Kernel-phase analysis: Aperture modeling prescriptions that minimize calibration errors

Frantz Martinache¹, Alban Ceau¹, Romain Laugier¹, Jens Kammerer^{2,3}, Mamadou N'Diaye¹, David Mary¹, Nick Cvetojevic¹, and Coline Lopez¹

¹ Université Côte d'Azur, Observatoire de la Côte d'Azur, CNRS, Laboratoire Lagrange, Nice, France
 e-mail: frantz.martinache@oca.eu

² European Southern Observatory, Karl-Schwarzschild-Str 2, 85748 Garching, Germany

³ Research School of Astronomy & Astrophysics, Australian National University, Canberra, ACT 2611, Australia

Received 24 October 2019 / Accepted 4 March 2020

ABSTRACT

Context. Kernel phase is a data analysis method based on a generalization of the notion of closure phase, which was invented in the context of interferometry, but it applies to well corrected diffraction dominated images produced by an arbitrary aperture. The linear model upon which it relies theoretically leads to the formation of observable quantities robust against residual aberrations.

Aims. In practice, the detection limits that have been reported thus far seem to be dominated by systematic errors induced by calibration biases that were not sufficiently filtered out by the kernel projection operator. This paper focuses on the impact the initial modeling of the aperture has on these errors and introduces a strategy to mitigate them, using a more accurate aperture transmission model.

Methods. The paper first uses idealized monochromatic simulations of a nontrivial aperture to illustrate the impact modeling choices have on calibration errors. It then applies the outlined prescription to two distinct data sets of images whose analysis has previously been published.

Results. The use of a transmission model to describe the aperture results is a significant improvement over the previous type of analysis. The thus reprocessed data sets generally lead to more accurate results, which are less affected by systematic errors.

Conclusions. As kernel-phase observing programs are becoming more ambitious, accuracy in the aperture description is becoming paramount to avoid situations where contrast detection limits are dominated by systematic errors. The prescriptions outlined in this paper will benefit from any attempt at exploiting kernel phase for high-contrast detection.

Key words. instrumentation: high angular resolution – methods: data analysis – stars: low-mass – binaries: visual

1. Introduction

Within the anisoplanetic field of an imaging instrument, and in the absence of saturation, an in-focus image I can formally be described as the result of a convolution product

$$I = O \star \text{PSF} \quad (1)$$

between the spatially incoherent brightness distribution of an object O and the instrumental point spread function (PSF). The careful optical design of telescopes and instruments assisted by adaptive optics (AO) attempts to bring the PSF as close as possible to the theoretical diffraction limit. Yet even for high quality AO correction, subtle temporal instabilities in the PSF make it difficult to solve for important problems, such as the following: the identification of faint sources or structures in the direct neighborhood of a bright object (the high-contrast imaging scenario) or the discrimination of sources that are close enough to one another to be called nonresolved (the super-resolution scenario). Weak signals of astrophysical interest compete with time-varying residual diffraction features that render the deconvolution difficult.

The overall purpose of interferometric processing of diffraction-dominated images is to provide an alternative to the otherwise ill-posed image deconvolution problem. The technique takes advantage of the properties of the Fourier transform, which turns the convolution into a multiplication. One must,

however, abandon the language describing images and instead manipulate the modulus, which is also referred to as the visibility, and the phase of their Fourier transform counterpart. This Fourier transform can be sampled over a finite area of the Fourier plane traditionally described using the (u, v) coordinates, whose extent depends on the geometry of the instrument pupil.

Nonredundant masking (NRM) interferometry uses a custom aperture mask featuring a finite number of holes that considerably simplifies the interpretation of images. Accurate knowledge of the mask's subaperture locations unambiguously associates particular complex visibility measurements in the image's Fourier transform to specific pairs of subapertures forming a baseline. The Fourier phase Φ at the coordinate (u, v) is the argument of a single phasor:

$$\phi(u, v) = \text{Arg}(v_0(u, v)e^{i(\phi_0(u, v) + \Delta\varphi(u, v))}) \quad (2)$$

$$= \Phi_0(u, v) + \Delta\varphi(u, v), \quad (3)$$

where $v_0(u, v)$ and $\phi_0(u, v)$ represent the intrinsic target visibility modulus and phase for this baseline, respectively, and $\Delta\varphi(u, v)$ is the instrumental phase difference (i.e., the piston) experienced by the baseline at the time of acquisition. The same geometrical knowledge also makes it possible to combine complex visibility measurements by baselines forming closing-triangles, which lead to the formation of closure phases, which are observable quantities engineered to be insensitive to differential piston errors affecting the different baselines.

The closure phase was first introduced in the context of radio interferometry by [Jennison \(1958\)](#) and eventually exploited in the optical starting with [Baldwin et al. \(1986\)](#). This useful observable enables NRM interferometry to detect companions at smaller angular separations than a coronagraph can probe.

Kernel-phase analysis attempts to take advantage of the same property without requiring the introduction of a mask. The description of the full aperture requires a more sophisticated model that reflects the intrinsically redundant nature of the aperture. Any continuous aperture can be modeled as a periodic grid of elementary subapertures, resulting in a virtual interferometric array where every possible pair of subapertures forms a baseline. Whereas the NRM ensures that each baseline is only sampled once, the regular grid results in a highly redundant scenario. For a baseline of coordinate (u, v) and redundancy R , the Fourier phase is that of the sum of R phasors that all measure the same $\phi_0(u, v)$ but experience different realizations of instrumental phase $(\Delta\varphi_k)_{k=1}^R$:

$$\phi(u, v) = \text{Arg} \left(\sum_{k=1}^R v_0(u, v) e^{i(\phi_0(u, v) + \Delta\varphi_k)} \right). \quad (4)$$

In the low-aberration regime provided by modern AO systems, the impact the residual pupil aberration φ has on the Fourier phase can be linearized and Eq. (4) can be rewritten as:

$$\phi(u, v) = \phi_0(u, v) + \frac{1}{R} \sum_{k=1}^R \Delta\varphi_k. \quad (5)$$

