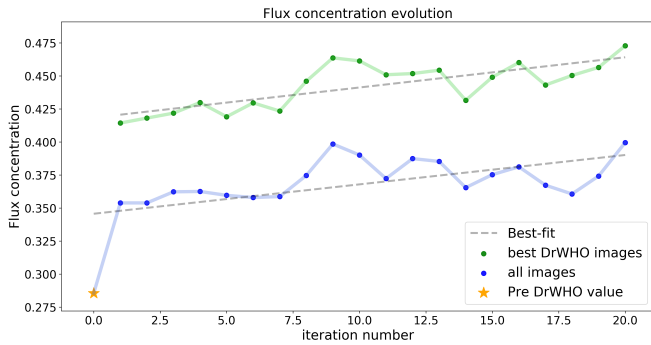


Fig. 7. Evolution of the FC over 20 iterations of DrWHO, corresponding to a period of 10 min, including the nFC of the PSF preceding the run. The blue curve corresponds to the evolution of the nFC of all the PSF averaged over the DrWHO iteration, while the green one corresponds to the evolution of the nFC of the 10% PSF chosen by DrWHO. The best linear fit is presented.

Based on the resulting PSF of size 70×70 pixels for each iteration, we calculated the nFC over a smaller window of 40×40 pixels (0.25 arcsec field of view), focused on the central part of the PSF. The images were background-subtracted and recentered. Figure 7 shows the evolution of these nFC over the 20 DrWHO iterations. We plotted the evolution of the FC for the 10% best PSFs that were selected by the algorithm for the new reference computation, as well as the FC for the average of all the PSFs for each iterations. The difference of those nFC is shown in Fig. 8.

Furthermore, the Modulation Transfer Function (MTF) was calculated for the PSFs before and after the DrWHO run, as shown in Fig. 9. The MTF of the ideal simulated PSF was added to the figure. This figure shows as well that most of the correction is done at the first iteration.

Five main points are worth noting. (1) The major correction is applied at the first iteration, with the nFC jumping from 24.3% to 35%, comparably to simulations; (2) The FC keeps increasing throughout the run, as shown by the best linear fit of the two curves, which is something that was not observed in simulations, where the evolution was smoother and reached a plateau. (3) Considering PSFs averaged over each iteration, nFC increases from 28%, to 35% at the first iteration, to then 40%: the overall improvement is 15.7%. (4) The best DrWHO images present a sharper increase than all the images together: Figure 8 presents the difference of these two nFC over the algorithm run. Despite selecting only the best 10% frames, DrWHO seems to be improving both the selected and unselected frames.

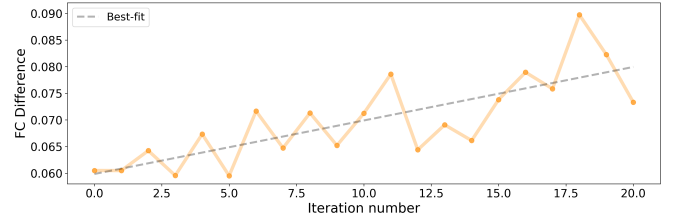


Fig. 8. Difference between the nFC of the mean of the best 10% images, and the nFC of the mean of all the PSFs, over 20 iterations of DrWHO. We calculated the best linear fit to better understand the increase of the difference.

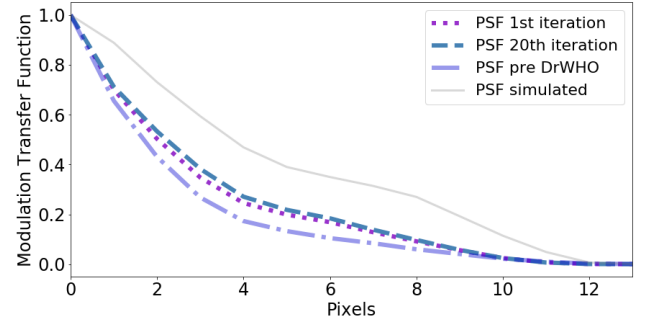


Fig. 9. MTF of the PSF before running DrWHO, of the 1st iteration and the last iteration (cf. Fig. 6), compared to the ideal simulated Vampires PSF.

