UCLA
UCLA Previously Published Works

Title

Speckle Space–Time Covariance in High-contrast Imaging

Permalink

https://escholarship.org/uc/item/7ht8j9k3

Journal

The Astronomical Journal, 165(2)

ISSN

0004-6256

Authors

Lewis, Briley
Fitzgerald, Michael P
Dodkins, Rupert H
et al.

Publication Date

2023-02-01

DOI

10.3847/1538-3881/aca761

Copyright Information

Peer reviewed

arXiv:2301.01291v1 [astro-ph.IM] 3 Jan 2023

Speckle Space-Time Covariance in High-Contrast Imaging

BRILEY LEWIS [ID],[1] MICHAEL P. FITZGERALD [ID],[1] RUPERT H. DODKINS [ID],[2] KRISTINA K. DAVIS [ID],[2] AND JONATHAN LIN [ID][1]

[1]*Department of Physics and Astronomy, UCLA, Los Angeles, CA 90024 USA*
[2]*Department of Physics, UCSB, Santa Barbara, CA 93106 USA*

ABSTRACT

We introduce a new framework for point-spread function (PSF) subtraction based on the spatio-temporal variation of speckle noise in high-contrast imaging data where the sampling timescale is faster than the speckle evolution timescale. One way that space-time covariance arises in the pupil is as atmospheric layers translate across the telescope aperture and create small, time-varying perturbations in the phase of the incoming wavefront. The propagation of this field to the focal plane preserves some of that space-time covariance. To utilize this covariance, our new approach uses a Karhunen-Loéve transform on an image sequence, as opposed to a set of single reference images as in previous applications of Karhunen-Loéve Image Processing (KLIP) for high-contrast imaging. With the recent development of photon-counting detectors, such as microwave kinetic inductance detectors (MKIDs), this technique now has the potential to improve contrast when used as a post-processing step. Preliminary testing on simulated data shows this technique can improve contrast by at least 10–20% from the original image, with significant potential for further improvement. For certain choices of parameters, this algorithm may provide larger contrast gains than spatial-only KLIP.

Keywords: exoplanet detection; high contrast imaging; atmospheric effects; instrumentation: adaptive optics; methods: data analysis; methods: statistical; techniques: image processing

1. INTRODUCTION

Direct imaging of exoplanets is a challenging endeavor, given the extreme contrasts that must be achieved to detect faint planets. Although significant starlight suppression can be achieved through optics and instrumentation, such as coronagraphs, adaptive optics (AO) systems, interferometers, and more, that alone is insufficient to detect analogs of planets in our solar system (Oppenheimer & Hinkley 2009; Guyon 2005). Improving contrast expands the space of the types of planets that can be directly detected and characterized.

Existing instruments, such as the Gemini Planet Imager (Macintosh et al. 2008) and VLT's SPHERE (Beuzit et al. 2019) are able to image giant planets and brown dwarfs, reaching contrasts (in the astronomical sense, meaning the detectable planet-star flux ratio) of around 10^{-6}. This is enabled by a combination of wavefront sensing, control, and post-processing, which re-

Corresponding author: Briley Lewis
blewis@astro.ucla.edu

duce the impact of noise by distinguishing between the planet signal and residual noise; this noise arises from uncorrected wavefront aberrations, resulting in quasi-static fluctuations in the focal plane known as "speckles." Generally, these algorithms use the data themselves to create a model of the speckle noise which can then be subtracted from the data to recover the target planet signal in a process known as point-spread function (PSF) subtraction. Previously developed algorithms include LOCI (Locally Optimized Combination of Images; Lafreniere et al. (2007)), KLIP (Karhunen Loéve Image Processing; Soummer et al. 2012), and more (Gebhard et al. 2022). Many directly imaged planet discoveries to date have relied on such algorithms, such as the famous HR 8799 planets (Marois et al. 2008).

Improvements to data processing pipelines and methods are one way in which we can push forward and improve contrast for future high-contrast imaging instruments. Other approaches to improving high-contrast imaging methods focus on wavefront sensing and control, such as predictive control techniques, which aim to improve adaptive optics corrections (Guyon & Males 2017; Guyon et al. 2018; Males & Guyon 2018), and sen-

sor fusion, both currently in development at multiple facilities, including Subaru's SCExAO facility (Guyon et al. 2017) and Keck Observatory van Kooten et al. (2021); Wizinowich et al. (2020); Jensen-Clem et al. (2019); Calvin et al. (2022). Other recent work such as Guyon & Males (2017) focuses on using on Empirical Orthogonal Functions (EOFs), a similar mathematical framework, to analyze spatio-temporal correlations; their work is in the context of predictive control, whereas our work applies to image processing. New advances in detector technology also affect both wavefront sensing and post-processing. High-speed, low-noise detectors will provide multiple opportunities for improvements, including focal-plane wavefront sensing, which eliminates non-common-path wavefront errors (Vievard et al. 2020). Of particular interest are arrayed photon-counting devices, such as Microwave Kinetic Inductance Detectors (MKIDs) (Schlaerth et al. 2008; Mazin et al. 2012; Meeker et al. 2018; Walter et al. 2020) and Infrared Avalanche Photodiodes (IR APDs) (Goebel 2018; Wu et al. 2021). Electron Multiplying CCDs (EMCCDs) are a functional equivalent in the optical (Lake et al. 2020). Photon arrival times have already been used to distinguish speckles from incoherent signals, such as planets (Walter et al. 2019; Steiger et al. 2021), and MKIDs have been used for high contrast imaging with the DARK-NESS instrument at Palomar (Meeker et al. 2018) and with MEC, the MKID Exoplanet Camera for high contrast astronomy at Subaru (Walter et al. 2018).

This new regime of photon-counting detectors and more advanced adaptive optics presents many opportunities. With the improved temporal resolution of next-generation detectors, we will be able to resolve the spatial and temporal evolution of atmospheric speckles. Some prior work has investigated use of spatio-temporal correlations on longer timescales, such as Mullen et al. (2019) and Gebhard et al. (2022), but this work focuses the shorter timescale changes of atmospheric speckles.

There is a rich history of theory and measurements of space-time atmospheric speckle behavior in the past decades, which this work builds off of. Since the 1970s–1980s, speckle patterns and intensity distributions have been measured (Dainty et al. 1981; Scaddan & Walker 1978; Goebel 2018; Odonnell et al. 1982), demonstrating agreement with models based in Rician statistics (Cagigal & Canales 2001; Canales & Cagigal 1999) and the importance of speckles as the limiting noise source in the high-contrast regime (Racine et al. 1999). The space-time covariance was even directly measured in Dainty et al. (1981), indicating that speckle boiling has a directionality related to turbulence. Speckle intensity patterns have been modeled as a modified Rician distribu-

tion (Aime & Soummer 2004; Gladysz et al. 2010), and speckle lifetimes have been constrained through models and direct measurements (Aime et al. 1986; Vernin et al. 1991; Glindemann et al. 1993). In fact, models of speckle boiling directly relate the lifetime of speckles to atmospheric parameters related to wind and turbulence, as in Roddier et al. (1982), estimating speckle lifetimes on the order of tens of milliseconds.

This work is a new addition to the variety of time-domain algorithms that have been developed in recent years. For example, the PACO algorithm uses temporal information from the background fluctuations of Angular Differential Imaging data (Flasseur et al. 2018), and the TRAP algorithm uses temporal information of the speckle pattern to improve contrast specifically at close separations (Samland et al. 2021). Another algorithm, from Gebhard et al. (2022), uses half-sibling regression on time-series data. These are all examples of the possibilities for temporal information in post-processing, in addition to the AO control improvements described earlier.

In this work, we aim for a second-order characterization of the statistical behavior of atmospheric speckles in the high-contrast regime, described by the space-time covariance, which we then leverage for improving contrast in post-processing with the eventual goal of improving exoplanet detection capabilities. As previously mentioned, this goal is not without its challenges — with kHz readouts, these detectors can produce large datasets and lead to computationally intensive post-processing methods. While developing this new technique, we must also contend with data storage and computational limitations.

In this paper, we first provide an analytical justification for the existence of these covariances in the high-contrast regime, observe their occurrence in test simulations focusing on millisecond time sampling, and then present an initial algorithm to exploit these covariances for PSF subtraction. Specifically, we are testing this algorithm in a regime dominated by atmospheric speckles at short exposures (where the timescale of our exposures is short compared to that of changes in atmospheric residual wavefront error, so atmosphere is essentially frozen).

Here in Section 2, we describe the process of baseline Karhunen-Loéve Image Processing (KLIP), the origins of space-time speckle covariances, and the extension of KLIP to space-time covariances. Following, in Section 3, we describe the models used to create datasets for initial testing of this processing framework. Section 4 presents results of this new algorithm implemented on simulated data. Finally, in Sections 5 and 6, we discuss

the promise of this new technique, as well as its current challenges/limitations and future work.

2. SPACE-TIME COVARIANCE THEORY

Speckles can limit contrast, but can also be subtracted to some extent to improve contrast. One of the most successful post-processing algorithms has been KLIP, described in Section 2.1, which exploits spatial correlations in long-exposure images. We motivate our extension of this technique to include space-time correlations on shorter timescales in Section 2.2 by describing how these correlations arise in imaging through the atmosphere. This extension of KLIP, referred to as space-time KLIP or stKLIP, is demonstrated in Section 2.3, exploiting spatio-temporal correlations between short-exposure images.

2.1. *Karhunen-Loéve Image Processing*

Karhunen-Loéve Image Processing is a data processing technique that uses principle component analysis (PCA), where data are represented by a linear combination of orthogonal functions. In high-contrast imaging, KLIP is used to build a model, used for PSF subtraction, that accounts for spatial correlations between speckles and other PSF features, first described in Soummer et al. (2012). This technique takes advantage of spatial covariances of the speckles in the image, because strong correlations exist in high eigenvalue modes and can be suppressed. This is a data-driven approach, which uses available information from the data itself to provide an approximation of the noise, by using a subset of the data as "reference images" from which to build the model of the noise while using another subset of the data as the "target image" for PSF subtraction.

To increase readability, all variables for the following mathematics are described in Appendix A. As described in Soummer et al. (2012), we assume we observe a point spread function $T(k)$, where k is the pixel index, that contains the stellar point spread function $I_\psi(k)$ and may also contain some faint astronomical signal of interest $A(k)$. Therefore, our target image can be described as

$$T(k) = I_\psi(k) + \epsilon A(k), \qquad (1)$$

where ϵ is 0 when there is no astronomical signal of interest, or 1 if there is. The goal of PSF subtraction is therefore to recreate $I_\psi(k)$ in order to isolate $A(k)$. Without an infinite number of references, though, we cannot exactly infer $I_\psi(k)$; instead, we approximate the PSF $\hat{I}_\psi(k)$. For consistency in our notation, herein we represent $T(k)$, $A(k)$, and $\hat{I}_\psi(k)$ as vectors $\boldsymbol{t}$, $\boldsymbol{a}$ and $\hat{\boldsymbol{\psi}}$ respectively.

In order to approximate $\hat{\boldsymbol{\psi}}$, KLIP computes a Karhunen-Loéve Transform based on the covariance matrix of the mean-subtracted reference images.

A sequence of reference images are first unraveled into one-dimensional vectors, each as $\boldsymbol{r}$. Note: henceforth vectors are denoted as bold, matrices with uppercase variables and subscript matrix elements. These image vectors $\boldsymbol{r}$ are then stacked into an $n_p \times n_i$ matrix R, where $n_p = n_x \times n_y$ and n_i is the number of images, as follows:

$$R = \begin{bmatrix} R_{1,1} & R_{1,2} & \dots & R_{1,n_i} \\ R_{2,1} & R_{2,2} & \dots & R_{2,n_i} \\ \vdots & \vdots & \ddots & \vdots \\ R_{n_p,1} & R_{n_p,2} & \dots & R_{n_p,n_i} \end{bmatrix} \qquad (2)$$

We then subtract the mean image of the set (summing over the matrix columns) from the reference set R, in order to produce a set of mean-subtracted images M to use throughout the process of KLIP:

$$\boldsymbol{x}_i = \frac{1}{n_i} \sum_{j=1}^{n_i} R_{i,j} \qquad (3)$$

$$M_{i,j} = R_{i,j} - \boldsymbol{x}_i \qquad (4)$$

The resulting covariance matrix (5) C has size $n_p \times n_p$.

$$C = MM^T = \begin{bmatrix} C_{1,1} & C_{1,2} & \dots & C_{1,n_p} \\ C_{2,1} & C_{2,2} & \dots & C_{2,n_p} \\ \vdots & \vdots & \ddots & \vdots \\ C_{n_p,0} & C_{n_p,1} & \dots & C_{n_p,n_p} \end{bmatrix} \qquad (5)$$

Note: in practice, this implementation is computationally expensive, so the covariance is instead often computed in image space on n_i by n_i images and then re-projected into pixel space, as is done in the Soummer et al. (2012) implementation. The ideal implementation depends on which dimension is larger / more computationally expensive, e.g. Long & Males (2021). In this work, the mathematics for KLIP and stKLIP, as written here, will be in pixel space.