The list of what pairs of subapertures contribute to the complex visibility of a redundant baseline is kept in the baseline mapping matrix $\mathbf{A}$. It contains as many columns as there are subapertures (n_A) and as many rows as there are distinct baselines (n_B). Elements in a row of $\mathbf{A}$ are either 0, 1, or -1 (see Fig. 1 of [Martinache 2010](#)). The phase that is sampled at all relevant coordinates of the Fourier-plane and gathered into a vector Φ , can thus be written compactly as:

$$\Phi = \Phi_0 + \mathbf{R}^{-1} \cdot \mathbf{A} \cdot \varphi, \quad (6)$$

where $\mathbf{R}$ is the diagonal (redundancy) matrix that retains the tally of how many subaperture pairs contribute to the Fourier phase for that baseline, φ is the aberration experienced by the aperture, and Φ_0 is the Fourier phase associated with the object being observed; it is related to the object function O of Eq. (1) by the Van-Cittert Zernike theorem. The redundancy $\mathbf{R}$ is expected to be directly proportional the modulus transfer function (MTF) of the instrument. The product $\mathbf{R}^{-1} \cdot \mathbf{A}$, referred to as the phase transfer matrix, describes the way pupil phase aberration propagate into the Fourier plane. The baseline mapping and the phase transfer matrices are rectangular and feature n_B rows (the number of baselines) for n_A columns (the number of subapertures in the pupil), with $n_B > n_A$.

As shown in [Martinache \(2010\)](#), selected linear combinations of the rows of the phase transfer matrix cancel the effect of the pupil phase φ out. These linear combinations, which are gathered into an operator called $\mathbf{K}$ (the left-hand null space or kernel of the phase transfer matrix), project the Fourier phase onto a subspace that is theoretically untouched by residual aberrations. The resulting observables, called kernel phases, are a generalization of the concept of closure phase, which can be found for an arbitrary pupil, regardless of the level of redundancy.

Practice suggests that the kernel and closure phase do not self-calibrate perfectly. Recently published studies using kernel

phase to interpret high-angular resolution, diffraction-dominated observations indeed lead to contrast detection limits, which are mostly constrained by systematic errors ([Kammerer et al. 2019](#); [Laugier et al. 2019](#); [Sallum & Skemer 2019](#)) instead of statistical errors ([Ceau et al. 2019](#)). The goal of this paper is to provide insights into the limitations of Fourier phase methods, in general, and to introduce the means by which to improve upon these limitations.

This difficulty affects the kernel-phase interpretation of images as well as that of NRM interferograms. Despite the generalized assistance of AO during NRM observations ([Tuthill et al. 2006](#)), the need for long integration times and the use of subapertures that are not infinitely small means that interferograms are not simply affected by a simple and stable piston, but by a time-varying higher order amount of aberration ([Ireland 2013](#)). Closure phases, which are thus acquired on a point source, therefore rarely reach zero and are biased.

Nevertheless, even when evolving over time, if the spatial and temporal properties of the perturbations experienced by the system remain stable across the observation of multiple objects, the overall resulting bias is also expected to remain stable. It is therefore possible to calibrate the closure phases acquired on a target of interest with those acquired on a point-source. Thus calibrated closure phases have been used as the prime observable for the detection of low to moderate contrast companions ([Kraus et al. 2008](#)). There is, however, a limit to the stability of the observing conditions when moving from target to target: Changes in elevation, apparent magnitude for the AO, and telescope flexures result in the evolution of the closure-phase bias. Therefore, in practice, the observations never calibrate perfectly and the evolution of the calibration bias results in what is generally referred to as a systematic error.

For low-to-moderate contrast detections up to a few tens, systematic errors are often a minor contribution that do not significantly affect the interpretation of the data. However, as observing programs become more ambitious, attempting the direct detection of higher contrast companions ([Kraus & Ireland 2012](#)) in a part of the parameter space that cannot be probed by coronagraphic techniques yet, systematic errors become more important than statistical errors and one must resort to more advanced calibration strategies using multiple calibrators ([Ireland 2013](#); [Kammerer et al. 2019](#)) to better estimate the calibration bias that minimizes the amount of systematic error.

2. Fourier phase calibration errors

Kernel-phase analysis requires approximating the near-continuous aperture of the telescope by a discrete representation: A virtual array of subapertures, which are laid on a regular grid of a predefined pitch s , paves the surface covered by the original aperture. The computation of all the possible pairs of virtual subapertures in the array leads to the creation of a second grid of virtual baselines, the majority of which are highly redundant. An example is shown in Fig. 1 for the aperture of an 8 m telescope, which is discretized with a grid of pitch $s = 42$ cm. Keeping track of what pairs of subapertures contribute to all the baselines leads to the computation of the baseline mapping matrix $\mathbf{A}$ and the redundancy matrix $\mathbf{R}$. The two are used to eventually determine the kernel operator $\mathbf{K}$ that generalizes the notion of closure phase, as introduced by [Martinache \(2010\)](#).

The following simulation illustrates the interest of this line of reasoning. Using a single, simulated, monochromatic ($\lambda = 1.6 \mu\text{m}$), and noise-free image of a 10:1 contrast binary object (located two resolution elements to the left of the primary) that is

Fig. 1. Binary discrete representation of the SCEXAO instrument pupil for kernel-phase analysis. *Left panel:* discretized instrument pupil built from a square grid of pitch $s = 42$ cm. *Right panel:* Fourier coverage associated with this discretization. The color code used to draw the Fourier coverage reflects the redundancy associated with the virtual interferometric baselines.

Fig. 2. *Left:* simulated monochromatic ($\lambda = 1.6 \mu\text{m}$) SCEXAO image of a 10:1 binary in the presence of 100 nm of coma along the axis of the binary. *Right:* associated Fourier phase ranging from ± 1.5 radian (see also Fig. 3).

affected by a fairly large (100 nm rms) amount of coma is shown in the left panel of Fig. 2. The Fourier phase Φ extracted from this image (shown in the right panel of Fig. 2) appears to be completely dominated by the aberration. The plot of the same raw Fourier phase (the blue curve in the top panel of Fig. 3) compared to the predicted Fourier signature of the sole binary Φ_0 confirms this observation. By using the kernel operator $\mathbf{K}$ computed according to the properties of the discrete model¹ represented in Fig. 1, it is possible to project the 546-component noisy Fourier phase vector Φ onto a subspace that results in the formation of a 414-component kernel. The bottom panel shows how, despite the drastic difference between the raw and theoretical Fourier phase, the two resulting kernels match one another: The kernel operator effectively erases the great majority of the aberrations affecting the phase present in the Fourier space, while leaving enough information to describe the target in a meaningful manner, such that:

$$\mathbf{K} \cdot \Phi = \mathbf{K} \cdot \Phi_0. \quad (7)$$

Although the application of $\mathbf{K}$ strongly reduces the impact of the aberration, the match between the kernel curves is not perfect. The small difference between the two example curves is what is generally referred to as the calibration error. This error can be independently measured using the image of a point source

¹ The model was computed using a python package called XARA, which was developed in the context of the KERNEL project, and is available for download <http://github.com/fmartinache/xara/>

Fig. 3. Demonstration of the impact of the kernel processing. *Top panel:* experimental Fourier phase extracted from a single aberrated image shown in Fig. 2 (the blue curve) bears little resemblance with the theoretical expected binary signal (in orange). Contrasting with the raw Fourier phase, the *bottom panel* shows how the projection onto the kernel subspace efficiently erases the impact of the aberration and brings the experimental kernel-phase curve ($\mathbf{K} \cdot \Phi$), which is also plotted with a solid blue line, much closer to its theoretical counterpart ($\mathbf{K} \cdot \Phi_0$), which is now plotted with a dashed orange line so as to better distinguish them.