We note that the algorithm has not yet fully converged and would have performed even better over a longer period of time. (5) Finally, Figs. 7 and 8 show that even if the PSFs are subject to the variations in seeing over a few minutes, they generally get better in terms of the concentration of flux. Hence, DrWHO corrects the NCPAs even if they are not as dominant as the dynamical wavefront residuals from atmospheric turbulence.

6.3. Evolution of the WFS reference

Figure 10 compares the WFS reference prior to the DrWHO run, Ref 0, which is simply the WFS mask showing which pixels are active (i.e. respond to DM pokes) in the WFS image, and the average of the last 3 WFS references applied by DrWHO, during the on-sky 10-min run. The averaging minimizes the noise contribution. The difference between the two references is also shown in Fig. 10. This difference was then converted to wavefront mode coefficients by multiplication by the control matrix. Modal coefficients are shown in Fig. 11 where the first three modes are tip, tilt, and focus and the following modes correspond to orthogonal modes, optimized for SCExAO's AO control law.

The total contribution of all modes is measured to be 116 nm rms. However, as seen on Fig. 11, the high order modes, corresponding to the higher spatial frequencies, seem to be noise more than actual correction. As described in Sect. 7, it is expected that a single metric cannot be sufficient to correct all 1217 modes.

Hence, we focus on the low order modes, from mode 2 to 20. When calculating the total contribution of those modes, we measured a correction of 17.2 nm rms.

Figure 12 (top) shows the first 20 wavefront mode values, excluding tip-tilt. Focus is the dominant contribution to the aberrations, with an amplitude of nearly 14 nm, while the other modes exhibit values below 5 nm.

We furthermore note that the higher modes, corresponding to higher spatial frequencies, show a non-negligible contribution

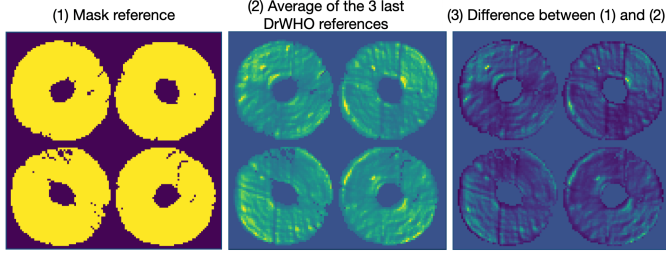


Fig. 10. Comparison of PyWFS references. *Left*: images of the PyWFS reference before the first iteration, referred to as Ref 0, set equal to the WFS mask. *Middle*: average of the last three iterations of DrWHO WFS references. The contrast scale for this image has been modified for better visualization. *Right*: difference between the two images in the left and in the middle.

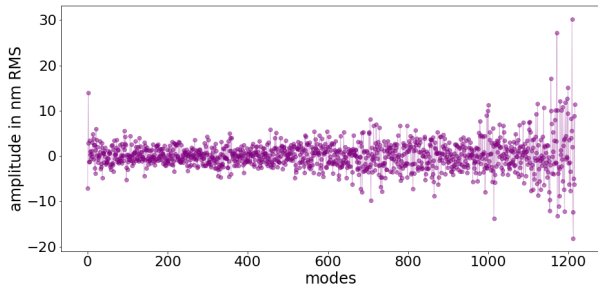


Fig. 11. Modal decomposition of DrWHO correction, computed from the difference between the pre-DrWHO WFS reference and the average of the WFS references for the last three DrWHO iterations.

Table 3. Contribution of the modes of the difference of Fig. 8 in nm rms.

Modes	2 to 1217	2 to 900	2 to 20	Focus
Contribution (nm rms)	116.3	70.3	17.2	13.9

in the overall correction. The way the selection is done is less efficient on high dimensions. DrWHO optimizes the PSF with over a thousand wavefront modes using a single scalar metric, here the FC. The low orders modes are improved, however the high orders follow a random walk. There is a global selection, but that creates noise on the higher order modes, as can be seen in Fig. 11. This is further developed in Sect. 8.

Table 3 presents the contribution of the modes in nm rms, in the following cases: in considering all the modes except tip-tilt (from 2 to 1217); then in removing the higher order modes (from 2 to 900 only) to show what the contribution would be without the higher spatial frequencies; then the low order modes (from 2 to 20), and finally the focus contribution, which is 13.9 nm.