An eigendecomposition of the covariance matrix C, mathematically described as solutions to the equation

$$C\boldsymbol{v}_j = \lambda_j \boldsymbol{v}_j, \qquad (6)$$

with

$$\lambda_1 > \lambda_2 > \lambda_3 > \dots \lambda_{n_p}, \qquad (7)$$

produces a length n_p vector of eigenvalues ($\boldsymbol{\lambda}$) and size $n_p \times n_p$ (or $n_m \times n_p$ if fewer than n_p eigenvectors/modes

are used) matrix of eigenvectors/eigenimages (V) containing n_m rows of individual eigenvectors $\boldsymbol{v}$ each of length n_p, such that $V_{i,j} = (\boldsymbol{v}_j)_i$.

$$V = \begin{bmatrix} V_{1,1} & V_{1,2} & \dots & V_{1,n_p} \\ V_{2,1} & V_{2,2} & \dots & V_{2,n_p} \\ \vdots & \vdots & \ddots & \vdots \\ V_{n_m,1} & V_{n_m,2} & \dots & V_{n_m,n_p} \end{bmatrix} \quad (8)$$

The eigenvalues order the eigenimages by their importance to rebuilding the original image and are used to construct the basis of the new subspace of greatest variation onto which we project our images. Assuming the vectors are sorted by decreasing eigenvalue, the first coordinate corresponds to the direction of greatest variation. The lowest-order (first coordinate) eigenimages are selected to represent $\hat{\boldsymbol{\psi}}$, while leaving the high-order terms to hopefully contain our astrophysical signal.

We select a given number n_m of the eigenimages as our number of modes of variation. The inner product of the matrix of eigenvectors V with the one-dimensional vector of the target image $\boldsymbol{t}$ (which has length n_p), is described mathematically as

$$\boldsymbol{t} = \begin{bmatrix} t_1 \\ t_2 \\ \vdots \\ t_{n_p} \end{bmatrix} \quad (9)$$

$$\boldsymbol{q} = V \cdot \boldsymbol{t} = \begin{bmatrix} q_1 \\ q_2 \\ \vdots \\ q_{n_m} \end{bmatrix} \quad (10)$$

and creates a vector of coefficients $\boldsymbol{q}$ of length n_m — each of these can be thought of as how much of each mode (or each eigenvector, v_j) is in the image, or equivalently, the coordinates in the new rotated principle axis space.

Lastly, we can project back into our original pixel space by taking the product of this vector of coefficients with the chosen eigenvectors, recovering a vector of length n_p, the same as our target image:

$$\hat{\boldsymbol{\psi}} = \boldsymbol{q}^T \cdot V \quad (11)$$

The resulting array is our image projected into the subspace of greatest variation, an estimation of the original PSF $\hat{\boldsymbol{\psi}}$, and what we will subtract from our target image for PSF subtraction. Note that the tuneable parameter here is the number of eigenvectors used in the basis (the number of "modes").

The planet signal is also projected onto a distribution of these modes, and it is assumed that the planet signal is primarily projected onto modes with lower eigenvalue. However, as we subtract more modes, the projection of the planet onto these modes is also subtracted. Therefore, a larger number of modes might lead to oversubtraction of a planet signal, but too few may not sufficiently subtract out the speckle noise. As a result, we must correct for this throughput effect and optimize the number of modes to attain the largest possible contrast gain.

2.2. *Space-Time Covariances*

Whereas KLIP harnesses spatial covariances of speckle noise, we propose to expand the scope of such projection methods to take advantage of space-time covariances in speckle noise. For bulk flow in a turbulent atmosphere, phase errors in the pupil, from atmospheric disturbances, translate across the telescope with wind motion, resulting in changes in phase and amplitude in the image plane. Atmospheric perturbations evolve across a broad set of spatial frequencies. Since the perturbations at these different spatial frequencies are correlated, we will illustrate that the speckles at the locations that correspond to those spatial frequencies in the image plane will be correlated as well. Similarly to the above section, all variables for the following mathematics are described in Appendix B.

The covariance of intensity in the image plane for points separated in space and time is characterized through the second moment $\langle I(\boldsymbol{x}_1,t)I(\boldsymbol{x}_2,t-\tau)\rangle$, where I is the intensity in the image. Angle brackets ($\langle\rangle$) denote averaging over a statistical ensemble. Suppose we have a perfect coronagraph and only phase aberrations are present, ignoring polarization as well as static phase errors, and treating electric field as a scalar. Also, we presume the phase aberrations are small, a reasonable assumption for the high-contrast imaging limit. In this case, the pupil amplitude is

$$\Psi_{\text{pup}}(\boldsymbol{u},t) = P(\boldsymbol{u})e^{i\phi(\boldsymbol{u},t)}, \quad (12)$$

approximated as

$$\Psi_{\text{pup}}(\boldsymbol{u},t) \approx [1 + i\phi(\boldsymbol{u},t)]P(\boldsymbol{u}), \quad (13)$$

where $P(\boldsymbol{u})$ is the pupil function, ϕ is the phase, and $\boldsymbol{u}$ is the coordinate in the pupil plane ($\boldsymbol{x}$ is the coordinate in the focal plane, related by a Fourier transform). It is worth noting that departure from this assumption of linearity may affect results. The amplitude in the focal plane is

$$\Psi_{\text{foc}}(\boldsymbol{x},t) = \mathcal{F}\{P(\boldsymbol{u})\} + i\mathcal{F}\{\phi(\boldsymbol{u},t)P(\boldsymbol{u})\}, \quad (14)$$

$$= C(\boldsymbol{x}) + S_\phi(\boldsymbol{x},t). \quad (15)$$

$C(\boldsymbol{x})$ is the spatially coherent part of the wavefront, and $S_\phi(\boldsymbol{x}, t)$ comes from phase aberrations – $S_\phi(\boldsymbol{x}, t)$ corresponds to the "speckles" we want to remove (Aime & Soummer 2004; Roddier et al. 1982). In the case of a perfect coronagraph, $C(\boldsymbol{x}) = 0$ and the intensity in the image is only due to phase aberrations, and can be expressed as

$$I(\boldsymbol{x}, t) = |\Psi_{\text{foc}}(\boldsymbol{x}, t)|^2, \tag{16}$$

$$= |S_\phi(\boldsymbol{x}, t)|^2, \tag{17}$$

$$= |\mathcal{F}\{\phi(\boldsymbol{u}, t)P(\boldsymbol{u})\}|^2. \tag{18}$$

The covariance of the intensity is

$$\langle I(\boldsymbol{x}_1, t)I(\boldsymbol{x}_2, t - \tau)\rangle = \langle |S_\phi(\boldsymbol{x}_1, t)S_\phi(\boldsymbol{x}_2, t - \tau)|^2\rangle. \tag{19}$$

If we assume (complex) Gaussian statistics for S_ϕ (Soummer et al. 2007), then by Wick's theorem (e.g. Fassino et al. 2019) we have,

$$\langle I(\boldsymbol{x}_1, t)I(\boldsymbol{x}_2, t - \tau)\rangle =$$
$$\langle I(\boldsymbol{x}_1, t)\rangle \langle I(\boldsymbol{x}_2, t)\rangle + |\langle S_\phi(\boldsymbol{x}_1, t)S_\phi^*(\boldsymbol{x}_2, t - \tau)\rangle|^2. \tag{20}$$

Therefore to compute this covariance, we need the quantity $\langle S_\phi(\boldsymbol{x}_1, t)S_\phi^*(\boldsymbol{x}_2, t - \tau)\rangle$, which is the covariance of the phase-induced aberration in the focal plane. Accounting for the Fourier relationship between the focal plane aberration S_ϕ and the pupil plane phase ϕ as in Equations 14 and 15, we find

$$\langle S_\phi(\boldsymbol{x}_1, t)S_\phi^*(\boldsymbol{x}_2, t - \tau)\rangle =$$
$$\int d\boldsymbol{u} \int d\boldsymbol{\xi} \exp[2\pi i\boldsymbol{\xi} \cdot \boldsymbol{x}_2 - 2\pi i\boldsymbol{u} \cdot (\boldsymbol{x}_1 - \boldsymbol{x}_2)] \times$$
$$\langle \phi(\boldsymbol{u}, t)\phi(\boldsymbol{u} + \boldsymbol{\xi}, t - \tau)\rangle P(\boldsymbol{u})P(\boldsymbol{u} + \boldsymbol{\xi}) \tag{21}$$

where $\boldsymbol{\xi}$ is the coordinate of the displacement in the pupil plane. If $\phi(\boldsymbol{u}, t)$ is statistically stationary in the pupil plane position $\boldsymbol{u}$ (and time), then we can define the phase covariance function as

$$B_\phi(\boldsymbol{\xi}, \tau) = \langle \phi(\boldsymbol{u}, t)\phi(\boldsymbol{u} + \boldsymbol{\xi}, t - \tau)\rangle, \tag{22}$$

independent of $\boldsymbol{u}$ and t. Equation 22 for B_ϕ relates space-time covariance in the pupil to space-time covariance in the image, and can be simplified into the Kolmogorov phase covariance function for turbulence with an assumption about time.

Kolmogorov's theory of turbulence describes a cascade of large scale turbulent motions that dissipate energy onto smaller scales, following a power spectrum described by $\Phi_n(\boldsymbol{k}) \propto |\boldsymbol{k}|^{-11/3}$, where Φ_n is the variation in index of refraction and $|\boldsymbol{k}|$ is the magnitude of the turbulence (Kolmogorov 1941; Hickson 2008). Fluctuations in density correspond to fluctuations in the index of refraction. These variations in index of refraction lead to differences in path length for the incoming light, creating some of the phase and amplitude error that we observe. However, we assume the timescale of change for this turbulence is generally slow when compared to wind speeds, an assumption known as Taylor frozen flow (Taylor 1938). This assumption is valid so long as the turbulent intensity is low compared to the wind speed, generally accepted to be true for astronomical contexts with the possible exception of boundary layer turbulence (Bharmal 2015). The turbulence can be thought of then as a "phase screen" propagating horizontally across the telescope with the wind. This phenomenon is described mathematically as

$$\phi(\boldsymbol{u}, t) = \phi(\boldsymbol{u} - \boldsymbol{v}_{\text{wind}}\tau, t - \tau) \tag{23}$$

which states that the phase structure at one time is related to the phase structure at a different time, just shifted by the wind velocity times the time difference (Taylor 1938; Hickson 2008).

This shows that a single phase screen $\phi(\boldsymbol{u}, t)$ (which contains Kolmogorov turbulence Φ_n) under Taylor frozen flow is related to a phase screen at a different time $\phi(\boldsymbol{u}, t - \tau)$ via the wind speed $\boldsymbol{v}_{\text{wind}}$. Similarly, we can then say

$$B_\phi(\boldsymbol{u}, t) = B_\phi(\boldsymbol{u} - \boldsymbol{v}_{\text{wind}}\tau, t - \tau). \tag{24}$$

This implies the phase covariance function at one location and time $B_\phi(\boldsymbol{\xi}, t)$ in the pupil is related to the phase covariance function at that location at a previous time $B_\phi(\boldsymbol{\xi}, 0)$, where $B_\phi(\boldsymbol{\xi}, 0)$ is a covariance related to the Kolmogorov phase covariance function. Since we know the Kolmogorov phase covariance function is non-zero as long as turbulence is present, this demonstrates that the phase covariance function at an arbitrary location and time $B_\phi(\boldsymbol{\xi}, \tau)$ is non-zero. Even if frozen flow is violated, as long as there is non-zero space-time covariance in the pupil, we expect non-zero space-time covariance in the image, as shown in Equation 22.

Rearranging Equation 21,

$$\langle S_\phi(\boldsymbol{x}_1, t)S_\phi^*(\boldsymbol{x}_2, t - \tau)\rangle =$$
$$\int d\boldsymbol{\xi} \exp(2\pi i\boldsymbol{\xi} \cdot \boldsymbol{x}_2)B_\phi(\boldsymbol{\xi}, \tau)$$
$$\int d\boldsymbol{u} \exp[-2\pi i\boldsymbol{u} \cdot (\boldsymbol{x}_1 - \boldsymbol{x}_2)]P(\boldsymbol{u})P(\boldsymbol{u} + \boldsymbol{\xi}). \tag{25}$$

The latter integral is the Fourier transform of the overlap of displaced pupils. Defining this function,

$$\mathcal{P}(\boldsymbol{r}, \boldsymbol{\xi}) = \int d\boldsymbol{u} \exp(-2\pi i\boldsymbol{u} \cdot \boldsymbol{r})P(\boldsymbol{u})P(\boldsymbol{u} + \boldsymbol{\xi}), \tag{26}$$

we now have the space-time covariance of speckles as the product of the turbulence phase covariance function and $\mathcal{P}$, as follows:

$$\langle S_\phi(\boldsymbol{x}_1, t) S_\phi^*(\boldsymbol{x}_2, t - \tau) \rangle =$$
$$\int d\boldsymbol{\xi} \exp(2\pi i \boldsymbol{\xi} \cdot \boldsymbol{x}_2) B_\phi(\boldsymbol{\xi}, \tau) \mathcal{P}(\boldsymbol{x}_1 - \boldsymbol{x}_2, \boldsymbol{\xi}). \quad (27)$$

This mathematical framework illustrates how the focal plane covariance is intimately related to pupil plane covariance in the high contrast imaging regime, with a perfect coronagraph and small phase errors. Looking at the overlap of displaced pupils, $\mathcal{P}(\boldsymbol{x}_1 - \boldsymbol{x}_2, \boldsymbol{\xi})$, the form of the expression suggests that covariance will be strongest at smaller spatial separations. Similarly, Equation 24 suggests that covariance will be strongest at smaller temporal separations. Overall, if there is non-zero space time covariance in the pupil plane, then we will have non-zero space time covariance in the focal plane. We will test this further with simulations, as described in Section 3.

