



Kernel Phase and Coronagraphy with Automatic Differentiation

Benjamin J. S. Pope^{1,2,3,4,5,8} , Laurent Pueyo³ , Yinzi Xin⁶ , and Peter G. Tuthill⁷

¹ Center for Cosmology and Particle Physics, Department of Physics, New York University, 726 Broadway, New York, NY 10003, USA; benjamin.pope@nyu

² Center for Data Science, New York University, 60 Fifth Avenue, New York, NY 10011, USA

³ Space Telescope Science Institute, 3700 San Martin Drive, Baltimore, MD, 21218, USA

⁴ School of Mathematics and Physics, The University of Queensland, St Lucia, QLD 4072, Australia

⁵ Centre for Astrophysics, University of Southern Queensland, West Street, Toowoomba, QLD 4350, Australia

⁶ Department of Aeronautics and Astronautics, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

⁷ Sydney Institute for Astronomy (SIfA), School of Physics, The University of Sydney, NSW 2006, Australia

Received 2020 August 9; revised 2020 October 22; accepted 2020 November 17; published 2021 January 25

Abstract

The accumulation of aberrations along the optical path in a telescope produces distortions and speckles in the resulting images, limiting the performance of cameras at high angular resolution. It is important to achieve the highest possible sensitivity to faint sources, using both hardware and data analysis software. While analytic methods are efficient, real systems are better modeled numerically, but numerical models of complicated optical systems with many parameters can be hard to understand, optimize, and apply. Automatic differentiation or “backpropagation” software developed for machine-learning applications now makes calculating derivatives with respect to aberrations in arbitrary planes straightforward for any optical system. We apply this powerful new tool to the problem of high-angular-resolution astronomical imaging. Self-calibrating observables such as the “closure phase” or “bispectrum” have been widely used in optical and radio astronomy to mitigate optical aberrations and achieve high-fidelity imagery. Kernel phases are a generalization of closure phases valid in the limit of small phase errors. Using automatic differentiation, we reproduce existing kernel phase theory within this framework and demonstrate an extension to the case of a Lyot coronagraph, which is found to have self-calibrating combinations of speckles, which are resistant to phase noise, but only in the very high-wave-front-quality regime. As an illustrative example, we reanalyze Palomar adaptive optics observations of the binary α Ophiuchi, finding consistency between the new pipeline and the existing standard. We present a new Python package MORPHINE that incorporates these ideas, with an interface similar to the popular package POPPY, for optical simulation with automatic differentiation. These methods may be useful for designing improved astronomical optical systems by gradient descent.

Unified Astronomy Thesaurus concepts: [Direct imaging \(387\)](#); [Astronomy data analysis \(1858\)](#); [Optical interferometry \(1168\)](#); [Coronagraphic imaging \(313\)](#); [Astronomical simulations \(1857\)](#)

1. Introduction

Many questions in astronomy and other sciences can only be answered with diffraction-limited high-resolution imaging. The highest resolutions are typically achieved with the method of interferometry, in which waves detected at multiple receivers are combined physically or in postprocessing to obtain the Fourier transform of the source intensity distribution (van Cittert 1934; Zernike 1938). Even in the case of a single telescope or camera, it is often nevertheless helpful to think of them as an interferometer composed of many subapertures, which combine their signals directly onto a focal plane (a “Fizeau interferometer”: Fizeau 1868). In all of these cases, unknown path delays through each receiver or subaperture are the dominant source of aberration with low spatial frequency structure causing issues such as the varying “tip tilt” of the final image position, while higher-order distortions yield clouds of “speckles”.

Strategies to correct for these aberrations include the use of active systems such as adaptive optics or delay lines; calibrating these errors with comparison to a reference star; or self-calibration, using the physics of the noise process to correct for it in postprocessing. In the recovery of object phase information, among the most longstanding techniques is that of self-calibration with “closure phases” (introduced in the context of radio

astronomy by Jennison 1958), in which phases are summed around three interferometric baselines that form a closing triangle. The phase error terms local to each subaperture cancel, so that three low-signal-to-noise (S/N) baseline phases deliver one high-S/N observable. In a nonredundant array (one in which no baseline vector is repeated between different pairs of subapertures) of sufficient size, a large number of closure phases that anchor image reconstruction with great precision can be obtained (Chael et al. 2018). “Closure amplitudes,” which are resistant to fluctuations in input amplitude or gain, can also be obtained by a similar construction using four telescopes (Twiss et al. 1960; Blackburn et al. 2020). The closure phase is the argument of a quantity called the triple product or “bispectrum”, and it has recently been shown from the perspective of invariant theory in algebraic geometry that for a wide class of problems limited by phase noise, knowledge of the mean of a signal, its power spectrum, and its bispectrum is necessary and sufficient for optimal signal reconstruction (Bandeira et al. 2017).

On the other hand, for the case of direct imaging of high-contrast companions with coronagraphs, analytic self-calibrations such as those above are not yet known, and external calibration is necessary. The standard approaches to data analysis in coronagraphy rely on exploring a diversity of point-spread functions (PSFs) experimentally, constructing a basis covering some of the diversity in speckle patterns, and then subtracting out a linear combination of vectors in this basis from

⁸ NASA Sagan Fellow.

measured data (e.g., KLIP; Lafrenière et al. 2007; Soummer et al. 2012; Pueyo 2016). Advantages can also be gained from angular differential imaging (ADI; Marois et al. 2006) and spectral deconvolution with wavelength (Sparks & Ford 2002). In this paper, we will extend our understanding of analytic self-calibration to better include coronagraphs so that we can add an additional layer of precision to exoplanet imaging calibration.

1.1. Kernel Phases

The kernel phase method is a way of extending the closure phase from simple nonredundant arrays of telescopes to the densely filled pupils of real telescopes. If we describe the propagation of phase and amplitude noise in terms of a matrix (Lannes 1991), we can obtain powerful generalizations to closure phases and amplitudes. We will describe the effect of aberrations on baseline phases as

$$\Phi = \Phi_0 + A_\phi \cdot \phi, \quad (1)$$

where Φ is a vector of observed phases on each baseline in the u, v (focal-plane Fourier domain), Φ_0 the true astrophysical phases, ϕ the phase noise at each point in the pupil plane, and A_ϕ a transfer matrix from the pupil plane to the u, v plane phases. A similar expression can be written for amplitudes (Pope 2016). In the redundant pupil case such as a standard full telescope pupil, where each baseline appears more than once, the propagation of phase noise from the pupil to baselines is no longer linear, but in the limit where aberrations are small, it can be linearized. For sufficiently small phase perturbations ($\ll 1$ rad), error propagation to the focal-plane Fourier phases is approximately linear: this is simply a Taylor expansion to first order of the nonlinear phase transfer operation, in which A_ϕ is the Jacobian matrix of partial derivatives $\partial\Phi_j/\partial\phi_k$. For a single focal-plane imaging system, this operator A_ϕ can be determined analytically for a given discrete pupil model.

Martinache (2010) introduced the idea of “kernel phase” to generalize closure phase to a redundant aperture. While in this case closure phases no longer exactly cancel out the contributions of aberrations ϕ , nevertheless, there is an equivalent that can be found numerically. A left kernel operator K can be calculated via singular value decomposition (SVD) that annihilates A_ϕ , such that

$$K \cdot A_\phi = 0, \quad (2)$$

and, therefore, we can find self-calibrating kernel phases $K \cdot \Phi$, which are immune to phase noise to linear order:

$$K \cdot \Phi = K \cdot \Phi_0 + \cancel{K \cdot A_\phi \cdot \phi} \quad (3)$$

$$= K \cdot \Phi_0. \quad (4)$$

Meanwhile, the phases in the space complementary to the kernel space have the opposite property—they are especially sensitive to input phase aberrations, and for appropriate apertures can be used for wave-front sensing from the image domain (Martinache 2013; Pope et al. 2014a; Martinache et al. 2016).