7. Mapping between WFS and PSF

This section further explores the mapping between PSF and WFS images, in both directions, from PSF to WFS images and to WFS images to PSF. We summarize our two main points below.

From WFS to PSF: we expect the PSF to be uniquely defined by the WFS image, as the WFS is designed to unambiguously measure the incoming wavefront phase, and the PSF is the square modulus of the wavefront complex amplitude. We note that the Pyramid WFS raw images also encode the wavefront amplitude, which may also affect the focal plane image. From PSF to WFS

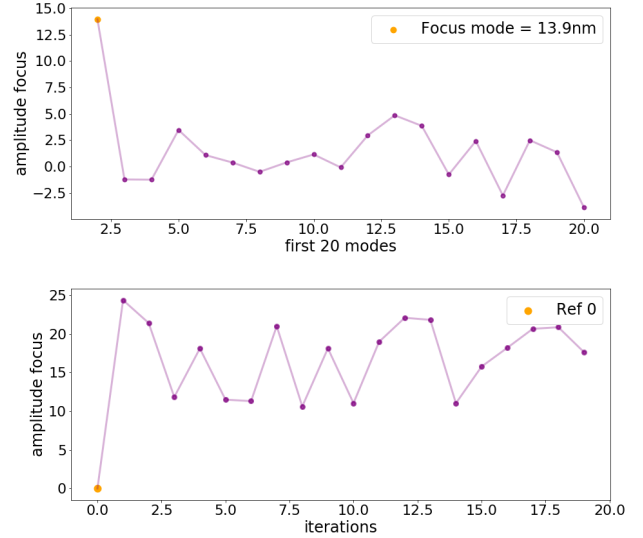


Fig. 12. Modal decomposition of DrWHO correction for the first 20 modes, excluding tip-tilt (*top*). Focus is the major contributor with an amplitude of 13.9 nm rms, which evolution during each iteration of the DrWHO run is shown on the *bottom* figure.

images: in general, the WFS image cannot be unambiguously derived from the focal plane image. For example, the wavefront focus mode exhibits a sign degeneracy, as positive and negative focus values produce the same focal plane image.

This latter aspect is problematic for DrWHO, as the algorithm requires a one-to-one mapping from PSF to WFS images; since we selected WFS images based on PSFs. Specifically, if two different wavefront maps produce the goal PSF image, DrWHO would converge to an average of the two wavefront maps. For example, a defocused PSF goal would be met by positive or negative focus values, but the average wavefront map produced by DrWHO would have both WF solutions cancel each other, yielding a zero = defocus solution that does not match the goal defocused PSF. Thankfully, we expect this degeneracy disappears if the goal image is the aberration-free PSF corresponding to a flat wavefront, as only a flat wavefront can produce this PSF. Hence, as DrWHO tries to optimize the PSF flux concentration, there should be a one-to-one mapping from PSF to WFS images.

To test this point, we analyzed the data cubes of synchronized frames from the observation night of Sect. 6, and used the approach described in Guyon et al. (2021) to identify image pairs that are similar, as measured by the Euclidean distance between images.

The result is shown in Fig. 13. Each point maps to a set of two epochs, from which we extract the corresponding pair of WFS images, and the pair of PSFs. The y axis is the Euclidean distance between the two PSFs, and the x axis is the Euclidean distance between the two WFS images. Points in the lower part of the graph therefore correspond to epoch-pairs for which the focal plane PSFs are similar, while the left part corresponds to epoch pairs for which the WFS images are similar. Because the number of pairs is very high, we adopted the approach of looking at a random sub-set of 100 000 points out of 400 million. The corresponding density plot is shown in Fig. 14.

Horizontal and vertical lines have been added for visualisation help in Fig. 13, both located at the same arbitrary value on each side. The horizontal line stakes out where the pairs of PSF are similar, and likewise for the vertical line and the pairs of

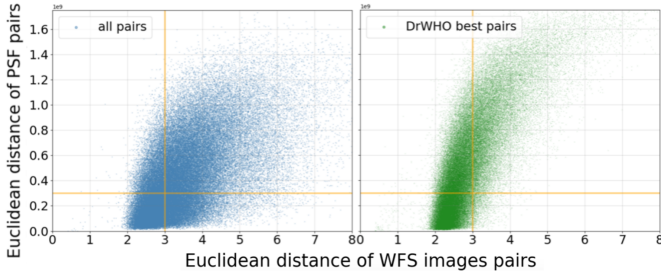


Fig. 13. Each point of that figure corresponds to a pair of two epochs, to which we can associate a pair of images, both in WFS space and PSF space. The left-hand side corresponds to all the pairs, without any selection. The right-hand side corresponds to the 10% best PSF selected by DrWHO.