(a calibrator), assuming that the system suffers from the same aberration. In this noise-free scenario, the subtraction of the kernel phase extracted from one such calibration image would result in a perfect match. An instrumental drift between the two exposures would result in a new bias. If the magnitude of this new bias becomes comparable to or larger than the statistical uncertainties; the interpretation of the kernel and closure phase typically requires invoking a tunable amount of systematic error added in quadrature to the uncertainty.

3. Kernel phase discretization prescriptions

Given that no noise was included in this ideal scenario, the fact that some aberration leaks into the kernel and results in the need for calibration suggests that the discrete model used to describe how pupil phase propagates into the Fourier plane is not sufficiently accurate. Thus, we look into ways to improve it.

3.1. Building a discrete representation

The discretization process is as follows: A high-resolution 2D image of the aperture is generated from the details of the pupil specifications (outer and inner diameter, spider thickness, angle, and offset). A square grid of subapertures of pitch s is laid atop the pupil image and the points that fall within the open parts of the aperture are kept in the model. To be counted amongst the virtual subapertures, the area of the transmissive part of the original aperture that overlaps with the region covered by the square virtual subaperture has to be greater than 50%. When building the model, it is critical to ensure that no virtual baseline is greater

Fig. 4. Example of discretized version of SCEXAO’s pupil using a square grid, aligned with the center of the aperture. Only the virtual subapertures for which the transmission, which was determined as the normalized intersection between the virtual subaperture and the underlying pupil, is greater or equal to 50% are kept as valid subapertures contributing to the model. The $s = 0.42$ m pitch value was chosen so as to fit an entire number (here 19) of subapertures across the pupil diameter.

than the actual telescope diameter: This would indeed result in attempting to extract information from outside the Fourier domain to which the true aperture gives access. To mitigate this risk, one needs to first ensure that an entire number of subapertures fits within one aperture diameter and then eliminate all of the computed baselines that are greater than the aperture diameter.

A regular grid is required so that the density of the discrete representation is as homothetic as possible to the original aperture: This translates into a model redundancy $\mathbf{R}$ that better matches the true instrument MTF. It is also important to align the grid with the aperture model so that the symmetry properties of the apertures are reflected in its discrete representation: one either uses a grid that is centered on the aperture, which features an integer odd number of apertures (the option retained to build the discretized aperture shown in this paper), or an offset grid with an integer even number of subapertures.

Figure 4 introduces the example that serves as a reference to compare the relative merits of different discrete models. It uses the Subaru Telescope pupil mask of the SCEXAO instrument (Jovanovic et al. 2015), which is characterized by its large (2.3 m diameter) central obstruction and nonintersecting thick spider vanes at the nontrivial angle of 51.75° (Lozi et al. 2009). This nontrivial aperture geometry makes it a rich test case. Using the aforementioned recommendations, that is, a centered grid with a $s = 0.42$ m pitch, fits almost exactly 19 virtual subapertures across the aperture nominal diameter of 7.92 m.

The $n_A = 272$ virtual subapertures of this array form $n_B = 562$ distinct baselines. As discussed in Martinache (2013), for a rotational symmetric of order 2 aperture², the Fourier phase and its kernels are insensitive to even order aberrations. This property is reflected in the properties of the linear phase transfer model: The number of nonsingular values of the baseline transfer matrix $\mathbf{A}$ should be equal to $n_E = n_A/2$, therefore leaving $n_K = n_B - n_A/2$ kernel phases. For the kernel analysis to lead to optimal results, it is important to ensure that these properties are

verified. An aperture that does not respect this symmetry condition results, in contrast, in less ($n_K = n_B - n_A + 1$) kernels.

3.2. Grid pitch and image size

The 42 cm pitch of the grid illustrated in Fig. 4 does not offer enough resolution to reflect the presence of the 25 cm thick spider vanes of the real aperture. This contrasts with models that have generally been used since the inception of the kernel phase (see for instance Fig. 2 of Martinache 2010) that have naively overemphasized the impact of spider, which in turn contributes, as is made clear below, in amplifying the calibration bias.

The pitch s of the grid is of course a free parameter that can be adjusted: The finer the grid, the more representative the details of the pupil are expected to become and the more capable of capturing higher spatial frequency are the components of the images. A discrete model with a finer pitch, however, implies the description of a wider effective field of view (of radius $0.5\lambda/s$) over which the kernel analysis can lead to meaningful results. The size of the image therefore imposes a limit on how fine the pitch can get.

For the wavelength ($\lambda = 1.6\mu\text{m}$) of the simulations used in this section, the plate-scale (16.7 mas per pixel) and size (128×128) of the images suggest that the pitch cannot be finer than $s > 206.265 \times 1.6 / (128 \times 16.7) \approx 0.15$ m. Beyond this simulation scenario, image noises induced by dark current, readout and photon noise, and a preference for computationally manageable problems guide the user toward using coarser models in practice.

3.3. Comparing models

To assess the relative merits of multiple models, we look at the impact the discretization strategy has on the magnitude of the calibration bias. We have seen that the pitch of the model impacts the overall dimension of the problem. It also impacts the associated redundancy $\mathbf{R}$ and therefore the overall magnitude of the kernel phases extracted from a given image. To enable a meaningful comparison of multiple models, we compare the root mean square of the calibration bias to the theoretical standard deviation of the theoretical signature (see, for instance, Eq. (27) of Ceau et al. 2019) induced by a 100:1 contrast companion that would be located two resolution elements to the left of the primary along the horizontal axis.

