



Universiteit
Leiden
The Netherlands

Improved companion mass limits for Sirius A with thermal infrared coronagraphy using a vector-apodizing phase plate and time-domain starlight-subtraction techniques

Long, J.D.; Males, J.R.; Haffert, S.Y.; Pearce, L.; Marley, M.S.; Morzinski, K.M.; ... ; Tolls, V.

Citation

Long, J. D., Males, J. R., Haffert, S. Y., Pearce, L., Marley, M. S., Morzinski, K. M., ... Tolls, V. (2023). Improved companion mass limits for Sirius A with thermal infrared coronagraphy using a vector-apodizing phase plate and time-domain starlight-subtraction techniques. *The Astronomical Journal*, 165(5). doi:10.3847/1538-3881/acbd4b

Version: Publisher's Version

License: [Creative Commons CC BY 4.0 license](https://creativecommons.org/licenses/by/4.0/)

Downloaded from: <https://hdl.handle.net/1887/3716902>

Note: To cite this publication please use the final published version (if applicable).



Improved Companion Mass Limits for Sirius A with Thermal Infrared Coronagraphy Using a Vector-apodizing Phase Plate and Time-domain Starlight-subtraction Techniques

Joseph D. Long¹ , Jared R. Males¹ , Sebastiaan Y. Haffert¹ , Logan Pearce¹ , Mark S. Marley² , Katie M. Morzinski¹ , Laird M. Close¹ , Gilles P. P. L. Otten^{3,4} , Frans Snik³ , Matthew A. Kenworthy³ , Christoph U. Keller^{3,5} , Philip Hinz^{6,7} , John D. Monnier⁸ , Alycia Weinberger⁹ , and Volker Toll¹⁰

¹ Steward Observatory, The University of Arizona, Tucson, AZ, USA; josephlong@arizona.edu

² Lunar & Planetary Laboratory, The University of Arizona, Tucson, AZ, USA

³ Leiden Observatory, Leiden University, Leiden, The Netherlands

⁴ Academia Sinica, Institute of Astronomy and Astrophysics, Taipei, Taiwan

⁵ Lowell Observatory, Flagstaff, AZ, USA

⁶ Steward Observatory, University of Arizona, Tucson, AZ, USA

⁷ University of California Observatories, Santa Cruz, CA, USA

⁸ University of Michigan, Department of Astronomy, Ann Arbor, MI, USA

⁹ Carnegie Institution of Washington, Earth and Planets Laboratory, Washington, D.C., USA

¹⁰ Harvard-Smithsonian Center for Astrophysics, Cambridge, MA, USA

Received 2022 October 12; revised 2023 February 10; accepted 2023 February 18; published 2023 April 26

Abstract

We use observations with the infrared-optimized Magellan Adaptive Optics (MagAO) system and Clio camera in $3.9\ \mu\text{m}$ light to place stringent mass constraints on possible undetected companions to Sirius A. We suppress the light from Sirius A by imaging it through a grating vector-apodizing phase plate coronagraph with a 180° dark region (gvAPP-180). To remove residual starlight in postprocessing, we apply a time-domain principal-components-analysis-based algorithm we call PCA-Temporal, which uses eigen time series rather than eigenimages to subtract starlight. By casting the problem in terms of eigen time series, we reduce the computational cost of postprocessing the data, enabling the use of the fully sampled data set for improved contrast at small separations. We also discuss the impact of retaining fine temporal sampling of the data on final contrast limits. We achieve postprocessed contrast limits of 1.5×10^{-6} – 9.8×10^{-6} outside of $0''.75$, which correspond to planet masses of 2.6 – $8.0\ M_J$. These are combined with values from the recent literature of high-contrast imaging observations of Sirius to synthesize an overall completeness fraction as a function of mass and separation. After synthesizing these recent studies and our results, the final completeness analysis rules out 99% of $\geq 9\ M_J$ planets from 2.5 to $7\ \text{au}$.