2.3. *Space-Time KLIP*

Recall that KLIP improves contrast by projecting away features that are spatially correlated in image sequences. We can extend the framework of KLIP (Soummer et al. 2012) to space-time covariances by using an image sequence instead of an image. Note that for the following mathematics we assume discrete time sequences, rather than continuous as in Section 2.2 above. Additionally, we assume regular and continuous time sampling for this implementation; however, this method can be extended easily to block-continuous sampling, which may be useful in future work.

All variables for the following mathematics are also described in Appendix A. Baseline KLIP uses an image vector of length n_p (number of pixels in image) as its target image and a $n_p \times n_i$ matrix as the set of reference images to determine covariance between pixels, find eigenvectors of covariance, and project out the largest eigenvalue modes from the image. Similarly, space-time KLIP (referred to as stKLIP) uses an image sequence of length $n_s \times n_\mathrm{p}$ (number of images in the sequence times number of pixels per image), as shown in Equation 28, to perform those steps. Note that this is transposed compared to KLIP, which uses $n_\mathrm{p} \times n_s$.

It is then necessary to create a block diagonal covariance matrix of size $n_s \times n_\mathrm{p}$ by $n_s \times n_\mathrm{p}$, as illustrated in Figure 1, from the mean-subtracted image sequence. Each block is the covariance at a given time lag, with the block diagonal as lag zero (spatial covariance). If only lag zero is used, the mathematics here reduces down to baseline (spatial) KLIP, as described in Section 2.1.

Lags should be chosen based on the translation time of the smallest relevant feature within the field of view at the focal plane up to the full crossing time of the wind across the telescope aperture. This is an additional tuneable parameter to consider when optimizing the algorithm, in addition to the number of modes.

The following computations mirror baseline KLIP, but, in practice, are potentially more computationally expensive due to the larger size of the covariance matrix used in the eigendecomposition. The steps of stKLIP are as follows:

1. Subtract the mean image over the whole reference set, then partition the reference set into image sequences. These image sequences have length $n_s = n_l = 2L + 1$ where L is the largest number of timesteps (lags) away from the central image and n_l is the total number of timesteps (lags) in the sequence. (The following steps will be repeated over each image sequence, such that every image, with the exception of L images at each end, is at some point the central image. Therefore, for n_i images, there will be $n_i - 2L$ image residuals at the end of this process.)

   Similarly to KLIP, the reference set/target image set S (which in this implementation are the same) are unraveled into one-dimensional vectors $\boldsymbol{s}$ of length $n_s \times n_p$, as seen below.

$$S = \begin{bmatrix} S_{1,1} & S_{1,2} & \dots & S_{1,n_p} \\ S_{2,1} & S_{2,2} & \dots & S_{2,n_p} \\ \vdots & \vdots & \ddots & \vdots \\ S_{n_s,1} & S_{n_s,2} & \dots & S_{n_s,n_p} \end{bmatrix} \quad (28)$$

$$\boldsymbol{s} = \begin{bmatrix} S_{1,1} \\ S_{1,2} \\ \vdots \\ S_{1,n_p} \\ \vdots \\ S_{n_s,n_p} \end{bmatrix} \quad (29)$$

2. Compute the $[n_s n_p, n_s n_p]$ size covariance matrix C of the image sequences. In practice, this is more straightforward when done by computing the covariance of each image pair (C_i) and then arranging them in the block diagonal ordering shown in Figure 1.

3. Perform an eigendecomposition on the covariance matrix, obtaining $n_s n_p$ eigenvalues ($\boldsymbol{\lambda}$) and a matrix eigenvectors (V) of size $[n_s n_p, n_s n_p]$ contain-

Pupil plane view of turbulence, leading to the below image sequence

Input image sequence with length n_s=5, lags=[0,1,2,3,4], n_{lags}=5

Space-time covariance matrix with shape $n_{lags}n_{pix}$ x $n_{lags}n_{pix}$

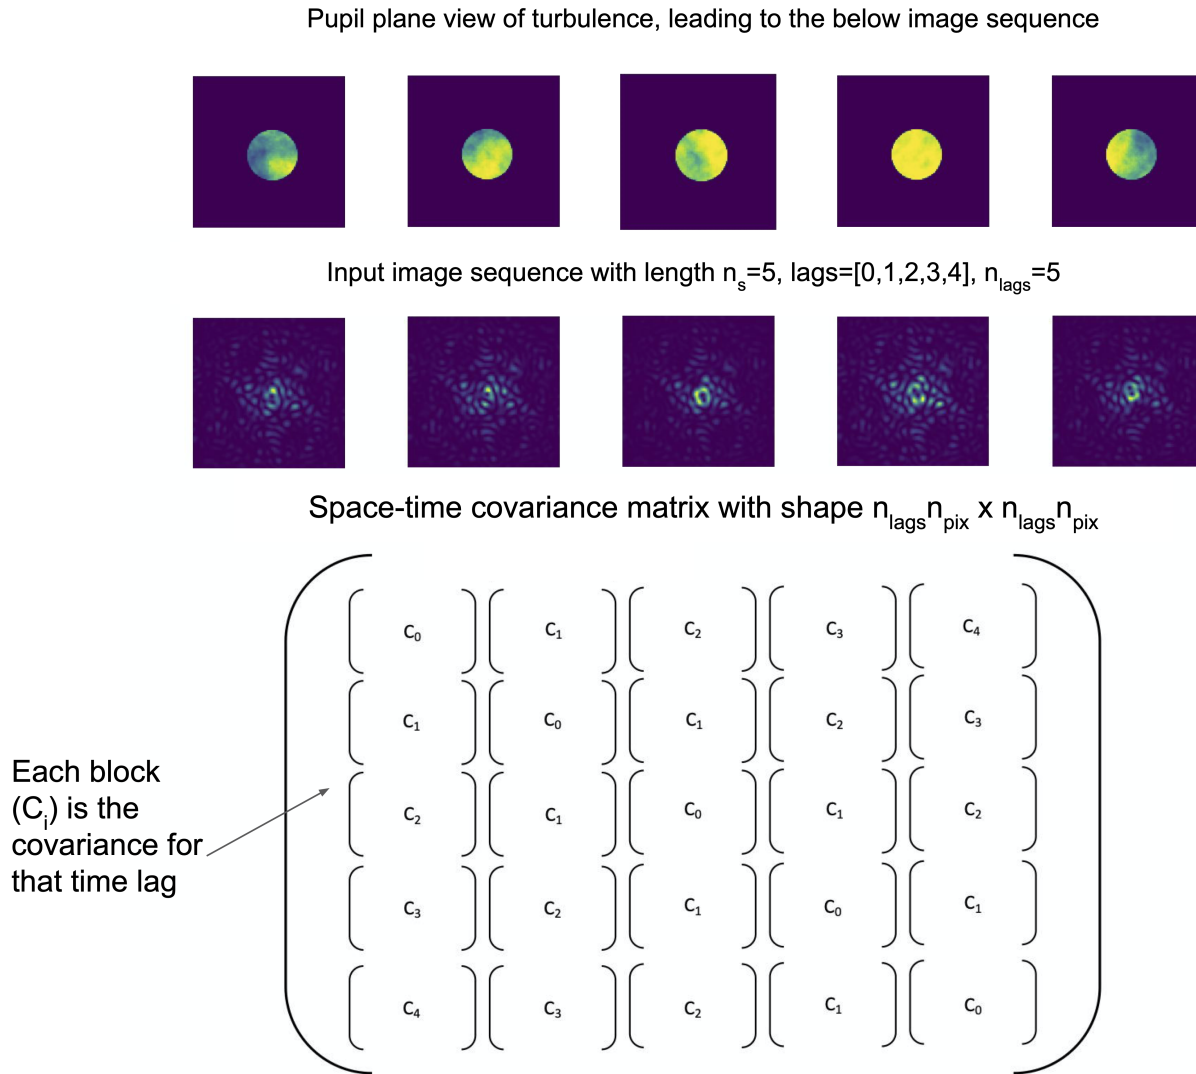Each block
(C_i) is the
covariance for
that time lag

Figure 1. Diagram of stKLIP input sequence setup – translating phase screens (top) and resulting image sequence (middle) – with the corresponding block diagonal space-time covariance matrix (bottom). Each covariance block C_i is the covariance for a single lag, with shape $n_p \times n_p$, and together they create a single space-time covariance matrix C with size $n_s n_p \times n_s n_p$. The covariance matrix takes this form because the 2d images are flattened into 1d vectors, which are then joined to make an $n_p \times n_s$ 1d vector, which is multiplied by its transpose to create this matrix.

ing individual eigenvectors $\boldsymbol{v}$.

$$Cv_j = \lambda_j v_j \qquad (30)$$

$$\lambda_1 > \lambda_2 > \lambda_3 > \dots \lambda_p \qquad (31)$$

each with length $n_s n_p$, such that $V_{i,j} = (\boldsymbol{v}_j)_i$.

$$V = \begin{bmatrix} V_{1,1} & V_{1,2} & \dots & V_{1,n_s n_p} \\ V_{2,1} & V_{1,1} & \dots & V_{2,n_s n_p} \\ \vdots & \vdots & \ddots & \vdots \\ V_{m,1} & V_{m,2} & \dots & V_{n_m,n_s n_p} \end{bmatrix} \qquad (32)$$

4. Choose a number of modes n_m, reducing the vector of eigenvalues and matrix of eigenvectors to sizes n_m and $[n_m, n_s n_p]$ respectively. The matrix of eigenvectors contains n_m rows of eigenvectors

5. Obtain image coefficients through inner product of chosen eigenvectors and image sequence, similar to

Equation 10.

$$\boldsymbol{q} = V \cdot \boldsymbol{s} = \begin{bmatrix} q_1 \\ q_2 \\ \vdots \\ q_{n_m} \end{bmatrix} \tag{33}$$

6. Project the image sequence back into pixel space to obtain a reconstructed sequence $\hat{\boldsymbol{s}}$ with central image $\hat{\boldsymbol{\psi}}_k$, again mirroring Equation 11. Note: For ease of implementation, we have calculated the entire sequence, but projecting only onto the central image may improve efficiency.

$$\hat{\boldsymbol{s}} = \hat{\boldsymbol{q}}^T \cdot V \tag{34}$$

$$\hat{\boldsymbol{\psi}}_k = [\hat{\boldsymbol{s}}_{n_p((n_l+1)/2-1)} \cdots \hat{\boldsymbol{s}}_{n_p(n_l+1)/2}] \tag{35}$$

7. Perform PSF subtraction using the central image.

$$\epsilon \boldsymbol{a}_k = \boldsymbol{s}_k - \hat{\boldsymbol{\psi}}_k \tag{36}$$

8. Iterate through the above steps such that each image is the central image of a sequence of length n_s, resulting in a set of residuals $\epsilon \boldsymbol{a}_{k,j} = [\epsilon_0 \boldsymbol{a}_{k,0}, \epsilon_1 \boldsymbol{a}_{k,1}, \ldots, \epsilon_{n_s} \boldsymbol{a}_{k,n_s}]$.

9. Compute mean of image sequence residuals to output an averaged residual, $\boldsymbol{r}_{k,\mathrm{avg}}$.

$$\boldsymbol{r}_{k,\mathrm{avg}} = \frac{1}{n_s} \sum_{j=0}^{n_s} \epsilon_j \boldsymbol{a}_{k,j} \tag{37}$$

Once our image sequence is projected into the new subspace in Step 6, we have two options for PSF subtraction: subtract the residuals from the whole sequence used, or subtract only from the central "target" image. We use a central target image to take advantage of speckle motions in timesteps both before and after. We then iterate through the full data set, as described in Step 8, performing stKLIP and PSF subtraction, so that each image is the central image of some image sequence with length $n_s = n_l = 2L + 1$. This outputs a sequence of image residuals that is of length $n_i - 2L$. In Step 9, we then average over the number of timesteps to output an averaged residual.

There are possibilities for improving the algorithm, such as by exploiting the symmetry in the covariance matrix C in order to hasten the process of updating the eigenimages; however, we leave this for future work. Further improvements are discussed in Section 5.

3. ALGORITHM DEVELOPMENT

In Section 2.2, we showed that we expect non-zero space-time covariance to exist in speckle noise. In Sections 2.1 and 2.3, we showed the mathematical framework for an algorithm to exploit these statistics for image processing and PSF subtraction.

In this section, we illustrate aberrations of increasing complexity to examine their covariance structure and test the application of stKLIP. These tests and simulations are described in 3.1, for initial proof of concept. Section 3.2 describes the algorithm application to simulated data and calculations of possible contrast gains in the algorithm's current form; here we also discuss selection criteria for the choices of number of modes and lags. Analyzing these data sets also requires some computational optimization, which is described in 3.3. In the following Section 4, we will discuss the results of these applications of stKLIP.

3.1. *Foundational Tests*

Our first step was to create and implement simple test cases in one and two dimensions to demonstrate that our theoretical expectations from Section 2.2 are valid and ensure that our algorithm reduced image variance as expected. A one-dimensional case allows us to directly compare a simulated covariance matrix with one calculated from the analytic theory in Section 2.2, serving as a test of the relationship between pupil plane covariance and focal plane covariance. Then, a two-dimensional case serves as a first in implementing the algorithm, ensuring that the algorithm reduces variance on a well-understood simple case before moving onto more complex atmospheric simulations.