The kernel phase method has been applied to space-based data from the Hubble Space Telescope (HST) NICMOS camera (Pope et al. 2013; Laugier et al. 2019; Martinache et al. 2020), ground-based data from the Palomar 200-inch adaptive optics (AO) equipped Pharo camera (Pope et al. 2016; Martinache et al. 2020), the Large Binocular Telescope (Sallum et al. 2015), the

VLT/NACO camera (Kammerer et al. 2019), and MagAO (Sallum et al. 2019), and explored theoretically for ground- and space-based telescopes (Ireland 2013; Martinache 2011; Ceau et al. 2019; Sallum & Skemer 2019).

The kernel phase method relies on a Taylor expansion of the optical propagation, which for a simple telescope is easy to do analytically. But many systems in reality cannot be treated this way—for example, problems involving diffraction between multiple planes like in coronagraphs. We shall see in the following that automatic differentiation can supply derivatives for arbitrary numerically-simulated imaging systems, extending self-calibration to a wider class of instruments and also offering new opportunities for optical design and optimization.

1.2. Automatic Differentiation

The practical use of neural networks in machine-learning applications is dependent on the efficient calculation of analytic gradients of often very complicated composite functions, for example, the matrix operations composed of nonlinear activation functions that are seen in neural networks (LeCun et al. 2015). This problem is usually referred to as algorithmic differentiation, automatic differentiation, or “autodiff”, and is solved simply by the chain rule. Autodiff is available in many implementations, such as the Python packages TENSORFLOW (Abadi et al. 2015), THEANO (Theano Development Team 2016), PYTORCH (Paszke et al. 2019), AUTOGRAD (Maclaurin et al. 2015), JAX (Bradbury et al. 2018), and many packages in the Julia language (Bezanson et al. 2012). The reverse-mode autodiff or “backpropagation” algorithm (Linnainmaa 1970; LeCun et al. 1988) has made this practical for neural networks, but it is typically most useful in cases where the output dimensionality is much smaller than the input, as for most neural networks. Forwards-mode autodiff, on the other hand, is usually better in the case where the output is of higher dimension than the input. For optics problems, both cases can be found—for example, for optimization, backpropagation is usually the best approach, while for kernel phase analysis, forwards mode is better suited.

In simulating physical optics, we normally consider the complex electric field on a 2D plane, which is pixelized and then flattened to a 1D vector. Optical propagation then consists of a series of Fourier (or Fresnel) transforms mapping between planes, and matrix or element-wise multiplications by phase and amplitude screens in those planes. Indeed, for practical purposes in astronomy and imaging science generally, optical propagation through a whole system is generally a linear operation that could be written as a single (large) matrix multiplication. However quantities of interest are usually not specified as real and imaginary electric field components, but rather as amplitudes and phases, so that the relations between (for example) input and output phases are in general nonlinear. Finding the derivatives of the u, v phase with respect to pupil aberrations is therefore an ideal problem for autodiff.

The analogy between the operations of optics and deep neural networks is so strong that not only have autodiff packages been used to simulate photonic systems (e.g., Hughes et al. 2018), photonic systems have in fact been used as analog computers implementing neural networks for machine learning (Guo et al. 2019; Hughes et al. 2019).

Several groups have applied autodiff to areas of optical science relevant to astronomy. Autodiff has been applied fruitfully to geometric optics or ray tracing (e.g., Werner et al. 2012; Sutin 2016), which is important in astronomy for understanding

gravitationally lensed systems. Chianese et al. (2020) and Morningstar et al. (2019) have applied this method to integrating differentiable forwards models of gravitational lensing with neural networks for image analysis. Autodiff methods have also been applied for image reconstruction from interferometric data, including gravitationally lensed systems (Morningstar et al. 2018) and protoplanetary disks (Czekala 2019).

The DEEPOPTICS project (Sitzmann et al. 2018) has used autodiff to optimize “computational cameras”. Building their model in TENSORFLOW, they couple a physical optics simulation, detector simulation, and a deconvolution postprocessing stage for a total end-to-end imaging simulation. Where normally you might optimize some heuristic of the PSF (such as the FWHM), this makes it possible to jointly optimize hardware and software with respect to figures of merit of the overall system such as final resolution or depth of field. Some designs arrived at in this way are exotic: for example, superresolution is achieved by finding a lens with three off-center Fresnel lens components, focusing on three separate spots, which the deconvolution stage shifts and stacks. Because it is built in TENSORFLOW, the diffraction simulation can be incorporated as a “physical layer” in neural network applications in microscopy, for example for optimizing hardware and software for image classification (Muthumbi et al. 2019), or with reinforcement learning for adaptively learning sample illumination (Chaware et al. 2019).

An approach similar to DEEPOPTICS is likely to be extremely valuable in designing, for instance, pupil masks for coronagraphy (e.g., Guyon 2003; Carlotti et al. 2011) or for diffractive-pupil astrometry (e.g., Guyon et al. 2012; Tuthill et al. 2018).

Close to the topic of this paper, autodiff has been applied to the problem of phase retrieval (Jurling & Fienup 2014; Paine & Fienup 2019), inferring a wave front from a PSF. In this context, a key advance of autodiff over previous methods is that we can trivially account for pixel sampling/binning and detector nonlinearity. These will be important issues when considering that the *JWST* mid-infrared imager MIRI will have significant detector nonlinearity (Rieke et al. 2015), or where we may wish to look at saturated sources. While they do not address more complex optical systems, this method may be straightforwardly applicable to sensing non-common-path errors in coronagraphic images. It may also be helpful to integrate phase retrieval simultaneously with flat-field calibration and light-curve extraction in photometric missions such as Kepler, K2, TESS, and their successors, combining the gradient descent flat-field optimization in halo photometry for bright stars (White et al. 2017; Pope et al. 2019) and PSF optimization for fainter stars in crowded fields (e.g., Feinstein et al. 2019; Nardiello et al. 2019).

While Jurling & Fienup (2014) derive analytic expressions for many optically relevant operations, machine-learning software that has become available since then have significantly widened the range of options and introduced more user-friendly APIs. These approaches have been taken up outside of optical imaging, for example, in X-ray-coherent diffractive imaging (Kandel et al. 2019; Nashed et al. 2019) and nanotomography (Du et al. 2020).

Rather than using autodiff gradients for optimization, in this paper we use them to understand optical systems and the information they propagate. In the following, we will show how autodiff can reproduce the existing state of the art in kernel

phase and demonstrate a way forward using autodiff to extend the kernel phase idea to coronagraphic instruments.

2. Simulations

To forwards-model systems, we adapt the popular physical optics library POPPY (Perrin et al. 2012) with the goal of compatibility with existing simulations built on POPPY such as WEBBPSF (Perrin et al. 2012). We use the matrix Fourier transform mode and avoid fast Fourier transforms (FFTs; Soummer et al. 2007) to more easily generate arbitrary image sampling and for consistency with the kernel phase code XARA (Martinache et al. 2020). We have adapted the low-level features of POPPY to use the Google autodiff library JAX (Bradbury et al. 2018) in place of NUMPY, and to distinguish it from the original version, we call this new POPPY derivatives library MORPHINE.

This has several advantages already over analytic methods for the kernel phase. MORPHINE can calculate monochromatic or polychromatic PSFs: the polychromatic capability allows us to explicitly construct broadband kernel phase operators. It is also possible to take derivatives with respect to wave fronts specified in bases other than the pixel basis, for example, Zernike or hexike modes, as we shall discuss in Section 2.5.