The simulations systematically include a 20 nm rms static (odd) aberration which is either a three-cycle horizontal sinusoid or coma along the same direction. These two examples were selected first because they are both perfectly odd, and therefore have full impact on the analysis, and second because they feature different structures: The impact of the sinusoidal modulation is more uniformly distributed across the aperture, whereas the impact of the coma (similar to that of most higher order Zernike modes) is stronger toward the edges. The same two images (128×128 pixels, one featuring coma and one featuring the sinusoidal aberration) were processed using the kernel-phase pipeline, using three discrete models. The results of these three analyses are summarized in Fig. 5, which features, side by side, a rendering of the discrete aperture model and the plot of the thus biased kernel-phase vector extracted from either image, as well as in Table 1, which summarizes the dimensions of the models and their intrinsic sensitivity to calibration error.

The first model, presented in the top panel of Fig. 5, is the reference using the $s = 0.42$ m pitch grid introduced earlier. Using this model, the magnitude of the calibration bias extracted from the images affected by either type of aberration represents a

² The aperture is identical to itself rotated by 180° relative to its center.

Fig. 5. Comparison of the self-calibrating performance of the kernel-phase analysis of a single image for three discrete models of the same aperture. Each of the three panels features, side by side, a 2D representation of the discrete aperture model used and a plot of the kernels extracted from the image of a point source (the calibration error) in the presence of either coma (the orange curve) or a three-cycle sinusoidal aberration (the red curve) and how they compare to the signal of a 100:1 contrast binary (the blue curve). *Top panels:* reference binary model of the SCEAO pupil, with a 42 cm pitch; *middle panel:* denser model with a 21 cm pitch that more accurately matches the fine structures of the telescope; *third panel:* model that uses the original 42 cm pitch grid, but it includes the transmission function.

Table 1. Summary of the model properties and their performance.

	Sparse binary	Dense binary	Sparse gray
Pitch	0.42	0.21	0.42
n_A	272	956	300
n_B	562	2132	562
n_K	426	1654	412
Ref. signal	0.435	1.583	0.402
Coma bias	0.180 (41%)	0.317 (20%)	0.036 (9%)
Sine bias	0.286 (66%)	0.263 (17%)	0.026 (6%)

Notes. We note that n_A , n_B , and n_K represent the number of subapertures, the number of baselines, and the number of kernels of each model, respectively. The coma and sine bias rows show the magnitude of the bias induced 20 nm rms of coma and the sinusoidal aberration in radians, respectively.

significant fraction (on the order of 50%) of the signature of the 100:1 binary companion. Under such circumstances, the contrast detection limits associated with these uncalibrated kernel phases are likely to be rapidly dominated by this systematic error. The middle panel of Fig. 5 illustrates the impact of a finer $s = 0.21$ m grid pitch: The model better reflects the presence of the spiders and the overall shape of the pupil. A larger number of kernels was extracted from the same image (almost four times as many), but more importantly, for this discussion, the relative magnitude of the calibration bias is reduced by a factor $\approx 2-3$: A kernel-phase analysis based on a finer and more accurate description of the original aperture would feature reduce model-induced calibration errors and would, therefore be less susceptible to calibration errors in general.

Increasing the resolution of the grid is not the only available option. One can indeed also refine its description by allowing for a variable subaperture transmission. In addition to deciding whether to keep or discard one virtual subaperture as part of the model, the information on the clear fraction of the subaperture, translated into a local transmission value, can be appended and taken into account when creating the phase transfer model. Such a “gray aperture” model makes it possible to more accurately describe the edges and high-spatial frequency features of the aperture without necessarily increasing the pitch of the model. One example using such a continuous transmission model is illustrated in the bottom panel of Fig. 5: despite using a grid pitch identical to the reference model, the discrete representation of the aperture clearly better outlines the finer features of the aperture as the trace of the spider vanes becomes visible. For this example, the transmission cut-off value was set to 10^{-3} : The gray model includes a slightly higher number of virtual subapertures than its binary counterpart for which the cut-off was set to 0.5. In the end, one forms a number of kernels (see Appendix A for a general discussion regarding the number of kernels) similar to the binary case. The improvement brought by the inclusion of this transmission model is substantial: The magnitude of the bias is brought well below 10% of the signature of the 100:1 binary companion.

The two effects of a finer resolution and a transmission model can be compounded to lead to even better performance. Generally, whether one uses a binary or a gray model, doubling the resolution of the grid leads to an improvement by a factor of ≈ 2 . The performance of the kernel phase reaches a point where the details of the implementation of the upstream simulation becomes critical.

Overall, there seems to be no significant difference between the two types of aberrations introduced. Sinusoidal modulation

seems to be better processed in general, which is likely because of the sharper edge structure of the coma that systematically requires more resolution. The impact of aberrations of higher spatial frequency, beyond what the chosen model can effectively describe, are filtered out either by adequate image cropping (following the recommendations given in Sect. 3.2) or by the application of an image mask. We can conclude that including the aperture transmission model is a major improvement that renders the kernel-phase analysis less susceptible to systematic errors.

4. Kernel-phase analysis revisited

In this section, we use the recommendations outlined in the previous section and apply them to a series of data sets whose kernel-phase analysis has already been published. We feature the following two applications: the analysis of a ground-based data set published by Pope et al. (2016) and an extended version of the data set used for the original kernel-phase publication by Martinache (2010). The review of these two applications further illustrates the importance of better aperture modeling practices for kernel-phase analysis.

4.1. Palomar/PHARO

The data consist of two data cubes of 100 images of the binary system α -Ophiuchi (Hinkley et al. 2011) and of the single star ϵ -Herculis that were acquired with the PHARO instrument (Hayward et al. 2001) at the focus of the Palomar Hale Telescope after AO correction was provided by the P3K AO system (Dekany et al. 2013).

The data cubes were recovered from an archive linked in the original publication. The preprocessed large 512×512 pixels original frames were first cropped down to a more manageable 64×64 pixel size. With a plate scale of 25 mas per pixel, the field of view is still ± 800 mas. To reach sufficient resolution in the Fourier space, a fast Fourier transform (FFT) based extraction algorithm, such as the one used in the original study, requires an adapted amount of zero-padding. The now standard complex visibility extraction method of XARA instead explicitly computes the discrete Fourier transform for the spatial frequencies of the discrete model, such as suggested by Soummer et al. (2007), and filters out subpixel centering errors as used by Kammerer et al. (2019). The cropping of the image not only filters out the higher level of noise brought out by weakly illuminated pixels, but it is also more computationally efficient as it requires the computation of smaller Fourier transform matrices.