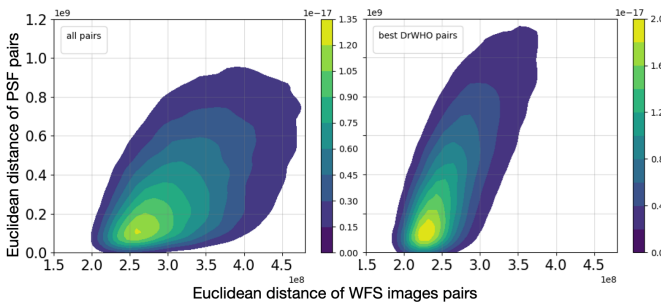


Fig. 14. Density plot corresponding to Fig. 13. As seen with the colorbar, the density plot of the best pairs selected by DrWHO is essentially focused near the origin, where the best PSF are located.

WFS images. To validate the mapping from PSF to WFS images, we consider at the region below the horizontal line. As seen on Figs. 13 and 14 on the left, for pairs of PSFs that are very close to each other (hence the part below the horizontal line) there are several pairs of WFS images possible (see the right part of the vertical line). This confirms our expectation that DrWHO should not work when not selecting for the sharpest PSF images (hence, it would fail when trying to maintain a defocused PSF for phase diversity, for example). However, on the figure on the right, the fraction of data points on the bottom right part of the cross is almost zero. Points are still scattered due to numerous sources of noise, and some points still fall to the right of the vertical line. However, as seen in Fig. 14, the density of points is much higher in the lower left part of the distribution, indicative of a one-to-one mapping from PSF to WFS when selecting for the best PSF. This confirms the validity of DrWHO.

8. Discussion

8.1. The problem of dimensionality

With DrWHO, two competing processes are at play: improvement by removal of dominant WF aberration modes and random walk due to noise in the wavefront sensor and focal plane image. Our data demonstrate that DrWHO is able to improve WF quality over low-order modes, where relatively large errors are concentrated over a small number of modes. At high spatial frequencies, our current implementation is inefficient and likely introduces wavefront noise due to the random walk process. The challenge is ultimately linked to the curse of dimensionality, where optimization in a high dimension space with a single scalar metric is inefficient and unstable: the probability

of all WF modes being simultaneously close to zero (no aberration) becomes vanishingly small, so the DrWHO solution follows a random walk around the optimal solution. To address this limitation, a multi-dimensional optimization metric should be adopted so that the optimization can be performed as parallel independent low-dimension optimizations. A natural implementation would be to separate the focal plane image in spatial zones, each corresponding to a set of wavefront modes according to the Fourier transform relationship between focal plane position and pupil plane spatial frequency.

8.2. DrWHO characteristics

Below are some noticeable characteristics of the DrWHO algorithm that are demonstrated in this paper:

The algorithm is robust: it slowly but surely converges towards a better compensation of static and slow aberrations. As it is based on lucky imaging, it relies on some realisations of residual atmospheric turbulence to reduce wavefront aberrations. DrWHO contributes to getting closer to an absolute sensor, driving the WFS reference to optimize PSF in the focal plane.

It does not rely on any model or on using the DM for probing, making it compatible with other wavefront control techniques. In particular, DrWHO's approach is different from active cancellation speckle because it simply does a statistical selection using the natural dynamical atmospheric residuals for probing, instead of adding artificial probes.

It does not require modification of the AO loop control scheme other than WFS reference updates. It also does not make any linear assumption, so it is not concerned with non-linearities of the system, such as WFS non-linearity.

It is flexible in terms of the parameters, whether it is the score (Strehl, contrast, intensity, etc.), or the selection fraction and the time of the iteration, making it adaptable to different weather conditions and AO systems with different characteristics and scientific goals. However, DrWHO needs a fast focal plane camera as close as possible to the science camera, to more accurately compensate the NCPAs and low aberrations.