3.1.1. *One-Dimensional Test of Pupil/Focal Covariance Relationship*

To begin, we created a simple one-dimensional model of two interfering speckle PSFs, which are simply two sinusoids with slightly different frequencies in the pupil plane. We first use this simple sinusoidal model to compare the simulated space-time covariance to the predicted behavior from theory, to show how a set of input aberrations in the pupil plane corresponds with the resulting focal-plane space-time covariance. Although the algorithm does not require pupil plane covariances, this test is done to further establish the existence of the focal plane covariances that we seek to harness.

To create the 1-d speckle model, first we must create a grid setup for evaluating the wavefront in the pupil and focal planes. These are parameterized in units of D/λ and λ/D respectively, where λ is our wavelength of observation, assuming monochromatic light. Keeping

these units preserves the Fourier duality relationship, and they can be converted to more conventional units if the focal length is known.

The next critical piece is to define the entrance aperture in the pupil plane. This pupil function sets the amplitude A of the electric field ($E = Ae^{i\phi}$), and is simply a top-hat function ($\Pi(u)$, 1 inside a given region and 0 outside). We also apply a translating phase screen (shown in the top panel of Figure 2) to the pupil, which is where phase aberrations are accounted for. We use a simple perturbation of two superimposed sinusoids with similar periods/frequencies, so that the wings of their PSFs overlap. This set-up is like simulating one layer of frozen flow translating across the telescope's aperture. These perturbations are small ($\ll 1$ radian), consistent with the high-contrast regime.

We then perform the necessary Fourier transform to retrieve the focal-plane electric field. By doing this for the pupil function with no perturbations, we retrieve what we would see in an ideal case for a uniformly illuminated pupil; this is also what would be blocked if we had a perfect coronagraph. We subtract this "perfect" case from the case with the sinusoidal perturbation, performing the action of the coronagraph and suppressing light from the unaberrated portion of the wavefront.

A one-dimensional case (Figure 2) illustrates the relative evolution of two neighboring speckles created from atmospheric perturbations. Atmospheric theory (as in Section 2.2), in particular the frozen flow assumption, predicts a symmetrical space-time covariance structure, which can be computed for a 1-d model with a top-hat pupil function ($\Pi(u)$), two sinusoidal functions in the pupil, and no uniform illumination in the pupil ($C(x) = 0$). We carried out these calculations in two ways. First, we solved the integrals in Section 2.2 for the simple two sinusoid situation using Fast Fourier Transforms (FFTs). Second, we began with an array describing the sinusoidal "phase screen" and simulated propagation through an optical system using FFTs.

The variation in pupil and focal plane covariance over various time lags, as shown in Figure 3, can be clearly interpreted based on the locations of the two interfering speckles. These matrices show a symmetric pattern that changes with the number of lags used, due to the change in the speckles' relative locations. At lags 0 and 100, the peaks are due to the alignment of the speckles' peaks, as marked in the top panel; lag 25 illustrates the lower covariance when the speckles are in slightly different places, and lag 50 shows two lower intensity peaks when the speckles are separated. Importantly, for a given non-zero lag, there are non-zero terms in both
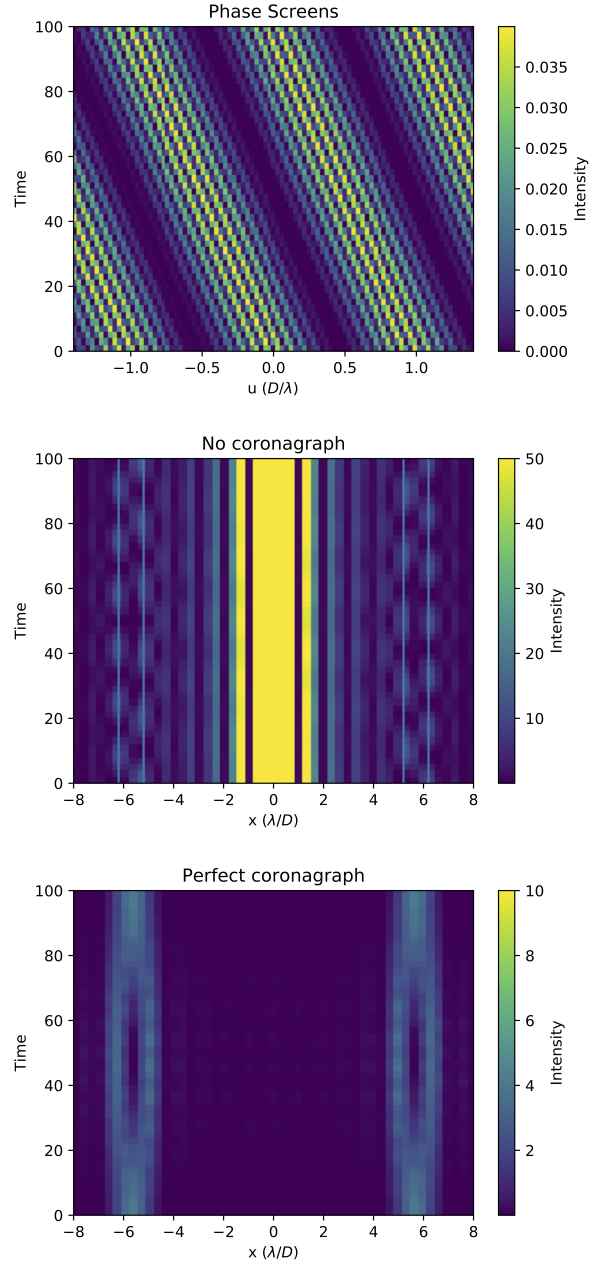

Figure 2. One-dimensional demonstration of speckle interference. Two sinusoidal perturbations in the pupil plane interfere to create moving speckles in the image plane. Top: 1d phase screen with interfering sinusoids over time. Middle: 1-d intensity over time without a coronagraph, showing the Airy pattern. Bottom: 1-d intensity over time with a coronagraph, with the speckles' relative evolution appearing more clearly due to the lack of coherent light, $C(\boldsymbol{x})$. This simulation is used as a test of the space-time speckle covariance theory in Section 2.2.
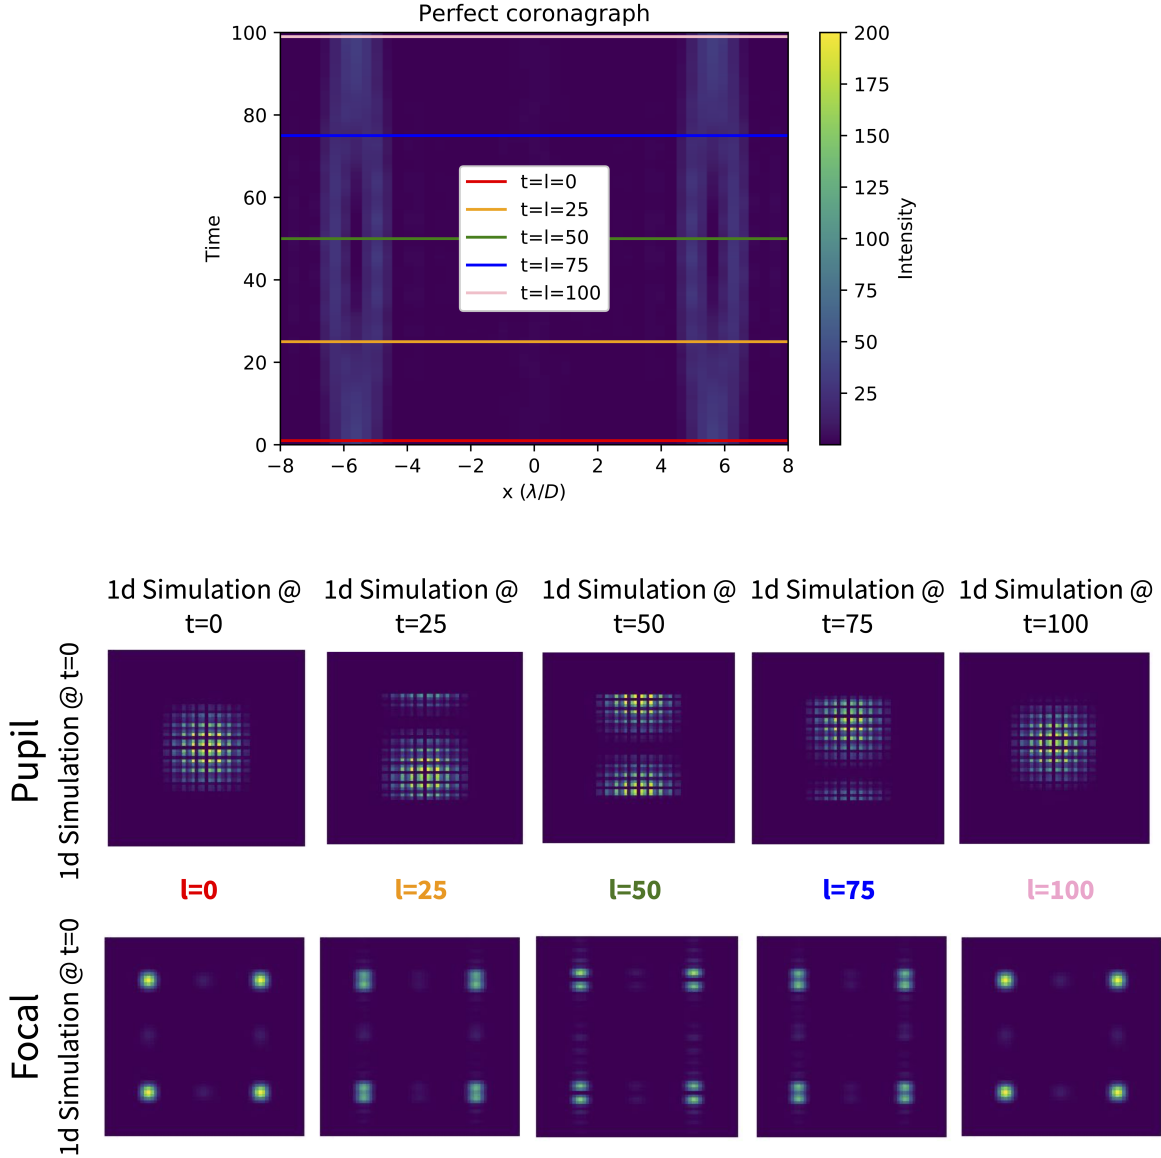
Figure 3. Space-time covariance matrices for pupil plane (middle) and focal plane (bottom) of a 1-d model of two sinusoids with different frequencies – as illustrated in the top panel of Figure 2 – with an annotated view of the simulation (top). These matrices show a symmetric pattern that changes with the number of lags used, due to the change in the speckles' relative locations. At lags 0 and 100, the peaks are due to the alignment of the speckles' peaks, as marked in the top panel; lag 25 illustrates the lower covariance when the speckles are in slightly different places, and lag 50 shows two lower intensity peaks when the speckles are separated. Importantly, for a given non-zero lag, there are non-zero terms, indicating that there are temporal correlations.

the pupil and focal plane covariances, indicating that there are temporal correlations.

This simulation further demonstrates the claim that a simplified frozen flow scenario in the pupil can create calculable space-time covariances in the focal plane, and validates our use of this simple test case to test stKLIP.

3.1.2. *Two-Dimensional Test Case for Algorithm Development*

In order to ensure that the algorithm is behaving according to our expectations – that it will reduce the image variance – we expand this one-dimensional test case into two-dimensions to make an image sequence of the two time-varying, interfering speckles. We use

this idealized test case as a check against our expectations for our stKLIP implementation, and for a first test of efficacy, comparing the reduction in image variance between three data processing methods: mean-subtraction, KLIP, and stKLIP. The setup is the same as the above one-dimensional test case, but in two dimensions, with a circular aperture instead of a top hat as the pupil function. We create a series of images at various time steps as the input to stKLIP, shown in Figure 4.

Although there are two tuneable parameters for stKLIP — number of modes (e.g. number of eigenimages used in the projection) and number of lags, as described in Sections 2.1 and 2.3 — we only test one set of modes and lags (10 modes, 2 lags) with this simple test case and leave further exploration of these parameters for later testing (see Section 3.2). We similarly use 10 modes for KLIP to make the comparison fair.

In this simple test case, KLIP and stKLIP reduce the variation in the image by factors of 6.8 and 5.7, respectively. Although stKLIP does not improve upon KLIP in this limited test case, it is important to remember that we have not optimized for modes and lags in this scenario; determination of performance is left for more rigorous and realistic tests in the following section, 3.2. They both outperform simple interventions, such as subtracting the mean of the image, in reducing the total variation in the image, as shown in Figure 5. To summarize, this 2d test was performed to demonstrate that the overall image variance decreases after projecting out modes of variation with stKLIP, as qualitatively expected, and in that sense the test can be considered successful.

3.2. *Simulated AO Residual Tests*

We then wanted to test stKLIP on a more realistic atmospheric phase screen and again measure potential contrast gains. To this end, we created a set of simulated observations to represent AO residuals and performed stKLIP on them for a variety of different modes and lags. We measure contrast curves and companion SNR for four methods of post-processing in order to understand the effectiveness of our new method: stKLIP, baseline/spatial KLIP, mean-subtraction, and no post-processing. Results from these tests are described in Section 4 and discussed further in Section 5. In this section, we first detail the methods used to create the simulated data set, then the methods for computing contrast curves and SNR on the processed data.