Unified Astronomy Thesaurus concepts: Coronagraphic imaging (313); High contrast techniques (2369); Direct imaging (387); Exoplanet detection methods (489); Exoplanet astronomy (486)

1. Introduction

As the brightest star in the sky, it is no surprise that Sirius has been observed since antiquity. Measurements of the proper motion of Sirius in the 1800s led to the prediction that a second body was affecting the motion of the Sirius system (Bessel 1844). Eighteen years later, this was confirmed observationally with the detection of Sirius B (Bond 1862). When a 6.25 yr periodic perturbation was noted and analyzed by Benest & Duvent (1995), there was a renewed search for a low-mass companion to Sirius A that could match their prediction of a mass $M \lesssim 50\ M_J$ object orbiting at $\sim 7.9\ \text{au}$. Studies have been undertaken with the Hubble Space Telescope (HST; Schroeder et al. 2000) and from the ground with adaptive optics systems (Thalmann et al. 2011; Vigan et al. 2015; among others), most recently by the VISIR team in $10\ \mu\text{m}$ light (Pathak et al. 2021). These observations have ruled out the existence of low-mass companions in the mass–separation regime predicted by Benest and Duvent (1995), but

Sirius remains an attractive target for exoplanet direct-imaging searches for other reasons.

The proximity of Sirius ($d = 2.67\ \text{pc}$; Bailer-Jones et al. 2021) means that smaller projected orbital radii are accessible to direct-imaging instruments than for more distant stars. Furthermore, extrasolar giant planets are statistically not uncommon around “retired A stars” (Johnson et al. 2007; Ghezzi et al. 2018), and we expect this to also be true for A-type progenitors like Sirius. Probing smaller separations and longer wavelength ranges increases the chance of finding a low-mass companion previous studies may have missed. The high $T_{\text{eff}} = 9910 \pm 130\ \text{K}$ (Liebert et al. 2005) of Sirius A means the contribution of instellation to the effective temperatures of planets at small separations will be significant, boosting their near-infrared luminosity and thus our ability to detect them. We summarize the properties of Sirius A and the relevant sources in Table 1.

2. Observations

We observed Sirius with the Magellan Clay telescope at Las Campanas Observatory on the night of 2015 November 29 UT using the Magellan Adaptive Optics (MagAO) system (Close et al. 2013) and the Clio infrared (IR) imaging instrument (Morzinski et al. 2015).



Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](https://creativecommons.org/licenses/by/4.0/). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

Table 1
Properties of Sirius A

Property	Symbol	Value	Source
Age	t	242 ± 15 Myr	Bond et al. (2017)
Luminosity	L_{Sirius}	$25.4 \pm 1.3 L_{\odot}$	Liebert et al. (2005)
Radius	R_{Sirius}	$1.711 \pm 0.013 R_{\odot}$	Kervella et al. (2003)
Effective temperature	T_{Sirius}	9910 ± 130 K	from L_{Sirius} and R_{Sirius}
Distance from Earth	d_{Sirius}	2.67 ± 0.001 pc	Bailer-Jones et al. (2021)
Magnitude in Clio [3.95] filter	$m_{[3.95]}$	-1.39	Using Bohlin et al. (2019)’s spectrum

Table 2
Observations of Sirius A and o Gruis: conditions

Target	Start UT	End UT	Sky Rotation	Seeing ($''$)	PWV (mm)
o Gruis	2015-11-29 00:38:01	2015-11-29 00:45:10	...	0.58	2.0
Sirius A	2015-11-29 05:58:17	2015-11-29 08:55:32	105°	0.39–0.53	1.9–2.2

Note . The PWV measurements are approximate.

Table 3
Observations of Sirius A and o Gruis: camera and MagAO configurations

Target	AO Modes	Loop Freq. (Hz)	Filter	Coronagraph	Purpose	Exposure (s)	Frames
o Gruis	300	989.6	[3.95]	gvAPP-180	Calibration	5	30
					Sky	5	30
Sirius A	300	989.6	[3.95]	gvAPP-180	Science	0.5	10,889
					Sky	0.5	1220
					PSF cal.*	0.28	8
					PSF cal. sky	0.28	8

Note. Frames with stray light artifacts impacting the dark-hole regions were excluded from the original 12,800 frames to obtain 10,889 science frames. *Not used due to saturation.