2.1. Simple Diffraction Example

First we want to ensure basic kernel phase calculations are reproduced in this new model. We propagate monochromatic $2.0\,\mu\text{m}$ light through a $2.0\,\text{m}$ diameter circular pupil that is imaged onto a detector with a $20\,\text{mas pixel}^{-1}$ scale and $4''$ field of view. The pupil plane and the baseline u, v plane are each calculated on a 64×64 grid. We evaluate the Jacobian of the u, v phases with respect to the input pupil phases using forwards-mode autodiff and then evaluate this Jacobian at zero phase to get our A_ϕ matrix. On a laptop computer, this takes a few minutes of processing time. Memory usage is the main overhead: with 16 GB RAM, grid size is limited to about 128×128 . Demand on computational resources presents a major limitation to widespread application; in the future, this may be overcome by more efficient refactoring of the code or simply by doing calculations on clusters with very large RAM allocations.

Although it is inefficient to calculate this full Jacobian with respect to a densely pixelized basis for the wave front in the pupil plane, we do so here for completeness. In practical cases, aberrations are typically dominated by low-order modes, and we could consider the Jacobian only with respect to (for example) some much smaller number of Zernike modes, or pixels in the u, v plane. Reducing the input dimensionality like this would have significant advantages in memory usage for forwards-mode autodiff.

In Figure 1, we display the results of our Jacobian calculation for a set of highlighted pupil samples and illustrating the influence of phase in the u, v plane. The results are similar to Figure 1 of Martinache (2010), consisting of the sums of translated positive and negative copies of the pupil aperture. Phase offsets outside the support of the pupil have identically zero derivative in the u, v plane as expected. There is a greater magnitude in the derivative at longer baselines compared to the flat maps in Martinache (2010) because we are rolling together the phase transfer A matrix and the redundancy




Figure 1. A representation of the Jacobian determined by JAX for monochromatic $2\ \mu\text{m}$ diffraction from a $2.0\ \text{m}$ diameter circular pupil. In each pair of images, on the left is the pupil, with a red dot highlighting a selected pixel; on the right is a map of the derivative of the phase in the u, v plane with respect to the phase at that point in the pupil. These Jacobian vectors are pairs of rows and columns of the A phase transfer matrix. We see the same maps as in Figure 1 of Martinache (2010).

matrix R , and short baselines are much more redundant than long ones.

We calculate the singular values of this matrix (restricted to the support of the pupil) by SVD (Figure 2), finding a sharp cutoff as expected, separating a subspace of kernel phases from phases contaminated by noise from the aberrations. In the nonsingular space, a selection of pairs of corresponding modes in the pupil and u, v plane are displayed in Figure 3, and a selection of kernel phase modes in the u, v plane are shown in Figure 4. These illustrate how at the edges of the pupil/long baselines there is extra noise. Because of the limited field of view, noise from outside the pupil or optical transfer function (OTF) is convolved inside, and large elements in the transfer matrix can dominate parts of the SVD. For numerical stability we have therefore excluded the very longest baselines in calculating kernel phases, but it would be preferable to use a more robust SVD that is tolerant of outliers. This is not a feature of the autodiff pipeline or MORPHINE, but also of real

data. Depending on the field of view of the simulated images, for the same pupil and u, v sampling, there can be slight differences in the autodiff-calculated kernel phases because of this effect. The normal analytic kernel phase derivation assumes an infinite field of view, but sees data with a finite field of view in which these window effects are apparent.

We have also repeated the above calculations for light with a 20% fractional bandwidth, with a uniform spectrum sampled ten times from $1.8\ \mu\text{m}$ – $2.2\ \mu\text{m}$. The results are very similar, except that the Jacobian maps equivalent to Figure 1 are slightly blurred. In the shape of the singular value curve between the broadband and monochromatic cases, a qualitatively similar behavior is found, although the broadband case has slightly fewer kernel phases and a flatter roll-off.

2.2. Stellar Imaging Data: α Ophiuchi

We now apply our new kernel phase formalism to real data: the A-type binary star system α Ophiuchi observed with the




Figure 2. Log-scale plot of the singular values of the phase transfer matrix plotted against the index ordered by decreasing singular value for simple monochromatic diffraction (blue), 20% bandwidth (green). Singular values are normalized to the first singular value. In both cases, there is a sharp cutoff separating significant singular values from a plateau at $\sim 5 \times 10^{-8}$ (zero to within machine precision). We take the vectors corresponding to these singular values in the plateau region to be the effective null space of the phase transfer operator and correspond to kernel phases.




Figure 3. Pairs of nonsingular vectors in the pupil and u, v planes from the SVD of the Jacobian in Figure 1. Some of these orthonormal basis vectors resemble Zernike modes, and all show some ringing structure at the edges of the pupil and OTF support. All eight images are dimensionless and the color scale is arbitrary.

Palomar 200-inch telescope using the extreme adaptive optics system PALM-3K and the PHARO camera. These data were first analyzed by Pope et al. (2016), successfully recovering the binary and properties of the system. These were subsequently improved by Martinache et al. (2020), who corrected a 2° offset in the clocking angle of the real Palomar pupil. Also employing the updated “XARA” kernel phase code, Martinache et al. (2020) obtained astrometry of the separation $\rho = 123.5 \pm 2.9$ mas, position angle $\theta = 86.5^\circ \pm 0.2^\circ$, and contrast $c = 25.1 \pm 1.1$. The well-analyzed nature of these data makes it something of a

useful test piece; we therefore seek to recreate this analysis using our new autodiff methods.

We first build a high-resolution model of the PHARO pupil using the same code as Martinache et al. (2020) and bin it down to a 64×64 pixel grid so that pixel values $\in [0, 1]$ indicate the fraction of the telescope aperture filling that binned pixel. We then propagate monochromatic $2.145 \mu\text{m}$ light through this onto a 128×128 pixel image plane and use a matrix Fourier transform to map this onto a u, v plane of the same size. We differentiate this using the forwards-mode autodiff as above,




Figure 4. Example kernel phase maps—singular vectors in the u, v plane from the SVD of the Jacobian in Figure 1. These are rows of the \mathbf{K} matrix. There is little apparent structure. All 16 images are dimensionless and the color scale is arbitrary.




Figure 5. A representation of the Jacobian determined by JAX for the PHARO pupil in the K band. In each pair of images, on the left is the PHARO pupil, with a red dot highlighting a pupil sample; on the right is a map of the derivative in the u, v plane with respect to phase at that pupil sample. These Jacobian vectors are pairs of rows and columns of the \mathbf{A} -phase transfer matrix. We see the same maps as in Figure 1 of Martinache (2010).

obtaining the Jacobian shown in Figure 5, and calculate a kernel phase transfer matrix.

We then extract kernel phases from both α Oph and a point-source calibrator ϵ Her using a version of XARA modified for compatibility with MORPHINE. Following Martinache et al. (2020),

we choose as our observables the median kernel phases across all 100 exposures for each and the standard error of the mean as our base uncertainty on each. To calibrate the kernel phases, we simply subtract those of the calibrator from α Oph and add their uncertainties in quadrature. To see the effect of the known optical