To create the simulated data set, we use a simulator specifically designed for high-contrast imaging with next-generation detectors, such as MKIDs, called MEDIS (the MKID Exoplanet Direct Imaging Simulator), the
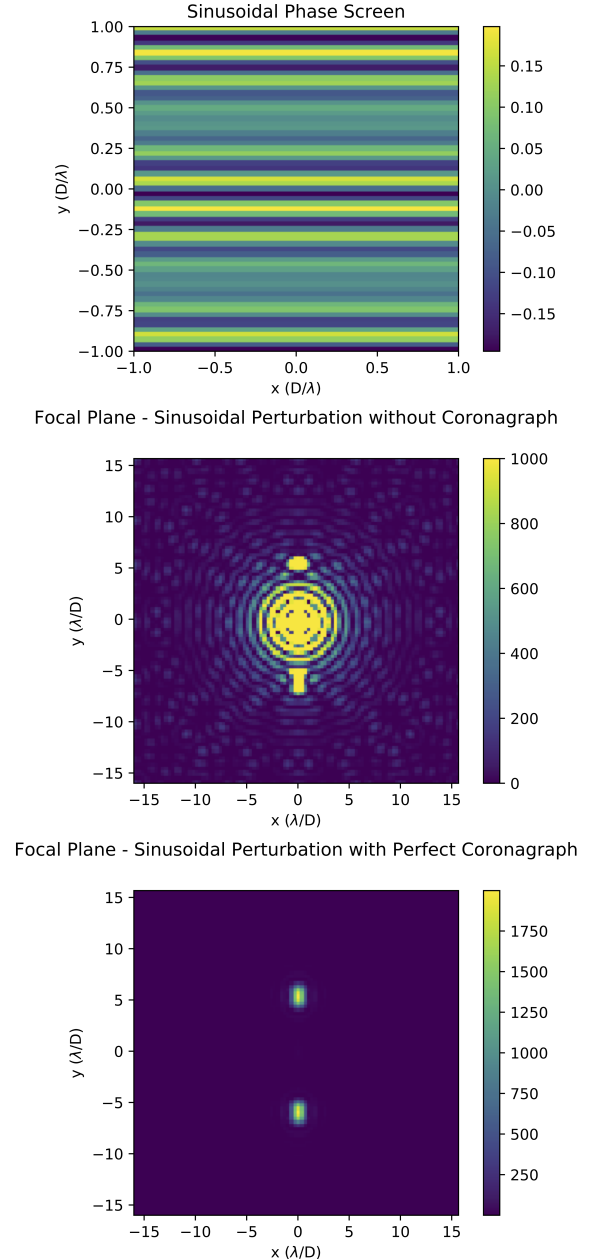

Figure 4. Two-dimensional test of speckle interference. A sinusoidal phase screen (top) produces a speckle pattern imposed on an Airy disk (middle). Subtracting the PSF of a model without perturbations, we simulate observations of this sinusoidal perturbation with a "perfect" coronagraph (bottom). All images depict the intensity ($I = |E|^2$). This simulation is used as a troubleshooting step for a first implementation of the stKLIP algorithm.

first end-to-end simulator for high contrast imaging instruments with photon counting detectors (Dodkins 2018; Dodkins et al. 2020).

MEDIS generates atmospheric phase screens with HCIPy (Por et al. 2018). These phase screens use mod-
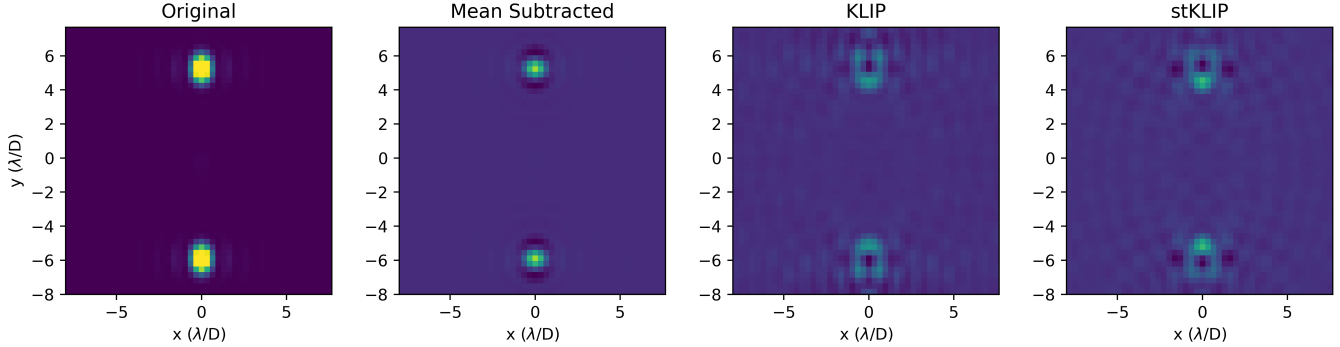
Figure 5. One frame of the input sequence (left) for the simple two-sinusoid test case with a coronagraph, with the residuals after PSF subtraction using mean-subtraction, KLIP, and stKLIP, showing a clear reduction in speckle intensity. Both stKLIP and baseline KLIP reduce image variance by a factor of at least 5.7 from the original image, an improvement over simple interventions like mean-subtraction. Although stKLIP does not improve upon KLIP in this limited test case, it is important to remember that we have not optimized for modes and lags in this scenario; this step was intended for troubleshooting, not rigorous characterization of the algorithm.

els of Kolmogorov turbulence, and we use the simplest option of a single frozen flow layer. Then, `MEDIS` uses `PROPER` to propagate the light through the telescope under Fresnel diffraction, including both near- and far-field diffraction effects (Krist 2007). Separate wavefronts are propagated for each object in the field — the host star, and any companion planets. `MEDIS` also includes options to introduce coronagraph optics, aberrations (like non-common path errors), and realistic detectors. `MEDIS` outputs the electric field or intensity at specified locations in the optical chain, such as the pupil and focal planes in our case, as shown in Figure 6.

Given the wide range of parameters available in `MEDIS`, we had to make decisions on what to use for the `MEDIS` simulations used to test stKLIP. For these simulations, we implement a telescope with 10 meter diameter, similar to the Keck Telescopes. We begin with a case without adaptive optics for simplicity. For this, the sampling rate needs to be a few milliseconds, a few times over-sampled compared to the smallest temporally resolvable features given the field-of-view (FOV) under consideration. The number of frames is chosen to create a total observation time of 30 seconds (6,000 frames at 0.005 second sampling) to recreate a realistic observation and attain a sufficient number of independent samples. The grid size is significantly larger than the area of interest (256×256 pixels) to avoid edge effects. However, we choose a region size / FOV that is significantly smaller than our whole grid (100×100 pixels) to make this problem more computationally tractable.

The simulation includes atmospheric parameters, such as the Fried Parameter (r_0), a length scale for coherence in the atmosphere, and the structure constant (C_n), a description of turbulence strength over multiple atmo-
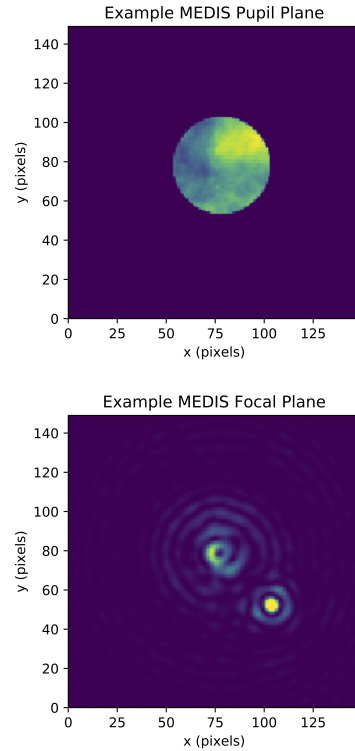

Figure 6. Examples of MEDIS simulations. (Top) Pupil plane, illustrating the phase screen. (Bottom) Focal plane, with a clearly bright companion object. These simulations are used as a preliminary test of stKLIP's efficacy and potential; however, there is a large parameter space to explore beyond the scope of this work.

spheric layers. The atmospheric model we use is a simple single layer of extremely mild Kolmogorov turbulence, with $r_0 > 10$ m, since we want $r_0 \gg D$ to stay in the high-contrast regime of small phase errors. Note:

this simulated atmosphere is not realistic in ground-based imaging, but we chose these parameters to approximate the high-contrast regime without simulating adaptive optics and introducing additional parameters. While our numerical experiments will depend on the input power spectrum, our primary aim was to assess the characteristics of a second-order statistical analysis of the linearized system (Equation 13), rather than impacts of the particulars of the wavefront error power spectrum. It is worth exploring how different atmospheric conditions (e.g. a smaller r_0 value) would change the effectiveness of this method, but that is beyond the scope of this initial investigation.

We choose a vortex coronagraph (Mawet et al. 2009), since it is the closest to an "ideal" coronagraph of the options available in MEDIS (e.g. closest to perfect cancellation of the spatially coherent wave), thanks to its small inner working angle (Guyon et al. 2006). We want an ideal detector since, for this initial investigation, we are not yet interested in how detector noise/error affects this method. We also include one companion object that would be readily detectable given current capabilities (a contrast of 5×10^3), in order to enable SNR measurements of an injected companion for various post-processing methods including stKLIP. As mentioned in Section 2.2, lags should be chosen based on crossing times and relevant features. In these simulations, this ranges from 2 to 10 timesteps (0.01 to 0.05 seconds) for a wind speed of 5 m/s and 5 millisecond sampling. Future work should test a further range of lags, up to 400 timesteps (2 seconds, or one full crossing time), but our current method is computationally limited as mentioned in Section 3.3. In this investigation, we also test a range of modes from 1 to 500.

Although these simulations are computationally expensive, MEDIS is capable of parallel processing, except in cases where AO parameters require serialization. We take advantage of this capability by using UCLA's Hoffman2 Cluster. The resultant data sets are quite large, and require inventive ways of computing the necessary statistics without loading the full array into memory, described further in Section 3.3. These simulations show us how realistic space-time covariance differs from the idealized case, and allow us to begin to test the effectiveness of our new method.

Metrics of efficacy used in this study are measurements of variance, signal, noise, signal-to-noise ratios (SNR), and contrast curves. Variance is simply computed over the whole 100×100 pixel residual image using numpy.var. Signal is computed using aperture photometry (via photutils), centered on the simulated companion. Noise is similarly computed using aper-

ture photometry by taking the standard deviation of a series of apertures in an annulus at the same separation as the simulated companion. SNR is then the ratio of these two measurements. Contrast curves are estimated using aperture photometry at various distances from the image center and dividing by the aperture photometry measurement of the unmasked (e.g. no coronagraph) peak, then adjusting by the signal throughput; the throughput here is estimated as the signal after processing divided by the signal before data processing. These various metrics are computed for the original images, as well as different post-processing scenarios, to understand the relative efficacy of stKLIP. Results are described in Section 4.

3.3. *Iterative Statistics Calculations*

There are two key computational challenges for large data sets such as those produced by MEDIS: memory access and computational complexity. Simulations with MEDIS for a realistic observing sequence based on our criteria above can be on the order of 100GB, which can pose challenges to RAM-based manipulation for the calculation of mean and covariance given our current computing resources. To address this problem, we implemented the framework for iterative statistics calculations set forth in Savransky (2015).

In order to perform a KLIP-style calculation, we first need to compute second-order statistical quantities for a data set of n samples $\boldsymbol{x}_i$, such as the mean and covariance. The formula for the calculating mean is:

$$\boldsymbol{\mu} \equiv \frac{1}{n} \sum_{i=1}^{n} \boldsymbol{x}_i \qquad (38)$$

When the mean μ is estimated from the data, the sample covariance can be calculated as follows:

$$C \equiv \frac{1}{n-1} \sum_{i=1}^{n} (\boldsymbol{x}_i - \boldsymbol{\mu})(\boldsymbol{x}_i - \boldsymbol{\mu})^T. \qquad (39)$$

These sums can be broken up into smaller iterative steps k, to make the calculation less memory intensive. For each step k, the mean can be updated with the formula

$$\boldsymbol{\mu}_k = \frac{(k-1)\boldsymbol{\mu}_{k-1} + \boldsymbol{x}_k}{k} \qquad (40)$$

and the covariance can be updated by

$$S_k = \frac{k-2}{k-1} S_{k-1} + \frac{k}{(k-1)^2} (\boldsymbol{x}_k - \boldsymbol{\mu}_k)(\boldsymbol{x}_k - \boldsymbol{\mu}_k)^T. \qquad (41)$$

However, Equation (41) is only applicable to the spatial covariance, e.g. a time lag of zero. The space-time covariance can be calculated as

$$S_l = \frac{1}{n-l-1}\sum_{i=1}^{n}(\boldsymbol{x}_i - \boldsymbol{\mu})(\boldsymbol{x}_{i-l} - \boldsymbol{\mu})^T. \qquad (42)$$

Following a similar protocol to Savransky (2015), we derived an update formula for the space-time covariance:

$$S_l = \frac{1}{n-l-1}\left[\sum_{i=l}^{n}\boldsymbol{x}_i\boldsymbol{x}_{i-l}^T - (n-l)\boldsymbol{\mu}\boldsymbol{\mu}^T \right.$$
$$\left. + \boldsymbol{\mu}^T\sum_{i=1}^{l-1}\boldsymbol{x}_i + \boldsymbol{\mu}\sum_{i=n-l-1}^{n}\boldsymbol{x}_i^T - 2l\boldsymbol{\mu}\boldsymbol{\mu}^T\right] \quad (43)$$

It is identical to Equation (41), except for the last 3 additional cross-terms. These cross-terms were directly calculated and determined to be negligibly small as the sample size becomes large relevant to the maximum lag, and thus would only be relevant in edge cases. For 1,000 samples, the error on the space-time covariance calculation is on the order of $10^{-4}\%$ or less. For 10,000 samples, the error decreases to 10^{-6} to $10^{-7}\%$, indicating a trend of decreasing error for an increasing number of samples. We do not plan to use fewer than 1,000 samples in a data set, so we consider this approximation to the space-time covariance acceptable and have implemented it for the tests described in Section 3.2.