The grating vector-apodizing phase plate 180° (gvAPP-180) coronagraph design we employed is a further refinement of the original apodizing phase plate of Kenworthy et al. (2007). By using a liquid crystal material imprinted with the apodizing phase pattern, the phenomenon of vector (or geometric) phase provides a larger range of wavelengths over which the apodization is effective (Snik et al. 2012). The grating separates oppositely circularly polarized light on the detector, producing two copies of the point-spread function (PSF) with complementary dark holes—and a nearly 360° accessible dark-hole region when the two are recombined. The APP family of coronagraphs are pupil-plane optics, meaning alignment and pointing accuracy tolerances are more forgiving than small inner working angle (IWA) focal plane coronagraphs (Doelman et al. 2021). The gvAPP-180 in MagAO/Clio was first characterized by Otten et al. (2017) and recently used by Sutcliffe et al. (2021) for high-contrast imaging of a brown dwarf companion.

Seeing conditions over the course of the observation ranged from $0''.39$ to $0''.53$ with a median seeing of $0''.44$ measured by a differential image motion monitor (DIMM). Direct precipitable water vapor (PWV) measurements at the time of these observations are not available, but measurements of ground-level relative humidity indicate an estimate of 2.15 mm PWV at the beginning of these observations, dropping to 1.92 mm by the end of them. The Sirius frames used for this study were taken from UT 05:58:17.11 to UT 08:55:32.35, encompassing

105° of sky rotation. Observations of o Gruis for PSF calibration were taken shortly beforehand, with a seeing of $0''.58$ and an approximate PWV of 2.0 mm. The conditions for both targets are summarized in Table 2.

We observed Sirius using the gvAPP-180 coronagraph with the [3.95] narrowband filter ($\lambda = 3.95 \mu\text{m}$, $\Delta\lambda = 0.091 \mu\text{m}$) to minimize radial smearing of the vAPP PSFs. We integrated for 500 ms each frame, obtaining a total of 12,800 frames. We identified a stray light artifact impacting 1911 frames, leaving 10,889 science frames to analyze encompassing ≈ 90 minutes of integration time.

The MagAO system was correcting 300 modes at 1 kHz, which resulted in residual wave front error of approximately 100–150 nm rms. At $3.95 \mu\text{m}$, this corresponds to a Strehl ratio of $S \approx 96\%$, before including noncommon-path aberrations. The field rotator was adjusted such that no stray light impacted the dark hole, and left in that position to enable the angular differential imaging (ADI) technique (Marois et al. 2006).

The Clio camera uses a legacy prototype Hawaii-I HgCdTe detector, sensitive out to $5 \mu\text{m}$. Due to the extreme brightness of Sirius, the typical “nodding” procedure whereby the star is placed on alternating halves of the detector for sky estimation was not used. Instead, the star was periodically offset to a location completely off chip, resulting in 1220 total sky frames at intervals of 24 minutes.

The camera and MagAO configurations are summarized in Table 3.




Figure 1. Single example frame after background subtraction, scaled to show structure. Bad pixels are cyan, the science PSFs and the cold stop diffraction spikes are outlined in blue, negative artifacts from amplifier crosstalk at a predictable separation from the star are shown in green, and other stray light artifacts that can move over the course of the observation are highlighted in orange. (Adapted from Long et al. 2018.)

3. Data Reduction

The extreme brightness of Sirius presented several hurdles to data reduction. Formerly unnoticed and negligible spurious reflections became nonnegligible, as shown in Figure 1. Not only did the PSF core saturate, but the first two diffraction rings of the PSF did as well. Short integration times translated into huge numbers of frames to incorporate in the reduction. In light of these challenges, we go into some detail about how we mitigated them.

3.1. Sky Background and Stray Light

Thermal IR is challenging to observe from the ground due to the rapidly varying background illumination from the sky and noncryogenic optics. Unlike stellar speckle noise, background counts are best modeled in the original pixel frame, before image alignment. We used the sky frames we recorded to construct a principal-components basis. To model and subtract the background signal in a given science frame, we first identified pixels within a certain radius of the twin gvAPP-180 PSFs or a number of Clio-specific artifacts and excluded them (Figure 1). To mitigate impacts on sensitivity, only the small subset of pixels with negligible light from Sirius A were used to estimate the background in the reconstruction process. We then performed a least-squares fit of the basis vectors to the pixels that remained, and used those coefficients to construct a background frame from the full unmasked principal component basis vectors (Long et al. 2018; and independently in Hunziker et al. 2018).