Images were acquired using the K_s filter (central wavelength: $2.14 \mu\text{m}$) and the medium cross pupil mask inside the PHARO camera was used to limit the risk of saturation in the image for the otherwise bright target of interest. An example of the image is shown in Fig. 6. The image presents a few noteworthy features as follows: The apparent companion that is clearly visible in the bottom left quadrant is a ghost induced by the 0.1% neutral density filter used at the time of the acquisition. This ghost is present in all of the frames, including those acquired on the calibrator. Strong diffraction spikes are also visible and they were induced by the very thick spider vanes of the medium-cross mask, whose orientation does not quite match the grid of pixels (the upper vertical spike slightly leans to the left).

We built a new discrete gray model of the medium-cross based on the specifications published by Hayward et al. (2001), which were confirmed by an image of the pupil enabled by one of the modes of the camera. In the image provided by the pupil imaging mode of the PHARO camera, the medium-cross mask

Fig. 6. Example of frame acquired on α -Ophiucus (nonlinear scale: power 0.25). Notable features of this image include the thick and slightly tilted diffraction spikes induced by the medium-cross pupil mask used at the time of acquisition as well as the ghost induced by the neutral density filter in the bottom left quadrant of the frame. The companion later recovered by the kernel-phase data reduction is buried underneath the first diffraction ring, which is to the left of the primary.

Fig. 7. Representation of the discrete model (*left*: aperture, *right*: Fourier coverage) of the PHARO med-cross aperture, using the transmission model of the true aperture. To further reduce the amount of systematic error, the model was built using a square grid that was rotated to match the orientation of the original pupil mask. The impact of the presence of the spiders in the model is revealed in the Fourier plane as four small depressions appear in the intermediate spatial frequency range.

appears to be rotated counterclockwise by two degrees. We used the recipe outlined in Sect. 3 to produce gray discrete representation of the true aperture using a square grid with a pitch $s = 0.16$ m that was then rotated to match the grid to the true aperture. To eliminate possible mistakes, we used a simulation reproducing the properties of the PHARO K_s -band images that included a fixed amount of aberration and rotated our mask until we found the orientation that minimizes the amount of calibration error. The optimum model that was thus identified is shown in Fig. 7.

The presence of the ghost in all of the images contributes to the calibration bias of the data. Pope et al. (2016) chose to further window the data so as to mask the ghost out before attempting to extract the kernels. This, however, leaves too few useful

pixels to lead to the formation of $n_K = 1048$ kernels of the model ($n_A = 528$, $n_B = 1312$). In this analysis, we keep all the information available in the image, under the assumption that the contribution of the ghost is erased when subtracting the kernel phase from the calibrator.

In the high-contrast regime, which in practice applies when the contrast is greater than 10:1, the amplitude of the kernel signature of a binary is expected to be directly proportional to the contrast (secondary/primary). This makes it convenient to compute the normalized dot product between the calibrated signal and the theoretical signal of a high-contrast binary, over a finite grid of positions around the primary. The use of such colinearity maps is introduced by Laugier et al. (2019): The presence of a clear maximum in this map shows where the input signal best matches the theoretical signature of a binary.

Figure 8 shows the result of this computation for the raw signal of both the target and calibrator as well as for the calibrated signal of α -Ophiuchi over a ± 500 mas field of view. The kernel phase, similar to the closure phase, is a measure of asymmetry of a target so the colinearity map is always antisymmetric. The two uncalibrated maps prominently feature the signature of the ghost present in all images in the bottom left quadrant as well as other structures at closer separations (up to ~ 200 mas). While the signature of the ghost is expected, these intermediate separation features (particularly on the map of the calibrator) suggest that the kernel phases are still affected by a calibration bias, despite our efforts to minimize the modeling induced errors. However, given that individual images were integrated over 1.4 s, which is many times the coherence time, the kernel phases are still affected by an additional bias induced by temporal variance described by Ireland (2013). We can observe that the subtraction of the kernel phases of the calibrator from those obtained on α -Ophiucus effectively erases these features along with that of the ghost. The bright bump (and its antisymmetric dark counterpart) that is clearly visible to the left (and the right) of the central star in the right panel of Fig. 8 is indicative of the quality of the detection of the companion around α -Ophiuchi.

We used the location of the maximum of colinearity as the starting point for a traditional χ^2 minimization algorithm using the Levenberg-Marquardt method that is shipped as part of the python package *scipy*. The uncertainties associated with the calibrated kernel phases were simply computed as the quadratic sum σ_e of the uncertainties deduced from variance between frames for the α -Ophiuchi and ϵ -Herculis. The result of this optimization is represented in the correlation plot of Fig. 9: The model fit looks very convincing and locates the companion in the area hinted at in the calibrated colinearity maps of Fig. 8.

The careful modeling of the aperture unfortunately does not suffice in eliminating the need for ad hoc systematic errors at the time of the optimization. A relatively large amount of systematic error ($\sigma_S = 1.2$ rad) still needs to be added quadratically to the experimental dispersion ($\sigma_E = 0.1$) in order to give a reduced $\chi^2 = 1$ for the following parameters: separation $\rho = 123.5 \pm 2.9$ mas, position angle $\theta = 86.5 \pm 0.2^\circ$, and contrast $c = 25.1 \pm 1.1$. It should be pointed out that these do not quite match the values reported (see Table 1 of Pope et al. 2016) for the NRM observation that usually set the standard. It should, however, also be pointed out that the new contrast estimate is in good agreement with the measurements reported in Table 3 of Hinkley et al. (2011), which were taken when the companion was at larger angular separation.

While the signature of the companion is more clearly visible in this analysis than in the results reported by Pope et al. (2016), the situation is still not fully satisfactory as our improved model

Fig. 8. Colinearity maps for the raw and calibrated kernel phases extracted from the data, over a ± 500 mas field of view. *Left panel:* map built from the uncalibrated kernel phases of α -Oph. *Middle panel:* same map built from the uncalibrated kernel phase of ϵ -Her. *Right panel:* colinearity map of the calibrated signal, which is the difference between the kernel phases of α -Oph and ϵ -Her. The two uncalibrated maps prominently feature the signature of the ghost present in all images in the bottom left quadrant as well as other structures at closer separations. The map of the calibrated signal on the other hand shows that most of these features are gone and reveals a positive bump on the left-hand side of the central star (for a separation of 100 mas and a position angle of 84°), which is taken as an indication of the presence of a companion.