Although the mathematics laid out in this section make covariance calculations possible, the resulting covariance matrices can be quite large, on the order of 10GB for even short test cases with small FOVs. Even with sufficient RAM for manipulation, these large covariance matrices can lead to long computation times for following steps of the algorithm. The image size and sequence length of data sets used in our stKLIP method is therefore still currently limited by memory requirements and prohibitively long execution times. This is mostly due to the eigendecomposition calculations, since the full space-time covariance matrix needs to be loaded into memory for input into `scipy.linalg.eigh`. As we proceeded with larger data sets, we chose to perform a standard eigendecomposition with `scipy.linalg.eigh` using the default backend (C LAPACK `evr`) but limited the maximum number of eigenvalues/eigenvectors computed, since many of the smaller eigenvalues only capture noise and are not necessary for this process. There may be more optimal choices for the eigendecomposition algorithm, but such optimization is left for future work.

Another possible solution to mitigate this bottleneck would be using an iterative eigendecomposition. This could theoretically be done with the NIPALS (Nonlinear Iterative Partial Least Squares) algorithm (Risvik 2007). However, applying the NIPALS algorithm is not straightforward for this problem; our space-time covariance matrix is currently assembled from various spatial covariance matrices, and considerable changes would need to be made to NIPALS to accommodate a space-time calculation instead of a solely spatial one, since the NIPALS algorithm relies on a data matrix as input instead of a covariance matrix. Future iterations of this algorithm could also make use of the `dask` package for parallelization of computations to help speed up run time, but as of this writing an eigendecomposition function (e.g. `dask.linalg.eigh`) was not yet implemented, although the similar `dask.linalg.svd` function could possibly be used. We leave such improvements in efficiency for future work.

4. ALGORITHM PERFORMANCE ON SIMULATED AO RESIDUAL DATA

We have confirmed through theory (§2.2) and simulation (§3.1) that space-time covariances exist for speckles in a simple high-contrast imaging system in the regime of small phase errors and short exposures. In Section 2.3, we defined a new algorithm, similar to Karhunen-Loéve Image Processing, to take advantage of space-time covariances and improve final image contrast, with the eventual goal of detecting fainter companion objects. As shown in Section 3.2, we have developed an initial implementation of this space-time KLIP (stKLIP) algorithm, and demonstrated it on simulated data. In this section, we present the results of those demonstrations. It is worth noting that these tests on simulated data only explore a small range of parameter space, and are not indicative of the absolute potential of using space-time covariance in data processing. Instead, we present this as a first proof-of-concept for the possibility of this new method.

An example of the images input to and output by the stKLIP processing algorithm is shown in Figure 7, along with a comparison to two other data processing interventions, mean-subtraction (as in Equation 3) and KLIP. For this simulated data, mean-subtraction makes such a slight improvement that in the following figures we omit it from comparison plots, as it would be almost precisely coincident with the original image's metrics.

To quantitatively measure the efficacy of our stKLIP data processing algorithm, we computed total image variance, signal-to-noise ratios, and approximate contrast curves, as described in Section 3. To further determine the utility of this algorithm and characterize its dependence on the tuneable parameters, we also in-
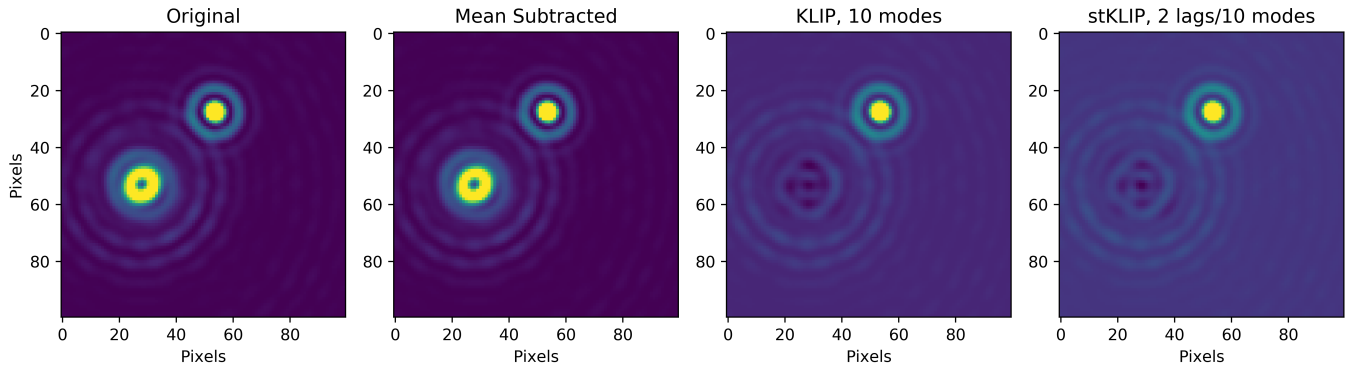
Figure 7. One frame of the input sequence (left) from `MEDIS`, with the residuals after PSF subtraction using mean-subtraction, KLIP, and stKLIP. Both stKLIP and baseline KLIP reduce image variance by a factor of ~1.85 from the original image for the listed case of 10 modes and 2 timesteps lag in stKLIP.

vestigated the relationships between these efficacy metrics, the number of KL modes used, and the number of stKLIP lags used. We leave adjustments of the residual wavefront error statistics and companion location, among other parameters, and their effects on stKLIP's efficacy for future work.

Image variance is a primary metric for subtraction efficiency. Total image variance is reduced by almost half for both spatial / baseline KLIP and stKLIP within the first 10 modes, and variance approaches 0 around 50 modes. In this test, spatial and stKLIP are similar in their variance reduction abilities, and are both improvements on mean-subtraction and the original image. Image variance drops off steeply within the first 20 modes, indicating that most of the power is removed with only a few eigenimages required in the reconstruction. Given that only a small number of modes are required to remove the majority of the variance in the image, future applications of this algorithm could exploit this fact to reduce the computational burden by only calculating the first n eigenvalues/eigenimages.

For both KLIP and stKLIP on these simulated data, signal starts to be lost around 5–10 modes and drops off more steeply after ~30 modes. Space-time KLIP with 4, 5, 6, or 8 lags in this scenario shows a slight edge over baseline KLIP in signal retention, as shown in Figure 8. It is worth noting that the choice of optimal number of lags depends on the wind speed and region in the image that we are most interested in. Recall from Section 3.2 that this test uses $v = 5\,\mathrm{m/s}$, and the companion location can be seen in Figure 7. Noise reduction capabilities appear very similar between KLIP and stKLIP; after about 40–50 modes, so much of the image has been removed that noise approaches zero and shows small random fluctuations, indicating that these higher modes contain less information.
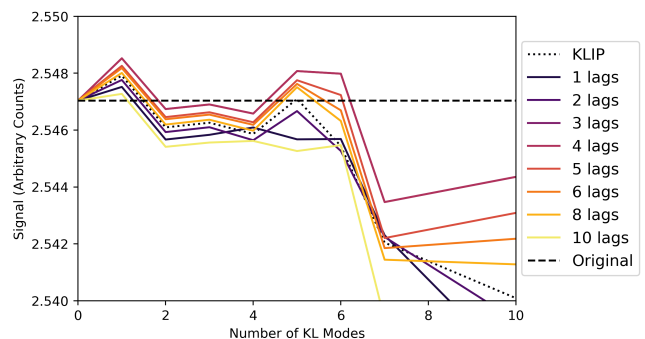

Figure 8. Companion signal over number of KL modes used in the model PSF subtraction; this figure shows that signal loss begins around 5 modes, indicating that future iterations of this algorithm would benefit heavily from implementing measures to prevent self-subtraction. Certain choices of lag (4, 5, 6, 8) show a minor improvement in signal retention over spatial (lag = 0) KLIP.

Signal-to-noise ratio (SNR) shows a 10–20% improvement over the original image within the first 40 modes, as shown in Figure 9. The 2nd peak in Figure 9 is possibly due to small number statistics (most of the signal has been removed by then) and not a real SNR improvement. It is worth noting that the SNR shown here could improve significantly if a method is implemented to retain signal and improve throughput, which we discuss more in the following section. We again see that there is a slight advantage for certain lags over spatial (lag=0) KLIP on the order of a few percent, indicating that there is possibility for properly tuned stKLIP to outperform KLIP.

Contrast curves (as shown in Figure 10) similarly show potential for up to 50% improvement depending on the number of modes, lags, and region of the image. Within 20 pixels, we see potential for up to 400% improvement, but with the caveat that this close to the coronagraphic
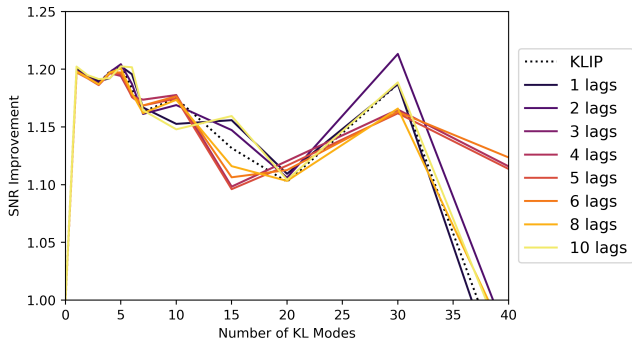
Figure 9. Companion signal-to-noise ratio (SNR) compared to the original image SNR over number of KL modes used in the model PSF subtraction; this figure shows a 10–20% improvement over the original image using stKLIP and KLIP, with stKLIP having a slight edge (on the order of a few percent) for certain choices of lag.

mask, measurements of SNR and contrast are less reliable. A slight spread in the contrast curves for various lags, such as that seen around 30–40 pixels for 5 modes in Figure 10, indicates that it is necessary to strategically choose the number of lags used in stKLIP depending on the image region in which we want to optimize contrast. We will discuss these results and future work further in the following section.

5. DISCUSSION

Overall, our tests on simulated data (Section 4) show that there is a demonstrated contrast gain (or equivalently, SNR improvement) of at least 10–20% from the original image using stKLIP with fewer than 40 modes. There is also evidence that stKLIP provides a slight advantage over spatial-only KLIP for certain choices of number of lags, number of modes, and location in image. However, the real potential for this method will be unlocked when the technique is safeguarded against self-subtraction and demonstrated on real data.

In this section, we first discuss how well the signal is retained for this new algorithm, and possibilities for future improvements to better avoid self-subtraction and retain signal in Section 5.1. Next, we discuss the relationship between the lag parameter and the optimized region of the target image in Section 5.2. Then we consider the addition of quasi-static speckles to our currently idealized, only atmospheric speckle regime in Section 5.3. Lastly, we propose other considerations for future work and implementations of this algorithm in Section 5.4.

5.1. *Signal Retention*

The signal clearly decreases beyond ∼5 KL modes as shown in Figure 8, indicating that we are not only sub-

tracting from the noise but also the companion (known as self-subtraction). If we can find a way to reduce this self-subtraction and retain signal, we could potentially further improve the contrast gain. This could possibly be accomplished by masking the location of the planet or excluding regions with high spatial covariance but low temporal covariance, but further development is needed to enable this functionality. Depending on the masking implementation, this data processing method could be used for blind searches or characterization observations. In fact, it may be particularly suited to characterization observations due to the dependence on a specific image region from the nature of atmospheric speckles.

Based on previous work on LOCI (Locally Optimized Combination of Images) (Lafreniere et al. 2007; Marois et al. 2014; Thompson & Marois 2021), we can expect additional contrast gains once masking is implemented. Additionally, there are other techniques used for KLIP to differentiate between signal and speckles, such as angular differential imaging (ADI, Marois et al. (2006a)), spectral differential imaging (SDI, Marois et al. (2005)), and reference differential imaging (RDI, Marois et al. (2003)). Similar efforts to increase the distance between the signal and the noise in the eigenimages may be useful for stKLIP.

Additionally, when the number of lags is zero, stKLIP simply reduces to baseline KLIP (Soummer et al. 2012) as mentioned in Section 2.3, and we have included spatial-only/baseline KLIP as a comparison for stKLIP in our analyses. It is worth noting, however, that KLIP is typically used on long-exposure images, a different regime than that for which stKLIP is useful. Additionally, we are comparing stKLIP to KLIP with no self-subtraction mitigation. Most current implementations of KLIP, such as pyKLIP (Wang et al. 2015), do have some sort of self-subtraction mitigation or method to increase spatial diversity implemented, such as forward modeling, angular differential imaging, or spectral differential imaging (Pueyo 2016; Marois et al. 2006b; Vigan et al. 2010). Therefore, in practice, KLIP would currently have a significant advantage over stKLIP as implemented in this work. However, future work can adapt many of the existing methods and techniques from KLIP to improve the implementation of stKLIP and its resulting performance.