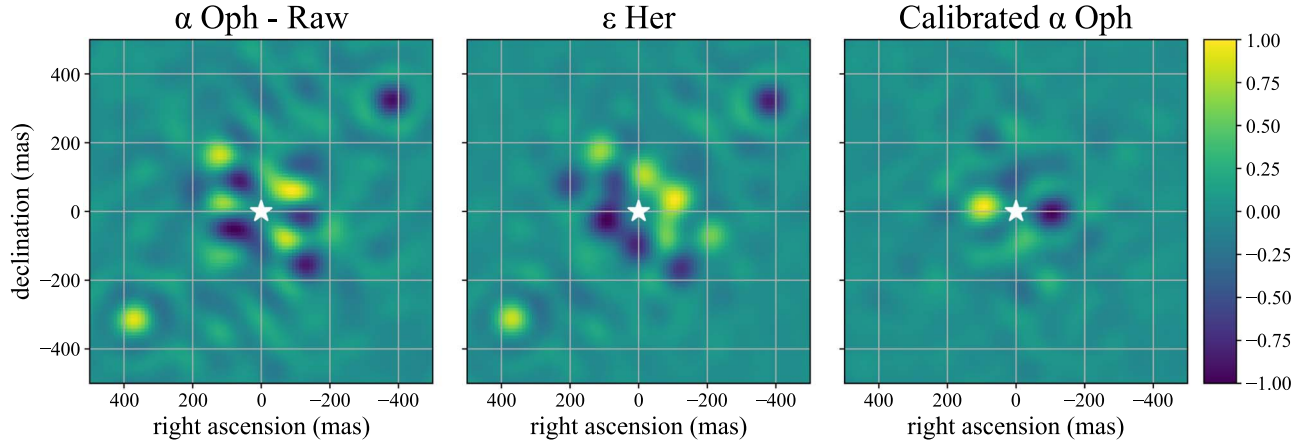


Figure 6. Kernel phase colinearity maps showing the normalized dot product of a binary model (pale is a better match) to the kernel phase signals from uncalibrated α Oph data (left), the calibrator ϵ Her (middle), and calibrated data (right). They closely resemble the maps in Figure 8 of Martinache et al. (2020). We can see that in uncalibrated data the filter ghost is clearly visible, but in calibrated data we extract the α Oph binary with no effect from the ghost.

“ghost” (due to an unwanted reflection from the neutral density filter used to observe bright stars such as these), we calculate kernel phase “colinearity” maps as in Martinache et al. (2020), shown in Figure 6. In these maps, we can see that the filter ghost shows up strongly as a false binary in the uncalibrated kernel phase maps, but calibrated kernel phases easily remove this and reveal the tight α Oph binary.

In Martinache et al. (2020), a large error term was added in quadrature in order to account for additional noise in the data unexplained by the diversity over exposures and uncalibrated by the kernel phase model. Using nonlinear least squares to fit the calibrated data, we also find that the best-fitting binary model has a high χ^2 . We therefore follow Martinache et al. (2020) and add a dimensionless kernel phase quantity of 0.0184 in quadrature to the uncertainties to make the reduced $\chi^2 = 1.0$.

We then use Markov Chain Monte Carlo (Metropolis et al. 1953) to sample from the posteriors assuming Gaussian likelihoods, using the `emcee` affine-invariant ensemble sampler (Foreman-Mackey et al. 2013). We obtain binary parameters of separation $\rho = 120.2 \pm 1.2$ mas, position angle $\theta = 86.2 \pm 0.3$, and contrast $c = 24.4 \pm 0.6$: all well within 1σ of the Martinache et al. (2020) figures. We likewise use `emcee` to sample from the likelihood using the Martinache et al. (2020) code. Posteriors for both inferences are shown in Figure 7.

2.3. Simulated Binary Star Data

In order to investigate the effect on both pipelines of wave-front errors leaking in, we generated a cube of 1000 PSFs at a wavelength of $2.145 \mu\text{m}$ with 20 Zernike modes of aberrations, this time with a 20 nm standard deviation, for a total peak-to-peak wave-front amplitude of 295 ± 100 nm. We then retrieved the system parameters with `emcee` for the cube frame by frame using weighted least squares (displayed in Figure 8). For a simulated binary with fiducial $\rho = 125$ mas, $\theta = 86^\circ$, and $c = 25$, there is no sign of the large systematic difference between the methods as seen in real data. Both methods, whether frame by frame or on average, retrieve the true parameters with similar accuracy and precision.

In order to establish the accuracy of both pipelines, we conducted an injection test. We took the first frame of the ϵ Her data cube as our “data” and use a subpixel Fourier shift to create a mock binary at a position angle of 90° , contrast 25, and

a range of separations from 90–300 mas. We then used least squares to fit these data using both models (initialized at the true values), determined in each case an error to add in quadrature to make reduced $\chi^2 = 1.0$, and reran the least-squares fit with new data, saving the best-fit parameters and uncertainties taken from the square root of the diagonal elements of the inverse covariance matrix. The results are shown in Figure 9. Both methods perform comparably well, achieving results reasonably close to the injected values. Systematics in all three parameters, of amplitude comparable to the 1σ statistical uncertainties, vary smoothly and similarly as a function of input separation. The autodiff pipeline performs somewhat better at smaller separations in all three parameters. This indicates that both kernel phase pipelines do a good overall job at binary imaging with PHARO, but that as noted by Martinache et al. (2020), not all sources of uncertainty are well understood or calibrated.

2.4. Complex Optical Systems: Nullers and Coronagraphs

It is desirable to find a priori self-calibration schemes for high-angular-resolution, high-contrast imaging systems more generally such as nulling interferometers and more optically complex coronagraphs intended to suppress light in a region around a target star. One particularly interesting extension of the kernel phase idea is the “kernel nulloer” concept (Martinache & Ireland 2018), in which the aberration transfer matrix and kernel operator idea is applied to the output of a nulling interferometer.

For arbitrary systems, to generate our Taylor expansion, we wish to be able to take the Jacobian of arbitrary features of the final speckle pattern with respect to phases in the input wave front or in intermediate planes. This could then be used in postprocessing to generate “kernel speckles” resistant to noise, or in wave-front control for adaptive optics, or for digging “dark holes” to search for high-contrast companions (Malbet et al. 1995).

Arbitrary-order Taylor series expansions for a PSF with respect to small phase perturbations can be derived analytically for a telescope with an arbitrary pupil brought to a single focus (Bloemhof et al. 2001; Sivaramakrishnan et al. 2002; Perrin et al. 2003). This expansion breaks down for more complex optical systems. The PSF of a coronagraphic imager far from




Figure 7. Posterior distributions for the binary parameters of α Ophiuchi as determined by the autodiff pipeline (black) and the Martinache et al. (2020) implementation (blue).

the occulted region is very similar to that of the primary PSF of the telescope, but close to the inner working angle, it is significantly affected by the occultation and especially strongly by low-order wave-front errors. This is likewise an issue where detector nonlinearity and saturation significantly distort the PSF. While the analytic series expansion is no longer applicable in these regimes, using matrix Fourier transforms and skipping saturated pixels or pixels inside the coronagraphic inner working angle has shown promise for pushing the kernel phase beyond its conventional limitations (Laugier & Martinache 2019), as have generalizations that exploit angular differential imaging for a further level of calibration (Laugier et al. 2020).

Here we will apply the formalism demonstrated above for simple diffraction to a new case: coronagraphy. We will find that the identical approach generates an analogous modal decomposition and kernel observables.

We consider a very simple Lyot coronagraph (Lyot 1930). We choose a $1\mu\text{m}$ wavelength and a 1 m diameter pupil sampled on a 96×96 grid. The central $4\lambda/D$ region of the resulting PSF is blocked out with an on-axis focal-plane stop (occulter), and the light is then propagated to a second pupil plane using an FFT rather than a matrix FT, following standard POPPY practice. This imposes tougher memory constraints and

in the future may be replaced with a Matrix Fourier transform. A Lyot stop (an iris mask undersized 10% relative to the input pupil) is then imposed, and the light propagated to a final focal plane sampled on a 50 mas pixel scale grid with a $4''$ field of view. Aside from the coarse gridding, this is intentionally identical to one of the standard test and verification examples supplied with POPPY.⁹ The PSF produced has a dark coronagraphic hole in the middle, and beyond an inner working angle around this hole has a pattern of diffraction rings.