We held back 25% of the sky frames in order to cross-validate the procedure, finding that a background model with the first six principal components reproduced background counts with an rms error of 13 counts ($\approx 0.1\%$ of the average background level) when the fit excluded a mask representative of those for science frames. Although the final starlight subtraction is equally capable of subtracting the sky background, we found separating these steps into two stages improved our ability to create aligned cutouts and simplified later stages of the pipeline. The error in the background reconstruction of 13 counts amounts to 7×10^{-8} in contrast units for a star of Sirius’ brightness, which we judge to be negligible.




Figure 2. Radial profiles of the “top” (upper) and “bottom” (lower) vAPP PSFs (see Figure 3) taken from median-combined data of Sirius and o Gruis, normalized to 1.0 at $\rho = 0''$. The o Gruis data allowed us to estimate the brightness of Sirius by fitting the wings of the PSF, and were used as an unsaturated PSF template when injecting fake planets.

3.2. PSF Saturation and Frame-to-frame Variation

In half-second exposures with the coronagraph splitting the incoming beam in two, Sirius still saturated the PSF core completely, along with some of the diffraction rings.

In order to model any companion accurately, we needed to infer the missing flux by fitting a model to the unsaturated portions of the Sirius PSF in each of the two gvAPP-180 halves. Shorter-integration Sirius data taken for calibration were also saturated, so we used a radial profile from stacked o Gruis observations with the same filter and coronagraph configuration to serve as our unsaturated reference. The o Gruis observations were taken the same night (2015 November 29 UT) under similar conditions in terms of PWV and with the MagAO system correcting 300 modes. The reported rms residual wave front error was approximately 100 nm, for a Strehl ratio $S \approx 97\%$ before noncommon-path errors. In an ideal system, there would be no flux difference between the two gvAPP-180 halves, but Otten et al. (2017) noted that one half of the split PSF was brighter than the other by several percent, and we also observed this in our data. As the PSFs are split by their left- or right-circular polarization, this discrepancy indicates a small amount of circular polarization is introduced by the instrument itself. As we observed the effect on multiple targets, we concluded the origin is instrumental, rather than any astrophysical polarization signal.

We divided each frame into two cutouts centered on the complementary PSFs, which we aligned to subpixel precision. We established alignment to a known center of rotation by simulating a gvAPP-180 PSF without wave front errors, using said PSF to align an unsaturated calibration target PSF (in this case, the o Gruis template), and using this second, more realistic, PSF to align the Sirius frame cutouts. The result was two data cubes with dark-hole regions on opposite sides of the core, aligned to subpixel precision. We computed a radial profile of the template PSF, normalized to unit flux, and each frame’s science PSFs. At each frame, we fit the scale factor that matched the radial profiles in the unsaturated region best, giving us the total Sirius flux. A comparison of the profiles in the top and bottom PSFs is given in Figure 2. The effect of the “lost” flux is illustrated by Figure 3.




Figure 3. The first row shows the “top” vAPP median PSF, with saturated Sirius data on the left and o Gruis data on the right. The second row shows the “bottom” PSFs. They have been scaled to relative units such that the peak of the unsaturated data is 1.0, and the saturated data are scaled to match the radial profiles in the unsaturated regions (see Figure 2).

Once in possession of two aligned data cubes, the amplitude difference between complementary gvAPP-180 PSFs was divided out using the per-frame scale factors.