Fig. 9. Correlation plot of the calibrated signal of α -Ophiuchi with the best binary fit solution: following the hint provided by the colinearity map of the calibrated data, the signal present in the data is fairly well reproduced by a binary companion located at angular separation $\rho = 123.4 \pm 1.7$ mas, position angle $\theta = 86.5 \pm 0.5$ and contrast $c = 25.1 \pm 0.1$. Residual structure in the data is accounted for by the introduction of an ad hoc amount of systematic error so that the reduced $\chi^2 = 1$.

of the aperture did not lead to a detection with uncertainties on the binary parameters driven solely by statistical errors. Several explanations were invoked in the original publication to justify the subpar performance, which still apply here. The substandard seeing conditions that induce variability in the AO correction on targets of distinct magnitudes and the fact that both sources were acquired in very different areas of the detector explain, in large part, how the statistical variance experienced during the observation cannot, on its own, be representative of all errors affecting the kernel phase. Since this new analysis uses a model

that is adapted to the information available in the data cubes, it draws a more favorable picture for kernel phase, which shows a much more convincing result here.

4.2. HST/NICMOS

As pointed out in Sect. 3.2, the seminal kernel-phase publication used a rather crude and discrete representation of the aperture of the *Hubble* Space Telescope and was nevertheless able to report the detection of a companion to the M-dwarf GJ 164 (Martinache et al. 2009). In attempting to accurately model the effective aperture of the NICMOS1 instrument used to acquire the data, we refer to the work of Krist & Hook (1997), which suggests the presence of an important ($\sim 10\%$) misalignment of the instrument cold mask relative to the original optical telescope assembly (OTA) that was completely overlooked by Martinache (2010).

Multiple data sets recovered from the HST archival were acquired on GJ 164 on epoch 2004-02-14 UT (proposal ID #9749) in several narrow band filters: *F108N*, *F164N*, and *F190N* (Viana et al. 2009). Our updated analysis also includes images of calibration star SAO 179809 observed at a single epoch (1998-05-01, proposal ID #7232) acquired in the *F190N* filter. The original 256×256 pixel images were bad-pixel corrected, recentered, and cropped down to 84×84 pixel size over which the *F190N* filters seem to feature good signal-to-noise ratio (S/N), before being gathered into data cubes. With a plate scale of 43 mas per pixel, the effective field of view is thus ± 1.8 arcsec. According to the image sampling constraints addressed in Sect. 3.2, the size of the field of view and the wavelength of acquisition set a limit to how fine the pitch can get, which for the *F190N* filter translates into $s = 0.109$ m. Although the model pitch could be updated for the shorter wavelengths, which for a fixed image size give access to an increasing number of spatial frequencies, we built a single discrete model (including transmission) for a homogeneous analysis across the entire data set that enables comparison across spectral bandpasses.

We, however, first need to demonstrate that the discrete model indeed benefits from the updated aperture description recommended by Krist & Hook (1997). For this, we used the images of the calibration star SAO 179809. In Sect. 3.1, we

Fig. 10. Comparison of the predicted redundancy with the experimental MTF (estimated from the modulus of the Fourier transform of $F190N$ images of calibration star SAO 179809) for two models: the first (*top panel*) assumes that the aperture is that of the HST Optical Telescope Assembly (OTA) and the second (*bottom panel*) takes into account the misaligned cold mask of the NICMOS camera. While the redundancy of the first model fails to reproduce the modulus of the Fourier transform that was effectively measured from the image, the second model convincingly matches the fine features of the instrument MTF.

introduce the idea that an accurate discrete model should translate into a predicted redundancy diagonal matrix $\mathbf{R}$ that matches the true instrument MTF, which for a calibration star should correspond to the modulus of the complex visibility extracted for the different baselines of the model. Figure 10 illustrates this property and compares the redundancy associated with models characterized by the same $s = 0.109$ m pitch for two apertures: one that includes the outline of the OTA only (top panel) and one that includes the misaligned NICMOS cold mask (bottom panel). While the OTA model should already be an improvement on the one originally used, we can observe that the associated redundancy fails to reproduce the modulus of the Fourier transform computed for the corresponding spatial frequencies. The gap is particularly visible for intermediate spatial frequencies that feature less power than what is predicted by the model. The more accurate model including the misaligned cold mask is a major improvement as the predicted redundancy $\mathbf{R}$ almost perfectly matches the fine features (in particular the dropped lobes visible for baseline indices ranging from 400 to 800) of the experimental MTF.

Fig. 11. Kernel-phase histograms computed from a set of images of calibration source SAO 179809 using two aperture models characterized by the same pitch $s = 0.109$ m. The first model assumes that the aperture is that of entire optical telescope assembly (OTA) and results in the blue histogram. The second model takes into account the misaligned cold mask inside the NICMOS camera and results in the red histogram.

Unlike any of the previously considered scenarios, the pupil here is clearly not rotation symmetric so we do not expect to form the optimal number of kernels (see the end of Sect. 3.1). The more accurate model is nevertheless expected to translate into a more accurate kernel phase. As SAO 179809 is a calibration source, we can verify that the magnitude of the calibration biases decreases by comparing (see Fig. 11) the histograms of the kernel phase computed using the two aforementioned models. The improvement is significant with a reduction of the standard deviation by a factor ~ 10 , despite a larger number of kernels in the better model (375 vs 320), demonstrating one more time that a more accurate model reduces the impact of calibration errors. With the accurate model, the magnitude of the calibration bias ($\sigma_S = 0.222$ radians) is now comparable to that of a 100:1 contrast ratio companion ($\text{rms} = 0.215$ radians) that is located two resolution elements away from the primary.

The images of SAO 179809 were acquired more than five years before those of GJ 164. We continue to use the same aperture model, but do not expect to use the kernel phase of SAO 179809 reliably to calibrate the kernel-phase signal of GJ 164. The magnitude of the calibration error estimated from the observation of SAO 179809 can, however, provide an order of magnitude for the expected fitting error for a binary such as GJ 164.