Again using JAX, we calculate the Jacobian of this PSF with respect to the input wave front, evaluated at uniform-zero phase, which again takes a few minutes on a laptop. The previous approach has been to directly propagate through the end-to-end simulation for a grid of small perturbations in the input plane (Riggs et al. 2018), whereas this is now taken care of by autodiff. The results are displayed in Figure 10 in a similar format to Figure 1. Pixels far out in the PSF correspond to sinusoidal modulation of the input wave front as expected, and this is symmetric about the origin. On the other hand, pixels near the inner working angle where the effects of the

⁹ github.com/mperrin/poppy/blob/master/notebooks/MatrixFTCoronagraph_demo.ipynb

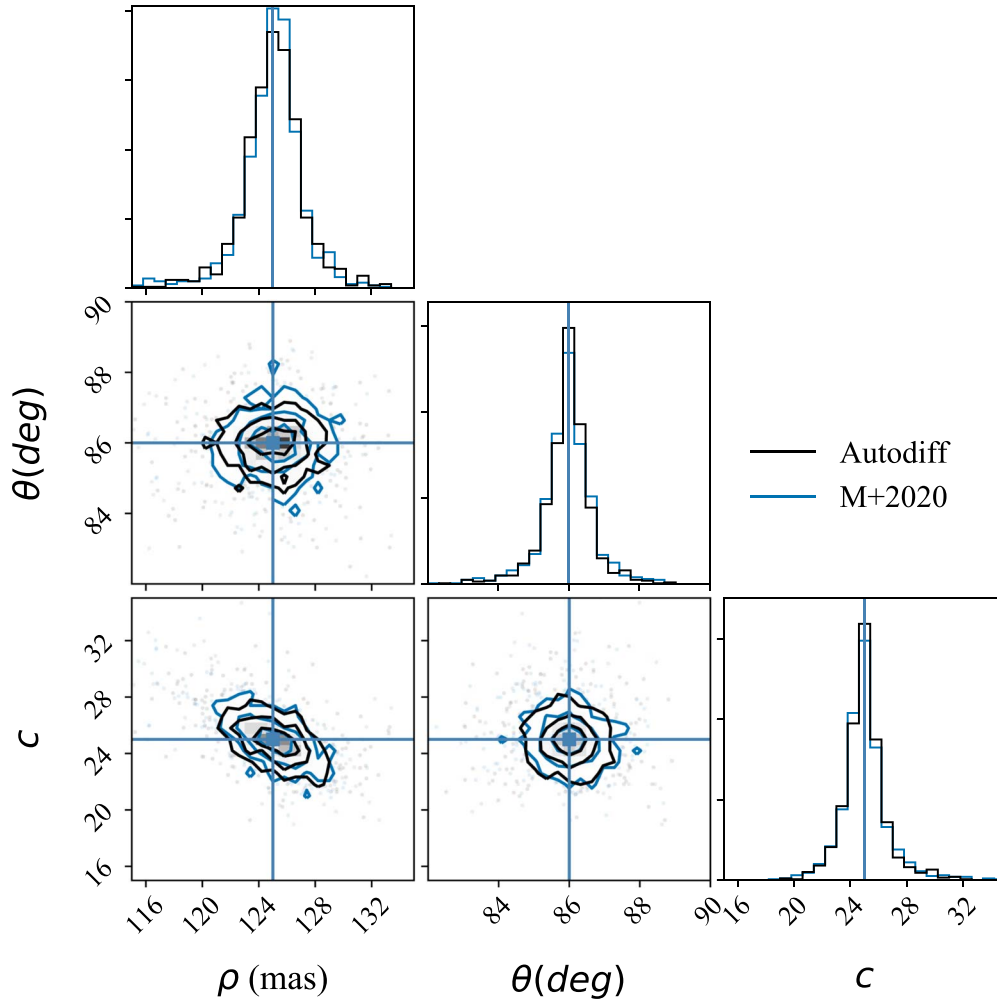


Figure 8. Comparison of Martinache et al. (2020, blue linetype) and autodiff (black) kernel phase pipelines on simulated data with true parameters $\rho = 125$ mas, $\theta = 86^\circ$, $c = 25$. We generate a cube of 1000 mock data sets using a 64 pixel pupil model and Zernike aberrations as described in Section 2.2. The corner plot displays the parameters retrieved using a simple least-squares fit applied to each individual frame. The methods have very similar distributions centered tightly on the true value.

coronagraph are evident correspond to more complex modes in the pupil, also containing contributions from the ring obscured by the Lyot stop when the pupil is reimaged.

Calculating the SVD of this Jacobian (again restricted to the support of the primary pupil), we see a decline to very small singular values. The ordered singular value curve is displayed in Figure 11. The nonsingular modes form a basis, like the Zernike basis, that maps between aberration patterns in the pupil and the speckle modes they generate. Some pairs of nonsingular modes in the pupil and image plane are shown in Figure 12.

For the coronagraph under consideration, the regime over which the linear approximation applies may be strictly limited. To illustrate this, we take one of the sinusoidal ripples that are the columns of the Jacobian in Figure 10, propagate this through a full optical simulation, and take the difference with respect to an unaberrated PSF. In Figure 13, we see the effects of a sinusoidal phase ripple in the nearly exact linear regime with a 0.1 nm amplitude, and with a small but nonnegligible 5.0 nm amplitude. In the linear regime (where the phase ripple is very small), we see that the speckle is positive on one side of the PSF and negative on the other, corresponding closely with the expected behavior from the Jacobian. On the other hand,

increasing the amplitude of the phase ripple to just 5 nm is enough to alter the outcome completely. The figure indicates a switch in behavior so that the speckles are now both positive, indicating the quadratic term in the Taylor series has become dominant. This quadratic term involves a rank 3 tensor of the second partial derivatives of each pixel with respect to the wave front—and there is no equivalent to a kernel operator for tensors of this dimension. The use of kernel phases as robust observables will therefore fail in this case, though we suggest that future work may search for a higher-dimensional kernel manifold.

We have therefore shown that under the same linear assumptions previously applied to yield kernel phases, analogous self-calibrating observables also exist for propagation through more complex optical systems such as coronagraphs. However, for the example configuration explored, the rapid onset of a nonlinear response indicates they may be of limited practical applicability to coronagraphic observations except in the extremely high-wave-front-quality regime. Subnanometer wave-front precision may nevertheless be achieved in future large space telescopes with high-order deformable mirrors.




Figure 9. Injection recovery results for PHARO using both the Martinache et al. (2020) and autodiff kernel phase models. From top to bottom, we show the error in the recovered separation, the recovered position angle, and the recovered contrast as a function of injected separation, with 1σ uncertainty error bars. Both methods show similar performance, with small systematic errors of comparable size to their statistical uncertainties.

2.5. Zernike Basis

Taking the Jacobian of one large pixel grid by another is expensive in terms of memory; it is therefore desirable to bring down the dimensionality of either the input or target space. One way to do this is to use a Zernike basis to represent the wave front (Zernike 1934). One advantage of this is that we can examine only up to a certain order in Zernike polynomials if we believe that wave-front errors are negligible below a certain spatial scale. We can then differentiate just with respect to Zernike coefficients, which is much more computationally efficient and allows us to run simulations on grids that are many times larger than otherwise. Similar bases (e.g., the “hexike” basis) can be constructed for arbitrary noncircular telescope apertures.

We have rerun the simple diffraction simulation from Section 2.1 with 200 Zernike modes on a 256×256 grid, and the coronagraph from Section 2.4 on a 128×128 grid using 300 Zernike modes. The calculations now take around 4 s, rather than many minutes, with peak memory use around 1.5 and 3 GB rather than 7 and 9 GB respectively.

The Jacobian derived from this basis for simple diffraction is an excellent approximation to the full calculation on a grid. Given that in problems such as HST data analysis we believe that PSF variations are dominated by low-order modes (Lallo et al. 2006), this allows us to calculate a larger kernel space

with a more accurate pupil model than has previously been possible.