3.3. Starlight Subtraction

The brightness of Sirius required short integration times, which in turn led to a large number of frames containing distinct realizations of the system PSF. The KLIP algorithm of Soummer et al. (2012) leverages the difference between the number of realizations n_{obs} and the number of random variables (i.e., pixels) p , under the assumption that $n_{\text{obs}} \ll p$, to reduce the size of the matrix to be eigen decomposed to obtain the Karhunen–L  ve basis. With $n_{\text{obs}} \sim 10,000$ and $p \sim 7000$ for our analysis, this assumption no longer held. In fact, when $n_{\text{obs}} \simeq p$, the cost of forming the covariance matrix outweighed any time savings compared to a direct computation of the singular value decomposition (SVD) of the data matrix (Long & Males 2021).

A common technique, adopted by Vigan et al. (2015) and many others, is to bin the resulting data temporally to reduce the dimensions of the problem. However, Males et al. (2021) found that we should expect speckle lifetimes of roughly 20–100 ms on bright stars, well under even our 500 ms Sirius exposures. As such, we expected that capturing the information from every frame within an interval rather than temporally binning will result in improved starlight-subtraction performance. This is, in fact, what we found for locations near the host star, as discussed in Section 3.6.

Initial attempts to apply the SVD downdate algorithm described in Long & Males (2021) to accelerate the reduction of the Sirius vAPP data proved insufficient. KLIP itself is vulnerable to self-subtraction in ADI data, as a planet PSF rotating through the field has enough overlap with itself from frame to frame that the resulting starlight-subtraction process is

able to remove or severely attenuate it. This is usually mitigated by use of a PSF reference star or an angular exclusion criterion, removing frames temporally adjacent to the one under analysis. However, applying an angular exclusion criterion undid some of the speed gains of the SVD downdate algorithm, as the inner SVD had to grow by a row and a column for each frame to be removed. In such a high frame rate data set, when probing locations of interest near the IWA large numbers of frames had to be excluded.

For this reason, we developed the PCA-Temporal (PCAT) algorithm detailed in Section 3.5. By operating on pixel time series rather than frames, it is straightforward to mask a planet’s path from the reference sample. This in turn means a single decomposition is required for all pixels (and therefore frames) to starlight subtract.

3.4. Fake Planet Injection and S/N Measurement

After normalizing the PSF cubes with the profile fitting procedure in Section 3.2, we conducted “fake planet” signal injection–recovery tests with the o Gruis template PSF as a realistic proxy for the unsaturated Sirius PSF. The paired PSFs of the gvAPP-180 required injecting the appropriate companion PSF template into each half of the data, as companions could rotate out of one dark hole and into another over the course of an observation.

To perform the injection–recovery test, a final focal plane location (ρ, θ) was selected (where ρ is the separation from the center of the stellar PSF and θ is the position angle (PA) in degrees E of N). For each frame and each half, the template PSF was scaled by amplitude A and translated to $(\rho, \theta + \phi_i)$ in the focal plane (where ϕ is the negative of the derotation angle that places north up and east left for frame i) before adding it to the original image. To mitigate the influence of the injected planet signals on each other and the potential biasing of the estimated contrast, we only injected one companion at a time to measure the contrast limit in a location.

The starlight in each frame was subtracted following the algorithm in Section 3.5, and the residuals were derotated to place north up and east left, and combined as the pixel-wise median of the stack of frames.

The signal and noise samples were computed from the final image by simple aperture photometry in λ/D diameter apertures, excluding apertures on either side of the measurement location (i.e., known injected planet). The signal-to-noise ratio (S/N) was then computed following the corrected equation for small-sample statistics from Mawet et al. (2014).

3.5. The PCAT Algorithm

Each focal plane pixel can be treated as a time series that can be represented in a basis of eigen time series. This technique is employed by the Temporal Reference Analysis of Planets (TRAP; Samland et al. 2021), which also includes a constant term and a model light curve in the design matrix to be fit. However, the unique constraints of the paired vAPP PSF required development of a bespoke pipeline. By excluding pixels at the same separation as our injected companion from the basis determination, we avoid self-subtraction from overlapping companion flux. In essence, we have transposed the problem defined by Soummer et al. (2012) from eigenimages that describe the column space of a data matrix with one column per frame, to eigen time series that live in the row space