The selection criteria for the PSF selection has to optimize the PSF, DrWHO will not work on a convergence point that is not the flat wavefront due to the fact that the mapping from PSF to WFS images cannot be established otherwise.

On-sky tests were run on a bright star, but the algorithm efficiency must be improved to run on fainter targets, and be able to optimize deep contrast for HCI. Such applications to the HCI with coronagraphic images that provide some adaptations of the quality criteria of selection (where it would be exactly the same as what is presented here), but while also adapting the quality criteria of the selection. In this case, instead of selecting PSFs with the maximum FC, the image contrast would be optimized. A first attempt at HCI application was presented in Skaf et al. (2021).

9. Conclusions

We implemented the DrWHO algorithm first using an ExAO simulator, then validating on-sky using the SCExAO instrument at the Subaru Telescope. The results presented in this paper demonstrate its ability to measure and compensate for NCPAs, especially the low order modes, considered to be a limiting factor in the detection and characterisation of exoplanets in high-contrast imaging observations. To quantify the quality of the PSF, we used the Strehl ratio for the simulations, and the Flux Concentration for the SCExAO run. For the simulation, we

observed a nearly perfect compensation of the added NCPAs at the first iteration.

For the on-sky test, we combined the PyWFS measurements and the focal plane images in visible from the VAMPIRES module. When DrWHO was running, the PSF significantly improved (15.7% relative improvement in flux concentration) over 20 iterations of 30 s, for a total of 10 min. The best PSFs have a sharper increase over the run of the algorithm, reaching 48% of the simulated ideal PSF flux concentration. The visible improvement in the image quality was confirmed with the calculation of the MTF before and after the DrWHO run. The data on this run were used to explore the mapping from PSF to WFS images, as a one-to-one mapping is a necessary condition for DrWHO to converge to the goal PSF image. We established the uniqueness of the mapping in the case where the selection metric is optimal for the aberration-free PSF.

This shows that DrWHO is able to improve the wavefront quality arriving on the science camera, and partially correct for the static and quasi-static NCPAs, over a relatively short period of time of a few seconds. We show that the algorithm converges rapidly, with about half of the correction achieved by the first iteration. This has been observed in simulation and on-sky. Furthermore, the correction of the NCPAs is effective even if it is considerably smaller than the dynamic of the atmosphere.

The characteristics of DrWHO are enumerated in Sect. 8. The main strong points are the robustness of the algorithm, its independence from any model, making it compatible with other wavefront control methods, and that it does not rely on any linear assumption, thus it is particularly fit for a PyWFS. In fact, some other techniques for compensating NCPAs mentioned in Sect. 1 usually work independently from the WFS. They do their own correction based on a model, for example, by applying probes on a DM channel that is parallel to the AO loop DM channel. DrWHO combines focal plane wavefront sensing and the WFS. Furthermore, there is a natural extension of this technique in high contrast imaging in the IR, to be more efficient by parallelizing the algorithm in terms of spatial frequencies, which would allow to correct speckles individually and to observe fainter targets.

This paper is the first proof of concept of the algorithm and presents some first results of its PSF quality and stability enhancement. DrWHO shows a wide potential of improvement, which will be presented in future work.

Acknowledgements. Based [in part] on data collected at Subaru Telescope, which is operated by the National Astronomical Observatory of Japan. The development of SCEXAO was supported by the Japan Society for the Promotion of Science (Grant-in-Aid for Research #23340051, #26220704, #23103002, #19H00703 and #19H00695), the Astrobiology Center of the National Institutes of Natural Sciences, Japan, the Mt Cuba Foundation and the director's contingency fund at Subaru Telescope. The authors wish to recognize and acknowledge the very significant cultural role and reverence that the summit of Maunakea has always had within the Hawaiian community. We are most fortunate to have

the opportunity to conduct observations from this mountain. N.S. acknowledges support from the PSL Iris-OCARV project. N.S. and V.D. acknowledge support from NASA (Grant #80NSSC19K0336). K.A. acknowledges support from the Heising-Simons foundation. We would like to thank the referee for their insightful comments.