While it is a huge advantage in simulation speed and fidelity, this basis does not accurately represent the high-spatial-frequency sine waves corresponding to speckles far from the core of the PSF. Nevertheless, it does an adequate job representing the effects of low-order wave-front error, and this may be sufficient for many purposes, for example in optimization of optical design.

3. Discussion

A phase transfer matrix constructed analytically with a discrete pupil model contains no a priori information about the pixel sampling, binning, or windowing. This means that the implicit convolutions in the u, v plane, and therefore the associated correlation between adjacent baselines, are not taken into account. Therefore, the normally orthonormal kernel phases (calculated analytically for an infinite field of view and fine sampling) are no longer fully linearly independent. For example, Martinache et al. (2020) note that in the analysis of Pope et al. (2016), windowing to avoid the ghost introduced by a neutral density filter cuts the number of pixels to fewer than the number of kernel phases, so that information is being lost. Martinache et al. (2020) also demonstrates the utility of a detailed model of pupil (mis)alignment to improve the extraction of both kernel phase and visibility information.

The nonindependence of kernel phases has been treated statistically by Ireland (2013), who uses principal component analysis applied to an ensemble of kernel phase observations to extract statistically independent kernel phases. When the advantage of kernel phase is knowing that some combinations of phases have high S/N a priori, it would be preferable to avoid this situation with a better kernel phase construction rather than calibration.

A further issue is that the u, v -plane calculation of earlier kernel phase work suffers from discretization noise from interpolating the FFT. This can be ameliorated by using a matrix Fourier transform adapted to the wavelength and exact pupil model (Martinache 2018).

By calculating kernel phases from differentiating an end-to-end optical simulation, even for trivial apertures, the exact pixel grid, including binning, windowing, and dead pixels can be included. As the matrix DFT is already used to calculate the FT, there is no sampling noise to separately incorporate, and all kernel phases are automatically orthonormal.

To avoid self-subtraction, we expect that the best approach to data analysis with kernel phases using these methods will be forwards-modeling the data through the optical system and kernel phase filter along the lines described by Pueyo (2016) or Martinache et al. (2020). These can then be augmented by diversity-based methods along the lines described by Ireland (2013) to achieve even higher levels of calibration. We have not attempted to use these methods on real data in this paper.

Because of its NUMPY-like API, we have used JAX to power the automatic differentiation in MORPHINE. We have not extensively optimized this code for fast and memory-efficient autodiff, nor have we benchmarked it against competing frameworks such as TENSORFLOW. We expect that integration with NUMPY and similarity to the existing standard POPPY will make MORPHINE more readily useful in astronomy, but for integration with neural networks, it may be more convenient to use a different autodiff package.




Figure 10. A representation of the Jacobian determined by JAX for a simple Lyot coronagraph. In each pair of images, on the left is the PSF, with a red dot highlighting the selected pixel; on the right is a map of the corresponding pixel’s derivative with respect to phase over the pupil, the pattern of pupil phases that linearly affect that speckle. Speckles far from the central dark hole are generated by sinusoids across the pupil, while speckles near the inner working angle correspond to more complicated patterns.

4. Open Science

In the interests of open science, we have made available the Python package MORPHINE, together with Jupyter notebooks used to generate the figures in this paper, under a BSD three-clause open-source license at github.com/benjaminpope/morphine. We encourage and welcome other scientists to replicate, apply, and extend our work.

5. Conclusions and Future Work

We have shown that automatic differentiation software offers significant benefits for simulating astronomical optics. Using a simple model of diffraction, we can recover the idea of kernel phases, and using real Palomar data, we obtain similar astrometric performance to the analytically derived standard approach. Crucially, this can then be extended to arbitrary systems for which analytic self-calibration may not be possible.

We have demonstrated this for a simple Lyot coronagraph, finding a modal basis and kernel observables analogous to kernel phases. However, this is unfortunately of limited practical utility, due to the strong nonlinearities in the Lyot coronagraph problem studied. Nevertheless, alternate coronagraphs and imaging systems such as vortex coronagraphs (Foo et al. 2005), phase-induced amplitude apodization (“PIAA” Guyon 2003), or nulling interferometers (Bracewell 1978) may derive greater benefit from analysis with this approach. In addition to this, second- or higher-order Taylor expansion of the aberrations may reveal higher-order kernel phases even in the Lyot case (Y. Xin et al. 2020, in preparation).

In this paper, we have shown that autodiff allows us to straightforwardly account for broadband effects as a linear combination of diffraction simulations, which are also differentiable. Likewise, linear mixing of polarization states via Mueller matrices could be included in this framework. This




Figure 11. Log-scale plot of the singular values of the phase transfer matrix vs. index for the coronagraph simulation in Section 2.4. Singular values are normalized to the first singular value.




Figure 12. Pairs of nonsingular vectors in the pupil and image planes from the SVD of the coronagraph Jacobian in 10.

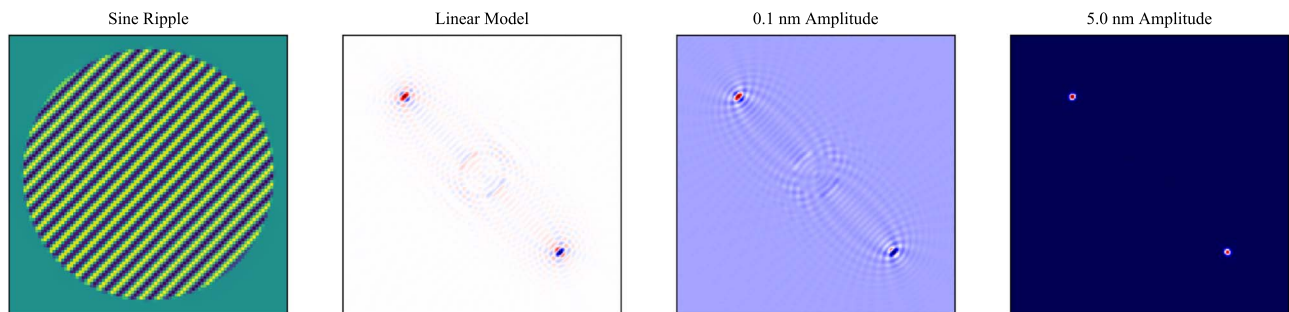


Figure 13. The effect of a sinusoidal phase ripple on the speckle pattern of the coronagraph in Section 2.4. Left: sinusoidal phase pattern. Center left: linear transfer map applied to the sinusoidal phase pattern, showing alternating positive and negative speckles on either side of the center of the PSF. Center right: the difference between a full optical simulation with an 0.1 nm amplitude phase ripple with respect to an unaberrated PSF, clearly close to the linear case. Right: the same difference simulation but with a 5 nm amplitude ripple, showing positive speckles on both sides of the detector—the quadratic term.

may be important for high-angular-resolution differential-polarimetry instruments such as SPHERE/ZIMPOL (Schmid et al. 2018) on the VLT, VAMPIRES on the Subaru SCEAO (Norris et al. 2014), or the GPI polarimeter (Perrin et al. 2015), where polarization measurements are intrinsic to the goal of the experiment, or to issues with polarization leakage affecting closure phases in holographic aperture masking experiments (e.g., Doelman et al. 2018).