One interesting feature of the GJ 164 data set is the availability of images in multiple filters: 80 in the $F190N$ filter, 40 in the $F164N$ filter, and 10 in the $F108N$ filter for this 2004-12-23 epoch. Figure 12 shows a snapshot of GJ164 for these three filters. In addition to the expected linear scaling of the diffraction with the wavelength, one observes the linear scaling size of the circular window, which matches the effective field of view induced by the choice of a unique model with a fixed pitch. The phase transfer model at the core of the kernel analysis is achromatic: The pupil coordinate points and baselines are indeed all expressed in meters and not in radians as is customary in long baseline interferometry. At the time of data extraction, however, the wavelength needs to be taken into account in the computation of a discrete Fourier transform matrix in order to match the sampling of the data.

Fig. 12. Snapshots of GJ 164 AB for the 2004-12-23 epoch in the three NICMOS narrow band filters: *F108N*, *F164N*, and *F190N*. The non-linear scaling of the image makes the window applied to the data more apparent. Given that a single discrete model with a fixed pitch is used to process this data set, the window, used to cover a finite number of resolution elements, must be scaled linearly with the wavelength. Here, the effect of the window is mostly to reduce the contribution of poorly illuminated pixels to the overall noise of the kernel phase.

The published analysis of the *F190N* images has revealed that a companion is present at an angular separation ~ 90 mas, which is on the order of $0.5\lambda/D$. In the high-contrast regime, the kernel-phase signature of a binary companion of contrast c at wavelength λ has a simple analytic expression:

$$\mathbf{K} \cdot \Phi_0(u, v) = \mathbf{K} \cdot c \times \sin \frac{2\pi}{\lambda}(\alpha u + \beta v), \quad (8)$$

where α and β are the angular Cartesian coordinates of the companion (in radians) and u and v are the vectors collecting the coordinates of the baselines (in meters). Assuming that the contrast of the companion is constant for the different filters, we can write the derivative of the binary kernel-phase signal relative to the wavelength as:

$$\frac{\partial}{\partial \lambda} \mathbf{K} \cdot \Phi_0(u, v) = \mathbf{K} \cdot c \times \cos \frac{2\pi}{\lambda}(\alpha u + \beta v) \times \frac{-2\pi}{\lambda^2}. \quad (9)$$

If the companion is unresolved, the cosine term varies slowly and the dominant wavelength dependant effect is the overall $1/\lambda^2$ scaling factor of Eq. (9). Thus by multiplying the kernel phase extracted in the different filters by associated wavelength (expressed in microns) squared, we expect to see signals of comparable structure and magnitude. Figure 13 shows the result of one such comparison for the median signal extracted from the three sets of images. The stability of the signature of the companion across the three bands, covering almost a decade, is striking and suggests that the contrast is indeed stable for the different filters. Additionally, this attests to the consistency of the kernel-phase data analysis.

However, moving from 1.9 to $1.08 \mu\text{m}$ almost doubles the resolving power. The signature of the companion, which is expected to be degenerate in the *F190N* filter, for which one observes a strong correlation between contrast and angular separation, is better constrained by the *F108N* observation. The three data sets were combined to feed a five-parameter model fit optimization algorithm, including two astrometric parameters and three contrasts, which lead to a total of 1120 degrees of freedom. The result of this global optimization is represented in the correlation plots of Fig. 14, split by filter. The best solution places the companion 89.3 ± 0.4 mas away from the primary at the 100.4 ± 0.1 degree position angle and leads to the following three contrasts: 6.23 ± 0.1 in the *F108N* filter, 6.19 ± 0.1 in the *F164N* filter, and 6.36 ± 0.1 in the *F190N* filter. Figure 14 also shows that the $1/\lambda^2$ signal scaling factor of the binary signal (see Eq. (9)) leads to an intrinsically higher S/N for the observation at the shorter wavelength.

Fig. 13. Representation of the median kernel-phase vector extracted for the data sets of the three filters (*F108N*, *F164N*, and *F190N*), rescaled by the wavelength of the filter (taken in microns) squared. The three signals were rescaled and are thus very consistent with one another, confirming the presence of a near constant contrast structure partly resolved from the central star.

Although the astrometric solution for the combined fit is generally consistent with the result published by [Martinache \(2010\)](#), the contrast in the *F190N* is revised and drops from 9.1 ± 1.2 to 6.36 ± 0.1 . The origin of the initial overestimation of the contrast, which was not captured by the uncertainty, is not clear and can likely be attributed to the following combination of causes: the use of an inappropriate aperture model, the strong contrast-separation degeneracy of the *F190N* observation, and an overall more mature data analysis process today. One can incidentally note that the revised *F190N* and *F164N* contrasts are in better agreement with the majority of the contrast measurements reported by [Martinache et al. \(2009\)](#) with NRM observations using broad band *H* and *K* filters.

In the absence of a calibrator, the computation of parameter uncertainties requires the introduction of a controlled amount of systematic error (added in quadrature to the measured statistical uncertainties) so that the reduced χ^2 is unity. In this case, $\sigma_S = 0.15$ rad amount of systematic error is required, which seems to be comparable to the magnitude of the calibration bias that was estimated (≈ 0.22 rad) after the analysis of the SAO 179809 data set. Unlike what was the case with the PHARO data set (see Sect. 4.1), it seems that our modeling of the aperture and the interpretation of the resulting data meets our expectations.

5. Discussion

While we are able to show that the modeling prescriptions outlined in this paper do indeed bring the closure and kernel phase closer to the true self-calibration, it seems that in order to reach the highest contrast detection limits one must always resort to calibration observations, which typically require telescope repointing and this is therefore a time-consuming option. If a target were to exhibit different properties at two nearby wavelengths, such as a strong absorption or emission spectral line caught by one filter and not the other, it seems that a powerful and more efficient calibration scheme would subtract the kernel phase acquired in the two filters from one another. One would then have to fit the thus calibrated data to a spectral differential kernel-phase model.

We have seen that with its stable contrast, GJ 164 AB does not really feature any noteworthy spectral behavior and that filters are reasonably far apart from one another so this data set is not ideal to test this idea. Nevertheless, because of the relative proximity of the *F164N* and *F190N* filters, we can still assess the potential of one such observing mode by looking for

Fig. 14. Correlation plots for the combined five-parameter model fit of the multifilter GJ 164 AB image data set, split back between the different filters (from left to right: *F108N*, *F164N*, and *F190N*). The detection of the companion by the kernel-phase analysis is unequivocal.