5.2. *Optimization for Lags and Image Region*

Despite KLIP's apparent advantages, it appears that, depending on the number of lags used and the location in the image, stKLIP can outperform KLIP by a few percent without self-subtraction implemented for either case as is done in our test. This is evident in Figure
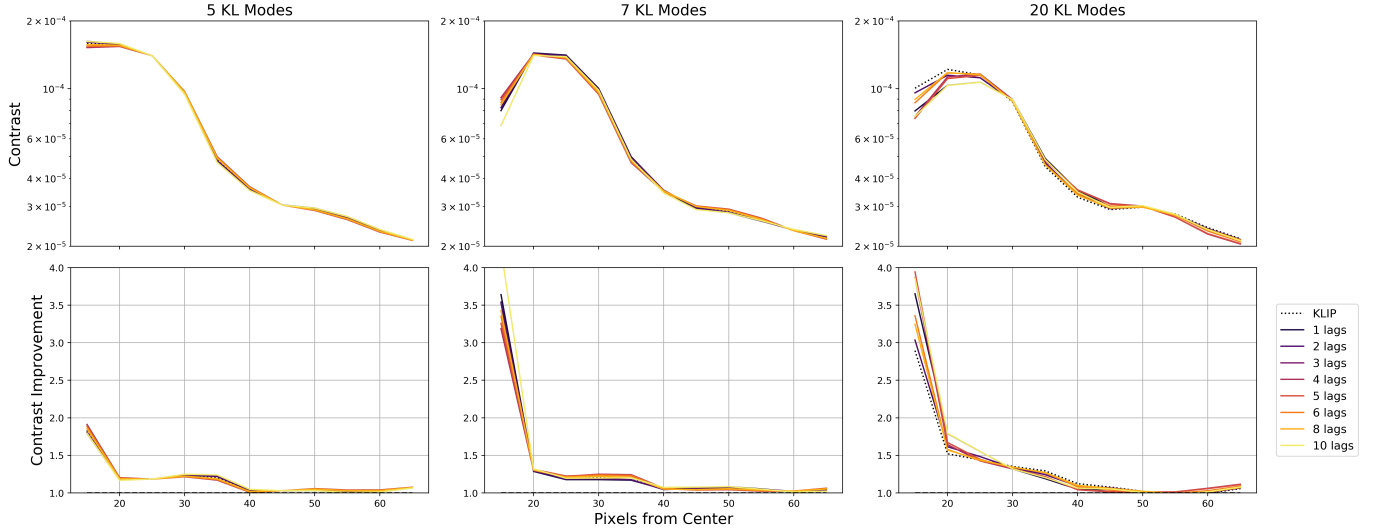
Figure 10. Contrast curves, as well as contrast improvement (a comparison to the original image's contrast curve), for three cases of KL modes: 5, 7, and 20. Each shows results for the image processed with baseline KLIP (0 lags) as well as stKLIP with a variety of lags. stKLIP is consistent with KLIP improvements, and in certain regions may show improvements depending on number of lags used.

10, showing detail of the region with highest contrast gain (other than near the central mask). The region of highest contrast gain will vary depending on the chosen lag as well as the atmospheric conditions creating the speckles in question. Optimization of input parameters is a notoriously tricky problem for KLIP (Adams et al. 2021), and it appears stKLIP is subject to the same challenges.

The variation of optimal lag and image region is due to the relationship between the wind speed and spatial frequency, since wind speed and telescope diameter combine to determine the crossing time for one cycle of the spatial frequency as $t_{\mathrm{cross}} = d_{\mathrm{telescope}}/v_{\mathrm{wind}}$. Spatial frequency in the pupil then corresponds to a location in the image plane. The effect of atmospheric parameters on speckle properties is further quantified in Guyon (2005) and speckle lifetimes are observed on shorter scales in Goebel et al. (2018). Empirical investigations of telescope telemetry and ambient weather conditions are also an ongoing area of study, especially with regards to predictive control (Guyon et al. 2019; Rudy et al. 2014; Hafeez et al. 2021), but that information may additionally be useful in determining optimal parameters for stKLIP on-sky. Additionally, using this information on the temporal/spatial locations of strongest correlations, it may be possible to reduce the matrix size or use only the most correlated images such as in T-LOCI (Marois et al. 2013).

For this work, we have been operating in the regime of milliseconds to track atmospheric speckle motions. However, in practice, the full 3D space-time correlation matrix will have power on multiple timescales, from that of atmospheric speckles to quasi-static speckles. It is outside the scope of this work to fully explore how space-time KLIP could be applied on multiple time domains, and there is additionally the caveat that computational complexity grows with longer timescales than those we have applied here.

5.3. *Including Quasi-Static Speckles*

As mentioned in Section 1, the scenario we have investigated is an idealized case — one in which quasi-static speckles are absent and our images are dominated entirely by atmospheric speckles. We are also working on short timescales, where the atmosphere is frozen at each time step. There is a timescale over which the intensity changes, which we are observing in this scenario, but there is also a timescale for changes in the electric field's phase. These phase changes will only result in changes in intensity if superimposed onto a constant electric field, such as the case of non-coronagraphic imaging, or when quasi-static speckles are significant ($C(\boldsymbol{x})$ in Equation 15). This is another regime in which to explore algorithm performance, wherein quasi-static and atmospheric speckles co-exist and interact, possibly even changing the speckle lifetimes (Soummer & Aime 2004; Fitzgerald & Graham 2006; Bloemhof et al. 2001; Soummer & Aime 2004). In this regime, there will likely be additional space-time variation as "pinned" speckles oscillate. Given that the presence of quasi-static speckles will make visible the additional space-time variations in phase, it is possible that stKLIP will operate even more effectively with this additional information to exploit. However, additional quasi-static speckles will lead

to additional photon noise, which may counteract any theoretical contrast gains from including phase information. (Note: recent work from Mullen et al. (2019) shows that using KLIP on shorter exposures may even help remove quasi-static speckles more effectively, further bolstering the case for the stKLIP's effectiveness in this regime.) Additionally, the presence of atmospheric residuals could even provide information about the phase of quasi-static speckles, allowing them to be effectively nulled with a deformable mirror (Frazin 2014; Frazin & Rodack 2021). Future simulations may explore this regime and determine if additional contrast gain is possible.

5.4. *Considerations for Future Work*

In this idealized test case, we also chose not to simulate adaptive optics corrections, instead leaving an investigation of how AO parameters affect space-time correlations and the resulting stKLIP processing for a future investigation. Since AO suppresses low frequencies and heaves high frequencies unchanged, although our total error is on par with an AO residual scenario, the overall shape of the power spectrum would be different. This would likely lead to weaker temporal correlations with AO. Previous work also shows that AO corrections do affect the lifetime of speckles (Males et al. 2021; Males & Guyon 2017), so this will be an important factor to consider in future work.

Currently, we have yet to demonstrate the full potential of this algorithm, in part due to the high computational costs. To run stKLIP on a 100×100 pixel window of a simulated 30-second data set (with the parameters specified in Section 4) over a range of KLIP parameters, we required 128GB of RAM and approximately 400 hours of computation time. The high memory requirement is due to the eigendecomposition, since the space-time covariance matrix can become extremely large when including a large number of lags and must be loaded in fully to the eigendecomposition. As mentioned in Section 3.3, there are possible solutions to this challenge to reduce computational costs in less memory intensive implementations, or even analytical gains in efficiency that exploit symmetries inherent in the covariance matrix (shown in Figure 1) or focus on only the strongest correlations depending on the temporal and spatial scales of interest, but those are beyond the scope of this paper.

It may also be possible to reduce the number of eigenvalues/modes computed, which will reduce computation time and possibly memory consumption as well, given that we now know that values beyond ∼50 KL modes aren't of much use in our tested scenario, but the ex-

act threshold will be dependent on the region of interest and number of lags used, among other factors. In future iterations, this code could also likely be improved by implementing this algorithm more optimally rather than in a high-level language, as the current implementation is in Python, and by using parallel processing.

6. CONCLUSION

Evolving atmospheric layers lead to time-varying speckles in the focal plane of an imaging system; for the high-contrast imaging regime, we have shown that spatio-temporal covariances in these speckles exist, and can be exploited for use in data processing to improve contrast. Our data processing tool has been implemented in Python, tested on a simple analytic test case to prove viability, and also tested on realistic simulations to understand the effectiveness of this technique. We have shown there is potential for a contrast gain (or equivalently, SNR improvement) of at least 10–20% from the original image, with significant potential for an even larger gain if self-subtraction is adequately addressed. Additionally, we have shown evidence that the space-time nature of our algorithm, in its current form, may provide a slight advantage over spatial-only KLIP in certain cases, with significant potential for stronger improvement under different conditions and with improvements to the algorithm implementation. Although the SNR gains for this new method aren't fully developed, this initial work on space-time KLIP opens the door for the use of space-time covariances in high-contrast imaging, especially in the short timescale regime of atmospheric speckle lifetimes.

Future work can use our data processing tool to further explore the dependence of the space-time covariances and the resulting contrast improvements on various parameters, such as the type of coronagraph, AO performance, strength of quasi-static speckles, and atmospheric conditions. It would be particularly interesting to determine how AO affects these covariances, since AO is important in a realistic scenario for exoplanet imaging and affects the resulting speckle lifetimes and structures.

Future implementations of this algorithm will also need to consider how to minimize self-subtraction of the companion object, and overcome the memory and computational demands in the eigendecomposition. Further optimization of the tunable parameters is also necessary to optimize algorithm performance and implement this as a refined tool for exoplanet imaging. It would also be interesting to apply this tool to on-sky data, such as that from MEC on SCExAO at Subaru (Walter et al. 2020, 2018; Jovanovic et al. 2015; Minowa et al. 2010),

to determine potential on-sky contrast gains from this technique. Although this current work focuses on the use of speckle space-time covariances in post-processing, these covariances could even be used in real-time predictive control (Guyon et al. 2018). Overall, the results in this work show that harnessing space-time covariances through "space-time KLIP" may be a promising technique to add to our toolkit for suppressing speckle noise in exoplanet imaging while retaining signal throughput.

REFERENCES

Adams, J., Wang, J., & Follette, K. I. 2021, in American Astronomical Society Meeting Abstracts, Vol. 53, American Astronomical Society Meeting Abstracts, 344.05

Aime, C., Borgnino, J., Martin, F., et al. 1986, Journal of the Optical Society of America A, 3, 1001, doi: 10.1364/JOSAA.3.001001

Aime, C., & Soummer, R. 2004, ApJL, 612, L85, doi: 10.1086/424381

Aime, C., & Soummer, R. 2004, The Astrophysical Journal Letters, 612, L85

Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, A&A, 558, A33, doi: 10.1051/0004-6361/201322068

Beuzit, J.-L., Vigan, A., Mouillet, D., et al. 2019, Astronomy & Astrophysics, 631, A155

Bharmal, N. A. 2015, in Journal of Physics: Conference Series, Vol. 595, IOP Publishing, 012003

Bloemhof, E. E., Dekany, R. G., Troy, M., & Oppenheimer, B. R. 2001, ApJL, 558, L71, doi: 10.1086/323494

Cagigal, M. P., & Canales, V. F. 2001, Optical Engineering, 40, 2690, doi: 10.1117/1.1417495

Calvin, B., Fitzgerald, M. P., van Kooten, M. A., et al. 2022, in Adaptive Optics Systems VIII, Vol. 12185, SPIE, 2542–2552

Canales, V. F., & Cagigal, M. P. 1999, ApOpt, 38, 766, doi: 10.1364/AO.38.000766

Collette, A. 2013, Python and HDF5 (O'Reilly)

Dainty, J. C., Hennings, D. R., & Odonnell, K. A. 1981, Journal of the Optical Society of America (1917-1983), 71, 490

Dask Development Team. 2016, Dask: Library for dynamic task scheduling. https://dask.org

Dodkins, R. 2018, PhD thesis, University of Oxford

Dodkins, R. H., Davis, K. K., Lewis, B., et al. 2020, Publications of the Astronomical Society of the Pacific, 132, 104503

Fassino, C., Pistone, G., & Rogantin, M. P. 2019, Mathematics, 7, doi: 10.3390/math7030263

Fitzgerald, M. P., & Graham, J. R. 2006, ApJ, 637, 541, doi: 10.1086/498339

Flasseur, O., Denis, L., Thiébaut, É., & Langlois, M. 2018, Astronomy & Astrophysics, 618, A138

Frazin, R. A. 2014, in Ground-based and Airborne Telescopes V, ed. L. M. Stepp, R. Gilmozzi, & H. J. Hall, Vol. 9145, International Society for Optics and Photonics (SPIE), 1368 – 1374, doi: 10.1117/12.2054356

Frazin, R. A., & Rodack, A. T. 2021, J. Opt. Soc. Am. A, 38, 1557, doi: 10.1364/JOSAA.426339

Gebhard, T. D., Bonse, M. J., Quanz, S. P., & Schölkopf, B. 2022, arXiv preprint arXiv:2204.03439

Gladysz, S., Yaitskova, N., & Christou, J. C. 2010, Journal of the Optical Society of America A, 27, A64, doi: 10.1364/JOSAA.27.000A64

Glindemann, A., Lane, R. G., & Dainty, J. C. 1993, Journal of Modern Optics, 40, 2381, doi: 10.1080/09500349314552401

Goebel, S. 2018, arXiv e-prints, arXiv:1807.04903. https://arxiv.org/abs/1807.04903

Goebel, S. B., Guyon, O., Hall, D. N. B., et al. 2018, PASP, 130, 104502, doi: 10.1088/1538-3873/aad8ed