In high-contrast imaging with adaptive optics, a major issue of residual aberration comes from differing optical paths: both those between the deformable mirror and the wave-front sensor, and the deformable mirror to the science camera. Non-common-path aberrations in this part of the optical system cannot be sensed and corrected inside the AO loop alone. Slowly evolving non-common-path errors cause quasi-static speckles that are one of the key sources of noise limiting the sensitivity of high-contrast instruments to exoplanets. In systems with an asymmetric pupil, the kernel phase formalism can also be used for focal-plane wave-front sensing (Martinache 2013). This has the advantage of avoiding non-common-path errors between the wave-front sensor and the science instrument. This can be made even more effective with an integral field spectrograph, as spectrally dispersed PSFs encode a hologram of the input wave front (Martinache 2016). The asymmetric pupil linear phase transfer approach has been used for focal-plane wave-front sensing in the lab (e.g., Pope et al. 2014a, 2014b) and on Subaru SCEAO (Martinache et al. 2016). Extending the domain of applicability of the linear phase transfer paradigm to coronagraphs using autodiff may allow for new approaches to focal-plane wave-front sensing.

The approach of Sitzmann et al. (2018) to end-to-end optimization of an optical system has clear relevance to optical astronomy. Going beyond using gradients for phase retrieval (Jurling & Fienup 2014), we can envision gradient-based phase design for coronagraphs, for diffractive-pupil telescopes, or for starshades. For example, it may be useful in rapidly finding deformable mirror settings to generate a “dark hole” in high-contrast imaging (e.g., Malbet et al. 1995; Currie et al. 2020). As noted above, autodiff has already been applied extensively in photonics; there may be many useful applications in astrophotonics, combining astronomical imaging optics with photonic devices such as beam combiners, photonic lanterns, and single-mode spectrographs (Bland-Hawthorn & Leon-Saval 2017; Minardi et al. 2020). It may also be useful for diffractive-pupil design: for example, the TOLIMAN space telescope concept (Tuthill et al. 2018; Bendek et al. 2018) aims to use a pupil-plane phase mask to achieve high-precision relative astrometry of α Centauri AB to search for exoplanets. Work so far in designing the phase mask has been in optimizing parameterized pupils with respect to the radially weighted gradient energy of the PSF as a proxy for astrometric precision. It will be in principle possible to calculate the Fisher information for the astrometric precision and to use this as an objective function for optimizing a nonparametric TOLIMAN pupil. By differentiating with respect to phase in intermediate planes, MORPHINE-like methods will also be useful in intermediate-plane wave-front control and mask design.

The generalization of kernel phase and optical gradient design permitted by autodiff also allows us to extend this work to the near-field (Fresnel) propagation regime. This may lead to improvements in modeling of components (such as deformable mirrors) that might occur in planes intermediate between pupil

and focus, or simply more accurate kernel phase models of instruments such as the HST. This may be most relevant outside of astronomy: for example, Fresnel coherent diffractive imaging (Williams et al. 2006) is a popular microscopy technique, in which gradient-based advances in phase retrieval have been applied (Du et al. 2020), and kernel phase and optimization may be valuable.

We would like to thank Anand Sivaramakrishnan, Marshall Perrin, Will Farr, David Hogg, Louis Desdoigts, Alison Wong, and Nour Skaf for their very helpful comments.

This work was performed in part under contract with the Jet Propulsion Laboratory (JPL) funded by NASA through the Sagan Fellowship Program executed by the NASA Exoplanet Science Institute.

This research made use of NASA’s Astrophysics Data System.

B.J.S.P. acknowledges being on the traditional territory of the Lenape Nations and recognizes that Manhattan continues to be the home to many Algonkian peoples. We give blessings and thanks to the Lenape people and Lenape Nations in recognition that we are carrying out this work on their indigenous homelands. We acknowledge and pay respect to the Gadigal people of the Eora Nation. It is upon their ancestral lands that the University of Sydney is built.

Software: This research made use of JAX (Bradbury et al. 2018); POPPY, an open-source optical propagation Python package originally developed for the James Webb Space Telescope project (Perrin et al. 2012); the IPYTHON package (Pérez & Granger 2007); NUMPY (Harris et al. 2020); MATPLOTLIB (Hunter 2007); SCIPY (Jones et al. 2001); emcee (Foreman-Mackey et al. 2013) and corner.py (Foreman-Mackey 2016); and Astropy, a community-developed core Python package for astronomy (Astropy Collaboration et al. 2013).

ORCID iDs

Benjamin J. S. Pope  <https://orcid.org/0000-0003-2595-9114>

Laurent Pueyo  <https://orcid.org/0000-0003-3818-408X>

Yinzi Xin  <https://orcid.org/0000-0002-6171-9081>

Peter G. Tuthill  <https://orcid.org/0000-0001-7026-6291>

References

- Abadi, M., Agarwal, A., Barham, P., et al. 2015, TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems, <http://tensorflow.org/>
- Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, *A&A*, **558**, A33
- Bandeira, A. S., Blum-Smith, B., Kileel, J., et al. 2017, arXiv:1712.10163
- Bendek, E., Tuthill, P., Guyon, O., et al. 2018, *Proc. SPIE*, **10698**, 106980G
- Bezanson, J., Karpinski, S., Shah, V. B., & Edelman, A. 2012, arXiv:1209.5145
- Blackburn, L., Pesce, D. W., Johnson, M. D., et al. 2020, *ApJ*, **894**, 31
- Bland-Hawthorn, J., & Leon-Saval, S. G. 2017, *OEExpr*, **25**, 15549
- Bloemhof, E. E., Dekany, R. G., Troy, M., & Oppenheimer, B. R. 2001, *ApJL*, **558**, L71
- Bracewell, R. N. 1978, *Natur*, **274**, 780
- Bradbury, J., Frostig, R., Hawkins, P., et al. 2018, JAX: composable transformations of Python+NumPy programs, 0.1.55, <http://github.com/google/jax>
- Carlotti, A., Vanderbei, R., & Kasdin, N. J. 2011, *OEExpr*, **19**, 26796
- Ceau, A., Mary, D., Greenbaum, A., et al. 2019, *A&A*, **630**, A120
- Chael, A. A., Johnson, M. D., Bouman, K. L., et al. 2018, *ApJ*, **857**, 23
- Chaware, A., Cooke, C. L., Kim, K., & Horstmeier, R. 2019, arXiv:1910.10209