Fig. 15. Correlation plot of the kernel-phase residuals after subtraction from the best-fitting binary model. The relatively good match between the two residuals (correlation coefficient ≈ 0.87) suggests that the use of the spectral differential kernel phase would be a valid way to solve the calibration problem.

similarities in the structure of the kernel-phase residuals after subtraction of the best fit model. Figure 15 features a correlation plot of these residuals for the *F164N* and *F190N* filters that include experimental uncertainties. The apparent good correlation between the two residuals suggests that a spectral-differential calibration scheme has some merit: The magnitude of the differential kernel-phase residual is ~ 0.11 rad, which is getting close to the associated experimental dispersion ($\sigma_E = 0.08$ rad). This approach should be further tested on images acquired in two filters that are not as far away or in the analysis of data cubes produced by an AO-fed integral field spectrograph, which will be the object of future work.

6. Conclusion

Kernel phase is a versatile adaptation of the idea of closure phase that can be used in a wide variety of configurations, assuming

that images are reasonably well corrected. With versatility however comes the need for care. The description of the aperture upon which the analysis is made must be thought through, requiring good knowledge of the original pupil, and matched to the constraints brought by images, in particular, the number of useful pixels as well as the scientific ambition.

We have seen that the inclusion of a transmission model for the description of the aperture required to build the pupil-Fourier-phase transfer model brings a major improvement in fidelity. Several examples using ideal numerical simulations and actual data sets from ground-based observations as well as from space have demonstrated that this overall higher fidelity reduces the impact of systematic errors and leads to better results. One should also note that the introduction of the gray transmission model now makes it possible to take advantage of pupil apodization masks used to reduce the contribution of photon noise over a finite area of the image, which, assuming good self-calibration, result in improved contrast detection limits.

Closure- and kernel-phase-based observing programs are becoming more and more ambitious with instruments that make it theoretically possible, in some cases, to probe for planetary mass companions (Sallum & Skemer 2019; Ceau et al. 2019) down to the diffraction limit without a coronagraph. The proper handling of systematic errors in both scenarios is becoming paramount. While efficient calibration procedures offer a way to recover from problematic solutions, the work described here is an attempt to avoid them in the first place.

Acknowledgements. This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (grant agreement CoG # 683029). It is based in part on observations made with the NASA/ESA *Hubble* Space Telescope, obtained from the data archive at the Space Telescope Science Institute. STScI is operated by the Association of Universities for Research in Astronomy, Inc. under NASA contract NAS 5-26555.

References

- Baldwin, J. E., Haniff, C. A., Mackay, C. D., & Warner, P. J. 1986, *Nature*, **320**, 595
- Ceau, A., Mary, D., Greenbaum, A., et al. 2019, *A&A*, **630**, A120
- Dekany, R., Roberts, J., Burruss, R., et al. 2013, *ApJ*, **776**, 130
- Hayward, T. L., Brandl, B., Pirger, B., et al. 2001, *PASP*, **113**, 105
- Hinkley, S., Monnier, J. D., Oppenheimer, B. R., et al. 2011, *ApJ*, **726**, 104

- Ireland, M. J. 2013, *MNRAS*, **433**, 1718
- Jennison, R. C. 1958, *MNRAS*, **118**, 276
- Jovanovic, N., Martinache, F., Guyon, O., et al. 2015, *PASP*, **127**, 890
- Kammerer, J., Ireland, M. J., Martinache, F., & Girard, J. H. 2019, *MNRAS*, **486**, 639
- Kraus, A. L., & Ireland, M. J. 2012, *ApJ*, **745**, 5
- Kraus, A. L., Ireland, M. J., Martinache, F., & Lloyd, J. P. 2008, *ApJ*, **679**, 762
- Krist, J. E., & Hook, R. N. 1997, in *The 1997 HST Calibration Workshop with a New Generation of Instruments*, eds. S. Casertano, R. Jedrzejewski, T. Keyes, & M. Stevens, 192
- Laugier, R., Martinache, F., Ceau, A., et al. 2019, *A&A*, **623**, A164
- Lozi, J., Martinache, F., & Guyon, O. 2009, *PASP*, **121**, 1232
- Martinache, F. 2010, *ApJ*, **724**, 464
- Martinache, F. 2013, *PASP*, **125**, 422
- Martinache, F., Rojas-Ayala, B., Ireland, M. J., Lloyd, J. P., & Tuthill, P. G. 2009, *ApJ*, **695**, 1183
- Pope, B., Tuthill, P., Hinkley, S., et al. 2016, *MNRAS*, **455**, 1647
- Sallum, S., & Skemer, A. 2019, *J. Astron. Telesc. Instrum. Syst.*, **5**, 018001
- Soummer, R., Pueyo, L., Sivaramakrishnan, A., & Vanderbei, R. J. 2007, *Opt. Exp.*, **15**, 15935
- Tuthill, P., Lloyd, J., Ireland, M., et al. 2006, in *Advances in Adaptive Optics II*, eds. B. L. Ellerbroek, & D. Bonaccini Calia, *Proc. SPIE*, **6272**, 2006
- Viana, A., et al. 2009, *Near Infrared Camera and Multi-Object Spectrometer Instrument Handbook for Cycle 17 v. 11.0*

Appendix A: Number of kernels for a symmetric aperture

Given the properties of the expected theoretical number of kernels n_K associated with the model of a circular obstructed

aperture, one can compute an approximate number of expected kernels that is only based on its dimension and the pitch s of the grid used to build the model. In neglecting the effect of spider vanes, the total number of virtual subapertures n_A that fit within a circular aperture of diameter D featuring a central obstruction of diameter d is proportional to the surface area of the aperture: $S_A = \pi \times (D^2 - d^2)/4$. The number of baselines n_B is itself proportional to the surface area of the Fourier coverage, which is half of a disk of diameter $2D$: $S_{UV} = \pi D^2$.

If the aperture is 180° rotational symmetric, we know (Martinache 2013) that the number of kernels n_K yielded by the model is $n_K = n_B - n_A/2$. We can therefore compute the ratio of the number of kernels and the number of baselines to find a number that only depends on the geometric properties of the aperture $r = n_K/n_A = 1 - 0.25 \times (1 - (d/D)^2)$. This makes it possible to evaluate the kernel efficiency of a circular aperture, which is on the order of 77% for the SCExAO pupil case discussed in the body of the paper.

Given the step s , one can estimate how many virtual baselines fit within the Fourier coverage, $n_B = S_{UV}/s^2$, and predict the number of kernels one can build, using the ratio determined above, $n_K = r \times n_B$. Applications for the reference model ($s = 0.42$ m) predict $n_K = 430$ kernels, and for the finer model ($s = 0.21$ m), $n_K = 1719$, which are reasonably close to the true values listed in Table 1.