Figure 4. The two complementary decompositions of a hypothetical data matrix with one column for each frame. The matrix can be decomposed into left singular vectors, shown here as $v_1, v_2, \dots, v_n$, or into right singular vectors as $u_1, \dots, u_n$. An interpretation of the left singular vectors, with the same number of entries as a single vector of pixels from an image, as eigenimages of the pixel-to-pixel covariance, is shown in the lower left. An analogous interpretation for the right singular vectors is as eigen time series from which each pixel’s time series can be composed, as shown on the right.

as shown in Figure 4. In this way, we may use a single basis of eigen time series to starlight subtract every frame, eliminating even the (reduced) per-frame cost of the SVD downgrade. Furthermore, by excluding an annulus of pixels from our reference set, we may reuse a single decomposition for multiple companion PA values at a given separation, as the reference sets are identical. We found that a more selective mask covering only the notional companion’s track, while better for including information on diametrically opposite speckles, also led to overfitting.

To begin, a mask is optionally applied to select pixels from an N -frames sequence of images, and each frame’s p selected pixels are unwrapped into vector $\mathbf{x}_i$, combined as the columns of the data matrix $\mathbf{X}''_{p \times N}$. Let $\tilde{\mathbf{x}}$ be a vector whose entries correspond to the median value of each row in matrix $\mathbf{X}''$. The median-subtracted data are then $\mathbf{X}'_{p \times N} = \mathbf{X}''_{p \times N} - \tilde{\mathbf{x}}_{p \times 1} \mathbf{1}_{1 \times N}$.

To reduce the effect of brighter pixels on the decomposition, we divide each row j of $\mathbf{X}'$ by a whitening factor given by the standard deviation σ_j of that row. Let matrix $\mathbf{W}$ be a diagonal matrix $\mathbf{W} = \text{diag}(\sigma_0^{-1}, \dots, \sigma_p^{-1})$. The preprocessed observation matrix $\mathbf{X}$ is then given by

$$\mathbf{X}_{p \times N} = \mathbf{W}(\mathbf{X}''_{p \times N} - \tilde{\mathbf{x}}_{p \times 1} \mathbf{1}_{1 \times N}).$$

3.5.1. Special Considerations for the gvAPP-180 Data

When applying starlight-subtraction algorithms to the gvAPP-180 data, there are many possible ways to combine (or not combine) the data from the two halves into a final image. After experimentation, we found that the best performance with the fewest artifacts came from forming the image vector from the dark-hole pixels from one half, unwrapped into a vector, and concatenated with the vector of dark-hole pixels from the other half such that the paired PSFs from a single frame map to a single observation vector $\mathbf{x}_i$. (This is equivalent to stitching the images beforehand, though this approach would also allow the use of the bright half of the PSF pixels in the reduction and using different masks for the estimation and the final combination steps.)

3.5.2. Reference Time-series Selection

Recall that to implement KLIP with ADI, one would exclude a observation columns from $\mathbf{X}$ to make $\mathbf{R}_{p \times (N-a)}$, a matrix of reference observations assumed not to contain the signal of interest. The analog in the time domain is excluding time series that would contain a planet “transit” signal. The rate at which a planet PSF would pass over a detector pixel in ADI

observations depends on its angular separation from the center of rotation (i.e., the host star). This means a planet-like signal “transiting” over any pixel at the same radius as our location of interest may cause it to be subtracted, if they are close enough in time. Therefore, we exclude b of the pixel time series in a ring of width $\Delta r = 2 \lambda/D$, centered on the separation of the companion (whether it is one we injected or one we are trying to detect). The reference data matrix is then of size $\mathbf{R}_{(p-b) \times N}$.

3.5.3. Eigen Time-series Computation

Having selected the reference pixel time-series vectors to form $\mathbf{R}$, the algorithm proceeds with an SVD to obtain the left and right singular vectors as well as their associated singular values. We use SVD to mean the “economy” SVD that omits the decomposition of the null space of the matrix. The SVD allows us to decompose $\mathbf{R}$ into

$$\mathbf{R} = \mathbf{U}_{(p-b) \times K} \mathbf{\Sigma}_{K \times K} (\mathbf{V}_{N \times K})^T,$$

where $K = \min((p-b), N)$. To prevent overfitting and improve starlight subtraction, we retain the $k < K$ column vectors in the decomposition corresponding to the greatest singular values. (If $k \ll K$, this may allow the use of faster algorithms that find the partial SVD.) The decomposition is now only approximate, and given by

$$\mathbf{R} \approx \tilde{\mathbf{R}} = \tilde{\mathbf{U}}_{(p-b) \times k} \tilde{\mathbf{\Sigma}}_{k \times k} (\tilde{\mathbf{V}}_{N \times k})^T.$$