- Chianese, M., Coogan, A., Hofma, P., Otten, S., & Weniger, C. 2020, *MNRAS*, **496**, 381
- Currie, T., Pluzhnik, E., Guyon, O., et al. 2020, *PASP*, **132**, 104502
- Czekala, I. 2019, iancze/MPoL: Base version, v0.0.1, Zenodo, doi:10.5281/zenodo.3647603
- Doelman, D. S., Tuthill, P., Norris, B., et al. 2018, *Proc. SPIE*, **10701**, 107010T
- Du, M., Nashed, Y. S. G., Kandel, S., Gürsoy, D., & Jacobsen, C. 2020, *SciA*, **6**, eaay3700
- Feinstein, A. D., Montet, B. T., Foreman-Mackey, D., et al. 2019, *PASP*, **131**, 094502
- Fizeau, H. 1868, *CRASE*, **66**, 932
- Foo, G., Palacios, D. M., & Swartzlander, G. A., Jr. 2005, *OptL*, **30**, 3308
- Foreman-Mackey, D. 2016, *JOSS*, **1**, 24
- Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, *PASP*, **125**, 306
- Guo, X., Barrett, T. D., Wang, Z. M., & Lvovsky, A. I. 2019, arXiv:1912.12256
- Guyon, O. 2003, *A&A*, **404**, 379
- Guyon, O., Bendek, E. A., Eisner, J. A., et al. 2012, *ApJS*, **200**, 11
- Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, *Natur*, **585**, 357
- Hughes, T. W., Minkov, M., Williamson, I. A. D., & Fan, S. 2018, *ACS Photonics*, **5**, 4781
- Hughes, T. W., Williamson, I. A. D., Minkov, M., & Fan, S. 2019, *SciA*, **5**, eaay6946
- Hunter, J. D. 2007, *CSE*, **9**, 90
- Ireland, M. J. 2013, *MNRAS*, **433**, 1718
- Jennison, R. C. 1958, *MNRAS*, **118**, 276
- Jones, E., Oliphant, T., Peterson, P., et al. 2001, SciPy: Open source scientific tools for Python, <http://www.scipy.org/>
- Jurling, A. S., & Fienup, J. R. 2014, *JOSAA*, **31**, 1348
- Kammerer, J., Ireland, M. J., Martinache, F., & Girard, J. H. 2019, *MNRAS*, **486**, 639
- Kandel, S., Maddali, S., Allain, M., et al. 2019, *OExpr*, **27**, 18653
- Lafrenière, D., Marois, C., Doyon, R., Nadeau, D., & Artigau, É. 2007, *ApJ*, **660**, 770
- Lallo, M. D., Makidon, R. B., Casertano, S., & Krist, J. E. 2006, *Proc. SPIE*, **6270**, 62701N
- Lannes, A. 1991, *InvPr*, **7**, 261
- Laugier, R., & Martinache, F. 2019, in SF2A-2019: Proc. of the Annual meeting of the French Society of Astronomy and Astrophysics, ed. P. Di Matteo et al. (Bordeaux: SFSA), **231**
- Laugier, R., Martinache, F., Ceau, A., et al. 2019, *A&A*, **623**, A164
- Laugier, R., Martinache, F., Cvetojevic, N., et al. 2020, *A&A*, **636**, A21
- LeCun, Y., Bengio, Y., & Hinton, G. 2015, *Natur*, **521**, 436
- LeCun, Y., Touresky, D., Hinton, G., & Sejnowski, T. 1988, in Proc. 1988 Connectionist Models Summer School (Pittsburgh, PA: Morgan Kaufmann), **21**
- Linnainmaa, S. 1970, Master's Thesis, University of Helsinki
- Lyot, B. 1930, *BuAst*, **6**, 305
- Maclaurin, D., Duvenaud, D., & Adams, R. P. 2015, in ICML 2015 AutoML Workshop (San Diego, CA: ICML)
- Malbet, F., Yu, J. W., & Shao, M. 1995, *PASP*, **107**, 386
- Marois, C., Lafrenière, D., Doyon, R., Macintosh, B., & Nadeau, D. 2006, *ApJ*, **641**, 556
- Martinache, F. 2010, *ApJ*, **724**, 464
- Martinache, F. 2011, *Proc. SPIE*, **8151**, 815111
- Martinache, F. 2013, *PASP*, **125**, 422
- Martinache, F. 2016, *Proc. SPIE*, **9907**, 990712
- Martinache, F. 2018, Repousser les limites de la diffraction pour l'astronomie à haute résolution angulaire, Habilitation, Université de Nice Sophia Antipolis, http://frantzmartinache.eu/static/share/hdr_no_papers.pdf
- Martinache, F., Ceau, A., Laugier, R., et al. 2020, *A&A*, **636**, A72
- Martinache, F., & Ireland, M. J. 2018, *A&A*, **619**, A87
- Martinache, F., Jovanovic, N., & Guyon, O. 2016, *A&A*, **593**, A33
- Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H., & Teller, E. 1953, *JChPh*, **21**, 1087
- Minardi, S., Harris, R., & Labadie, L. 2020, arXiv:2003.12485
- Morningstar, W. R., Hezaveh, Y. D., Perreault Levasseur, L., et al. 2018, arXiv:1808.00011
- Morningstar, W. R., Perreault Levasseur, L., Hezaveh, Y. D., et al. 2019, *ApJ*, **883**, 14
- Muthumbi, A., Chaware, A., Kim, K., et al. 2019, *Biomed. Opt. Express*, **10**, 6351
- Nardiello, D., Borsato, L., Piotto, G., et al. 2019, *MNRAS*, **490**, 3806
- Nashed, Y. S. G., Kandel, S., Du, M., & Jacobsen, C. 2019, *MiMic*, **25**, 62
- Norris, B., Tuthill, P., Jovanovic, N., et al. 2014, arXiv:1405.7426
- Paine, S. W., & Fienup, J. R. 2019, *Proc. SPIE*, **10980**, 109800T
- Paszke, A., Gross, S., Massa, F., et al. 2019, arXiv:1912.01703
- Pérez, F., & Granger, B. E. 2007, *CSE*, **9**, 21
- Perrin, M. D., Duchene, G., Millar-Blanchaer, M., et al. 2015, *ApJ*, **799**, 182
- Perrin, M. D., Sivaramakrishnan, A., Makidon, R. B., Oppenheimer, B. R., & Graham, J. R. 2003, *ApJ*, **596**, 702
- Perrin, M. D., Soummer, R., Elliott, E. M., Lallo, M. D., & Sivaramakrishnan, A. 2012a, *Proc. SPIE*, **8442**, 84423D
- Pope, B., Cvetojevic, N., Cheetham, A., et al. 2014a, *MNRAS*, **440**, 125
- Pope, B., Martinache, F., & Tuthill, P. 2013, *ApJ*, **767**, 110
- Pope, B., Thatte, N., Burruss, R., et al. 2014b, *Proc. SPIE*, **9148**, 914859
- Pope, B., Tuthill, P., Hinkley, S., et al. 2016, *MNRAS*, **455**, 1647
- Pope, B. J. S. 2016, *MNRAS*, **463**, 3573
- Pope, B. J. S., White, T. R., Farr, W. M., et al. 2019, *ApJS*, **245**, 8
- Pueyo, L. 2016, *ApJ*, **824**, 117
- Rieke, G. H., Ressler, M. E., Morrison, J. E., et al. 2015, *PASP*, **127**, 665
- Riggs, A. J. E., Ruane, G., Sidick, E., et al. 2018, *Proc. SPIE*, **10698**, 106982V
- Sallum, S., Follette, K. B., Eisner, J. A., et al. 2015, *Natur*, **527**, 342
- Sallum, S., & Skemer, A. 2019, *JATIS*, **5**, 018001
- Sallum, S., Skemer, A., Eisner, J., et al. 2019, *ApJ*, **883**, 100
- Schmid, H. M., Bazzon, A., Roelfsema, R., et al. 2018, *A&A*, **619**, A9
- Sitzmann, V., Diamond, S., Peng, Y., et al. 2018, *ACM Trans. Graph.*, **37**, 114
- Sivaramakrishnan, A., Lloyd, J. P., Hodge, P. E., & Macintosh, B. A. 2002, *ApJL*, **581**, L59
- Soummer, R., Pueyo, L., & Larkin, J. 2012, *ApJL*, **755**, L28
- Soummer, R., Pueyo, L., Sivaramakrishnan, A., & Vanderbei, R. J. 2007, *OExpr*, **15**, 15935
- Sparks, W. B., & Ford, H. C. 2002, *ApJ*, **578**, 543
- Sutin, B. M. 2016, *Proc. SPIE*, **9911**, 99112J
- Theano Development Team 2016, 02688, arXiv:1605.02688
- Tuthill, P., Bendek, E., Guyon, O., et al. 2018, *Proc. SPIE*, **10701**, 107011J
- Twiss, R. Q., Carter, A. W. L., & Little, A. G. 1960, *Obs*, **80**, 153
- van Cittert, P. H. 1934, *Phy*, **1**, 201
- Werner, J., Hillenbrand, M., Hoffmann, A., Sinzinger, S., et al. 2012, *Schedae Inform.*, **2012**, 169
- White, T. R., Pope, B. J. S., Antoci, V., et al. 2017, *MNRAS*, **471**, 2882
- Williams, G. J., Quiney, H. M., Dhal, B. B., et al. 2006, *PhRvL*, **97**, 025506
- Zernike, F. 1938, *Phy*, **5**, 785
- Zernike, V. F. 1934, *Phy*, **1**, 689