Guyon, O. 2005, The Astrophysical Journal, 629, 592

Guyon, O., Jovanovic, N., Lozi, J., et al. 2017, in Techniques and Instrumentation for Detection of Exoplanets VIII, Vol. 10400, International Society for Optics and Photonics, 1040017

Guyon, O., & Males, J. 2017, arXiv preprint arXiv:1707.00570

Guyon, O., Mazin, B., Fitzgerald, M., et al. 2018, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 10703, Adaptive Optics Systems VI, ed. L. M. Close, L. Schreiber, & D. Schmidt, 107030Z, doi: 10.1117/12.2314331

Guyon, O., Pluzhnik, E., Kuchner, M. J., Collins, B., & Ridgway, S. 2006, The Astrophysical Journal Supplement Series, 167, 81

Guyon, O., Sevin, A., Gratadour, D., et al. 2018, in Adaptive Optics Systems VI, Vol. 10703, International Society for Optics and Photonics, 107031E

Guyon, O., Lozi, J., Vievard, S., et al. 2019, in American Astronomical Society Meeting Abstracts, Vol. 233, American Astronomical Society Meeting Abstracts #233, 104.03

Hafeez, R., Archinuk, F., Fabbro, S., Teimoorinia, H., & Véran, J.-P. 2021, arXiv e-prints, arXiv:2112.01437. https://arxiv.org/abs/2112.01437

Hickson, P. 2008, The University of British Columbia: Vancouver, BC, Canada, 1

Hunter, J. D. 2007, Computing in Science and Engineering, 9, 90, doi: 10.1109/MCSE.2007.55

Jensen-Clem, R., Bond, C. Z., Cetre, S., et al. 2019, in Techniques and Instrumentation for Detection of Exoplanets IX, Vol. 11117, International Society for Optics and Photonics, 111170W

Jones, E., Oliphant, T., Peterson, P., et al. 2001, SciPy: Open source scientific tools for Python. http://www.scipy.org/

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

Kluyver, T., Ragan-Kelley, B., Pérez, F., et al. 2016, in ELPUB, 87–90

Kolmogorov, A. N. 1941, in Akademiia Nauk SSSR Doklady, Vol. 32, 16

Krist, J. E. 2007, in Optical Modeling and Performance Predictions III, Vol. 6675, International Society for Optics and Photonics, 66750P

Lafreniere, D., Marois, C., Doyon, R., Nadeau, D., & Artigau, É. 2007, The Astrophysical Journal, 660, 770

Lake, K., Isgar, V., Baker, I., et al. 2020, in Sensors, Systems, and Next-Generation Satellites XXIV, Vol. 11530, International Society for Optics and Photonics, 115300H

Long, J. D., & Males, J. R. 2021, AJ, 161, 166, doi: 10.3847/1538-3881/abe12b

Macintosh, B. A., Graham, J. R., Palmer, D. W., et al. 2008, in Adaptive Optics Systems, Vol. 7015, International Society for Optics and Photonics, 701518

Males, J. R., Fitzgerald, M. P., Belikov, R., & Guyon, O. 2021, PASP, 133, 104504, doi: 10.1088/1538-3873/ac0f0c

Males, J. R., & Guyon, O. 2017, Ground-based adaptive optics coronagraphic performance under closed-loop predictive control. https://arxiv.org/abs/1712.07189

—. 2018, Journal of Astronomical Telescopes, Instruments, and Systems, 4, 019001

Marois, C., Correia, C., Galicher, R., et al. 2014, in Adaptive Optics Systems IV, Vol. 9148, International Society for Optics and Photonics, 91480U

Marois, C., Correia, C., Véran, J.-P., & Currie, T. 2013, Proceedings of the International Astronomical Union, 8, 48

Marois, C., Doyon, R., Nadeau, D., et al. 2005, PASP, 117, 745, doi: 10.1086/431347

Marois, C., Doyon, R., Nadeau, D., Racine, R., & Walker, G. A. H. 2003, in EAS Publications Series, Vol. 8, EAS Publications Series, ed. C. Aime & R. Soummer, 233–243, doi: 10.1051/eas:2003067

Marois, C., Lafreniere, D., Doyon, R., Macintosh, B., & Nadeau, D. 2006a, The Astrophysical Journal, 641, 556–564, doi: 10.1086/500401

—. 2006b, The Astrophysical Journal, 641, 556

Marois, C., Macintosh, B., Barman, T., et al. 2008, science, 322, 1348

Mawet, D., Serabyn, E., Liewer, K., et al. 2009, The Astrophysical Journal, 709, 53

Mazin, B. A., Bumble, B., Meeker, S. R., et al. 2012, Optics express, 20, 1503

Meeker, S. R., Mazin, B. A., Walter, A. B., et al. 2018, Publications of the Astronomical Society of the Pacific, 130, 065001

Minowa, Y., Hayano, Y., Oya, S., et al. 2010, Society of
Photo-Optical Instrumentation Engineers (SPIE)
Conference Series, Vol. 7736, Performance of Subaru
adaptive optics system AO188 (SPIE), 77363N,
doi: 10.1117/12.857818

Mullen, W. I., Macintosh, B., Chilcote, J., & Ruffio, J.-B.
2019, in American Astronomical Society Meeting
Abstracts# 233, Vol. 233, 140–38

Mullen, W. I., Macintosh, B., Chilcote, J., & Ruffio, J.-B.
2019, in American Astronomical Society Meeting
Abstracts, Vol. 233, American Astronomical Society
Meeting Abstracts #233, 140.38

Odonnell, K. A., Brames, B. J., & Dainty, J. C. 1982,
Optics Communications, 41, 79,
doi: 10.1016/0030-4018(82)90321-2

Oppenheimer, B. R., & Hinkley, S. 2009, Annual Review of
Astronomy and Astrophysics, 47

Perez, F., & Granger, B. E. 2007, Computing in Science
and Engineering, 9, 21, doi: 10.1109/MCSE.2007.53

Por, E. H., Haffert, S. Y., Radhakrishnan, V. M., et al.
2018, in Proc. SPIE, Vol. 10703, Adaptive Optics
Systems VI, doi: 10.1117/12.2314407

Price-Whelan, A. M., Sipőcz, B. M., Günther, H. M., et al.
2018, AJ, 156, 123, doi: 10.3847/1538-3881/aabc4f

Pueyo, L. 2016, The Astrophysical Journal, 824, 117

Racine, R., Walker, G. A. H., Nadeau, D., Doyon, R., &
Marois, C. 1999, PASP, 111, 587, doi: 10.1086/316367

Risvik, H. 2007, Principal component analysis (PCA) &
NIPALS algorithm, May

Rocklin, M. 2015, in Proceedings of the 14th Python in
Science Conference, ed. K. Huff & J. Bergstra, 130 – 136

Roddier, F., Gilli, J., & Lund, G. 1982, Journal of Optics,
13, 263

Roddier, F., Gilli, J. M., & Lund, G. 1982, Journal of
Optics, 13, 263, doi: 10.1088/0150-536X/13/5/002

Rudy, A. R., Srinath, S., Poyneer, L., et al. 2014, in Society
of Photo-Optical Instrumentation Engineers (SPIE)
Conference Series, Vol. 9148, Adaptive Optics Systems
IV, ed. E. Marchetti, L. M. Close, & J.-P. Vran, 91481Z,
doi: 10.1117/12.2057194

Samland, M., Bouwman, J., Hogg, D., et al. 2021,
Astronomy & Astrophysics, 646, A24

Savransky, D. 2015, The Astrophysical Journal, 800, 100

Scaddan, R. J., & Walker, J. G. 1978, ApOpt, 17, 3779,
doi: 10.1364/AO.17.003779

Schlaerth, J., Vayonakis, A., Day, P., et al. 2008, Journal of
Low Temperature Physics, 151, 684

Soummer, R., & Aime, C. 2004, in Advancements in
Adaptive Optics, ed. D. B. Calia, B. L. Ellerbroek, &
R. Ragazzoni, Vol. 5490, International Society for Optics
and Photonics (SPIE), 495 – 503, doi: 10.1117/12.551985

Soummer, R., & Aime, C. 2004, in Society of Photo-Optical
Instrumentation Engineers (SPIE) Conference Series,
Vol. 5490, Advancements in Adaptive Optics, ed.
D. Bonaccini Calia, B. L. Ellerbroek, & R. Ragazzoni,
495–503, doi: 10.1117/12.551985

Soummer, R., Ferrari, A., Aime, C., & Jolissaint, L. 2007,
The Astrophysical Journal, 669, 642

Soummer, R., Pueyo, L., & Larkin, J. 2012, The
Astrophysical Journal Letters, 755, L28

Steiger, S., Currie, T., Brandt, T. D., et al. 2021, The
Astronomical Journal, 162, 44

Taylor, G. I. 1938, Proceedings of the Royal Society of
London. Series A-Mathematical and Physical Sciences,
164, 476

Thompson, W., & Marois, C. 2021, arXiv preprint
arXiv:2103.09252

van der Walt, S., Colbert, S. C., & Varoquaux, G. 2011,
Computing in Science and Engineering, 13, 22,
doi: 10.1109/MCSE.2011.37

van Kooten, M. A., Jensen-Clem, R., Bond, C., et al. 2021,
in Techniques and Instrumentation for Detection of
Exoplanets X, Vol. 11823, International Society for
Optics and Photonics, 118231F

Vernin, J., Caccia, J. L., Weigelt, G., & Mueller, M. 1991,
A&A, 243, 553

Vievard, S., Bos, S. P., Cassaing, F., et al. 2020, in
Adaptive Optics Systems VII, Vol. 11448, International
Society for Optics and Photonics, 114486D

Vigan, A., Moutou, C., Langlois, M., et al. 2010, Monthly
Notices of the Royal Astronomical Society, 407, 71

Walter, A., Mazin, B. B., Bockstiegel, C., et al. 2018, in
Ground-based and Airborne Instrumentation for
Astronomy VII, Vol. 10702, International Society for
Optics and Photonics, 107020V

Walter, A. B., Bockstiegel, C., Brandt, T. D., & Mazin,
B. A. 2019, Publications of the Astronomical Society of
the Pacific, 131, 114506

Walter, A. B., Fruitwala, N., Steiger, S., et al. 2020,
Publications of the Astronomical Society of the Pacific,
132, 125005

Wang, J. J., Ruffio, J.-B., De Rosa, R. J., et al. 2015,
pyKLIP: PSF Subtraction for Exoplanets and Disks.
http://ascl.net/1506.001

Wizinowich, P., Chin, J., Correia, C., et al. 2020, in
  Adaptive Optics Systems VII, ed. L. Schreiber,
  D. Schmidt, & E. Vernet, Vol. 11448, International
  Society for Optics and Photonics (SPIE), 49 – 63,
  doi: 10.1117/12.2560017

Wu, C., Gu, G., Chen, Q., et al. 2021, in Society of
  Photo-Optical Instrumentation Engineers (SPIE)
  Conference Series, Vol. 11763, Society of Photo-Optical
  Instrumentation Engineers (SPIE) Conference Series,
  117631J, doi: 10.1117/12.2586285

APPENDIX

A. NOTATION GLOSSARY – SECTIONS 2.1 & 2.3

Symbol	Definition
$I_\psi(k)$	Stellar PSF, as in Soummer et al. (2012)
k	Pixel index
$T(k), \boldsymbol{t}$	Target image as in Soummer et al. (2012), and as in this work unrolled to 1-d and represented as a vector
$A(k), \boldsymbol{a}$	Faint astronomical signal (as above)
ϵ	True/false binary parameter
$\hat{I}_\psi(k), \hat{\boldsymbol{\psi}}$	Approximated PSF as in Soummer et al. (2012) and as a vector in this work, respectively
R	Matrix of reference images before mean subtraction
$\boldsymbol{r}$	Individual reference image
X	Mean image from reference set
M	Mean subtracted reference images
n_i	Number of images in reference set
n_p	Pixel count $n_x \times n_y$
n_x	Dimension 1 size
n_y	Dimension 2 size
n_m	Number of modes / eigenvectors chosen
i	Used as an arbitrary index
j	Used as an arbitrary index
C	Covariance matrix
$\boldsymbol{\lambda}, \lambda$	Vector of eigenvalues, eigenvalue
V	Matrix of eigenvectors/eigenimages
$\boldsymbol{v}$	Eigenvector
$\boldsymbol{q}$	Vector of coefficients
$S, \boldsymbol{s}$	Mean subtracted image sequences (in matrix and vector form)
$\hat{\boldsymbol{s}}$	Reconstructed image sequence
n_s	Number of images in sequence
L	Largest number of timesteps/lags in use as measured from the central image
n_l	Total number of timesteps/lags used, equal to n_s
$\boldsymbol{r}_{k,avg}$	Averaged residual from stKLIP

B. NOTATION GLOSSARY – SECTION 2.2

Symbol	Definition
I	Intensity
$\boldsymbol{x}$	Location in image plane
$\boldsymbol{u}$	Location in pupil plane
t	Time
τ	Time step
Ψ_{pup}	Pupil amplitude
Ψ_{foc}	Focal amplitude
$\psi(\boldsymbol{u}, t)$	Pupil phase
$P(\boldsymbol{u})$	Pupil function
$C(\boldsymbol{x})$	Spatially coherent wavefront
$S_\phi(\boldsymbol{x}, t)$	Phase aberrations
ξ	Displacement in pupil
B_ϕ	phase covariance function
v_{wind}	Wind velocity