References

- Beuzit, J. L., Vigan, A., Mouillet, D., et al. 2019, *A&A*, **631**, A155
 Bos, S. P., Vieuvar, S., Wilby, M. J., et al. 2020, *A&A*, **639**, A52
 Chambouleyron, V., Fauvarque, O., Janin-Potiron, P., et al. 2020, *A&A*, **644**, A6
 Chambouleyron, V., Fauvarque, O., Sauvage, J. F., Neichel, B., & Fusco, T. 2021, *A&A*, **649**, A70
 Chauvin, G., Lagrange, A. M., Dumas, C., et al. 2004, *A&A*, **425**, L29
 Close, L. M., Males, J. R., Hedglen, A., Bouchez, A., & Guyon, O. 2020, ArXiv e-prints [arXiv:2004.06808]
 Codona, J. L., & Kenworthy, M. 2013, *ApJ*, **767**, 100
 Currie, T., Guyon, O., Lozi, J., et al. 2020, *SPIE Conf. Ser.*, **11448**, 114487H
 Deo, V., Gendron, É., Rousset, G., et al. 2019, *A&A*, **629**, A107
 Esposito, S., Puglisi, A., Pinna, E., et al. 2020, *A&A*, **636**, A88
 Ferreira, F., Gratadour, D., Sevin, A., et al. 2018, in *Adaptive Optics Systems VI*, eds. L. M. Close, L. Schreiber, & D. Schmidt, *SPIE, Int. Soc. Opt. Photon.*, **10703**, 1155
 Fitzgerald, M., Bailey, V., Baranec, C., et al. 2019, *BAAS*, **51**, 251
 Frazin, R. A. 2018, ArXiv e-prints [arXiv:1811.05096]
 Gonsalves, R. A. 1982, *Opt. Eng.*, **21**, 215829
 Guyon, O., Sevin, A., Gratadour, D., et al. 2018, in *Adaptive Optics Systems VI*, eds. L. M. Close, L. Schreiber, & D. Schmidt, *SPIE, Int. Soc. Opt. Photon.*, **10703**, 469
 Guyon, O., Norris, B., Martinod, M.-A., et al. 2021, *Int. Soc. Opt. Photon.*, **11823**, 134
 Jovanovic, N., Martinache, F., Guyon, O., et al. 2015, *PASP*, **127**, 890
 Kasper, M., Cerpa Urra, N., Pathak, P., et al. 2021, *The Messenger*, **182**, 38
 Kern, P., Lena, P., Gigan, P., et al. 1989, in *New Technologies for Astronomy*, ed. J. P. Swings, *SPIE Conf. Ser.*, **1130**, 17
 Lozi, J., Jovanovic, N., Guyon, O., et al. 2019, *PASP*, **131**, 044503
 Macintosh, B., Graham, J. R., Ingraham, P., et al. 2014, *Proc. Natl. Acad. Sci.*, **111**, 12661
 Males, J. R., Close, L. M., Miller, K., et al. 2018, *Proc. SPIE*, **10703**, 1070309
 Marois, C., Macintosh, B., Barman, T., et al. 2008, *Science*, **322**, 1348
 Minowa, Y., Hayano, Y., Oya, S., et al. 2010, in *Adaptive Optics Systems II*, eds. B. L. Ellerbroek, M. Hart, N. Hubin, & P. L. Wizinowich, *SPIE Conf. Ser.*, **7736**, 77363N
 Norris, B., Schworer, G., Tuthill, P., et al. 2015, *MNRAS*, **447**, 2894
 Paxman, R. G., & Fienup, J. R. 1988, *J. Opt. Soc. Am. A*, **5**, 914
 Potier, A., Galicher, R., Baudoz, P., et al. 2020, *A&A*, **638**, A117
 Roberts, L. C., Jr., Perrin, M. D., Marchis, F., et al. 2004, in *Advancements in Adaptive Optics*, eds. D. Bonaccini Calia, B. L. Ellerbroek, & R. Ragazzoni, *SPIE Conf. Ser.*, **5490**, 504
 Rousset, G., Fontanella, J. C., Kern, P., Gigan, P., & Rigaut, F. 1990, *A&A*, **230**, L29
 Sauvage, J.-F., Fusco, T., Rousset, G., & Petit, C. 2007, *J. Opt. Soc. Am. A*, **24**, 2334
 Skaf, N., Guyon, O., Boccaletti, A., et al. 2021, *Techniques and Instrumentation for Detection of Exoplanets X*, 11823
 Vidal, F., Ferreira, F., Déo, V., et al. 2018, *Proceedings of Adaptive Optics for Extremely Large Telescopes (AO4ELT5)*, 43
 Vieuvar, S., Bos, S. P., Cassaing, F., et al. 2020, *SPIE Conf. Ser.*, **11448**, 114486D
 Vigan, A., Gry, C., Salter, G., et al. 2015, *MNRAS*, **454**, 129
 Vigan, A., N'Diaye, M., Dohlen, K., et al. 2018, *Proc. SPIE*, **10703**, 107035O
 Vigan, A., N'Diaye, M., Dohlen, K., et al. 2019, *A&A*, **629**, A11

Appendix A: Algorithm description

Data:

- WFS frames
- fast science camera frames

Result: Update the reference of the WFS

Initialization: take the reference of the WFS with the internal light source;

while AO loop is running do

- acquire WFS and fast science focal plane images;
- synchronize data cubes;

while DrWHO iteration (30 s) is running do

- select the 10% best science data;
- select the corresponding WFS images;
- average selected WFS images;
- the resulting WFS frame replaces the WFS reference;
- WFS reference updated.

end

- restart DrWHO iteration

Change of AO point of convergence;

end

Algorithm 1: DrWHO algorithm

Appendix B: Expanding on data synchronization and re-sampling

It is fundamental for DrWHO to synchronize and re-sample PyWFS and short exposure focal plane camera data streams. This is done directly on the telemetry files that are saved to the file system, but not in real time. Indeed, DrWHO needs frames to be synchronized but not precisely in real time, since one iteration of the algorithm is several seconds long. A script for telemetry synchronization was implemented, for which four inputs are required:

- the PyWFS data telemetry timeseries;
- the focal plane camera data stream;
- absolute timestamps of each exposure on both timeseries, retrieved from a hardware latency calibration process
- the output sampling frame rate

For example, if the focal plane camera is run at 1 kHz, the PyWFS camera runs at 2 kHz, and a 3 kHz output frame rate is requested, both input streams will be resampled to a common 3 kHz frame rate. Once launched, the synchronization script waits for data cubes (FITS files) and the corresponding timing files to appear on the hard drive; then it reads them and re-samples them to the same time baseline. This script handles differential latency between the streams, potential missed frames, and possible interruptions in the input streams acquisition. A linear interpolation is performed to re-sample input streams to output data cubes. This script provides output data cubes of the required length (which is an input parameter of the script), that are saved on the file system. As described above, the operation is not real-time as files need to appear to generate the re-sampling. Choosing the length of the data cubes is linked with the latency we need for DrWHO. The script needs to be started once only, then it catches up as the AO loop is running. Figure B.1 gives a visual explanation of this file synchronization.

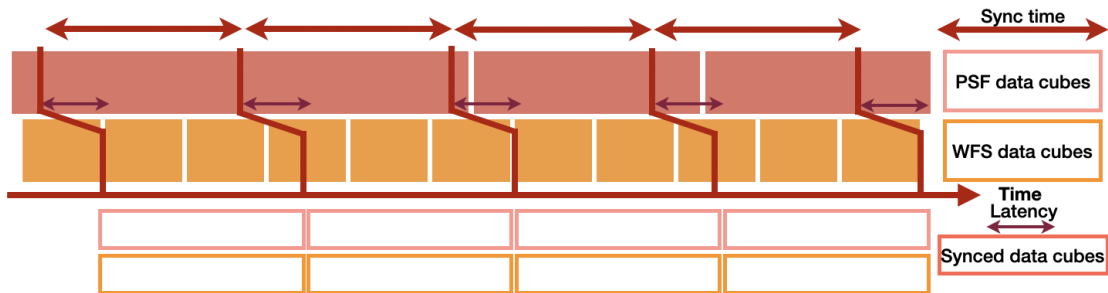


Fig. B.1. Schematic view of how the synchronization of files works. Given two data streams with a known latency between them, the algorithm generates two synchronized files for those data.