The question of the optimal choice of k is addressed in Section 3.6. The matrix $\tilde{\mathbf{V}}$ defined above gives us the eigen time series with which we will model each pixel’s evolution.

3.5.4. Subtraction of the Estimated Starlight

Up until now, we have referred to observation column vectors $\mathbf{x}_i \in \mathbb{R}^p$. Now we operate on the rows of $\mathbf{X}$, which we will call $\mathbf{y}_j \in \mathbb{R}^N$. For every pixel time series we use from the original data cube, we proceed to project it into the basis of eigen time series $\tilde{\mathbf{V}}$ and reconstruct it in terms of the first k eigen time series

$$\tilde{\mathbf{y}}_j \approx \tilde{\mathbf{V}}((\tilde{\mathbf{V}}^T)_{k \times N} \mathbf{y}_j).$$

The residual noise plus companion signal (if any) is a vector $\mathbf{s}_j \in \mathbb{R}^N$ given by $\mathbf{s}_j = \mathbf{y}_j - \tilde{\mathbf{y}}_j$.

By repeating this procedure for all p rows of $\mathbf{X}$, we obtain the starlight-subtracted signal $\mathbf{S}$ whose rows are $\mathbf{s}_1, \mathbf{s}_2, \dots, \mathbf{s}_p$

$$\mathbf{S}_{p \times N} = (\mathbf{X}^T - \tilde{\mathbf{V}}(\tilde{\mathbf{V}}^T \mathbf{X}^T))^T.$$

To reverse the whitening transformation, we use $\mathbf{W}^{-1}$ and obtain the final starlight-subtracted data $\mathbf{D}$

$$\mathbf{D} = \mathbf{W}_{p \times p}^{-1} \mathbf{S}_{p \times N}.$$

The entries of the columns of $\mathbf{D}$ may then be mapped back into pixel locations in the original data cube to construct a cube of residuals and further postprocessed (e.g., by derotating and stacking images for ADI).

3.6. Optimization of the PCAT Hyperparameters

Due to the unique structure of the gvAPP-180 data, the optimal number of modes to subtract k varied with both separation and PA. Additionally, when calibrating the $S/N = 5$ contrast floor, we observed that injecting and recovering a signal of amplitude A producing an S/N of $S \gg 5$ would




Figure 5. Median $S/N = 5$ contrast vs. separation shown at 10, 20, 30, 40, 50, 100, and 200 iterations of hyperparameter exploration by the Bayesian optimization algorithm. The Bayesian optimization procedure rapidly converges to near-optimal parameters.

predict a larger value $A_{S/N=5} = (5A)/S$ for the minimum contrast detectable at $S/N = 5$ than an injection–recovery test closer to the true $S/N = 5$ amplitude. Or, to put it another way, calibrating an accurate minimum contrast value depends on injecting and recovering a signal close to the minimum amplitude that is detectable at a good S/N , making this an iterative process. Calibrating the $S/N = 5$ contrast floor therefore required tuning of two hyperparameters for every location probed: the injected signal amplitude and the number of PCAT modes. The width of the annular mask excluding pixels from the reference sample was fixed at $b = 16$ pixels ($\approx 2\lambda/D$).

The hyperparameter-tuning problem is an area of ongoing research in the machine-learning literature (see, e.g., Yang & Shami 2020 for a review), with multiple techniques to explore the parameter space efficiently at our disposal. Since each parameter evaluation is comparatively expensive, we chose Bayesian optimization (Snoek et al. 2012), descended from the “kriging” technique (also known as Gaussian process regression) initially developed to predict the profitability of mining gold deposits from a small number of bore-hole samples (Krige 1951).

The optimization process attempts to learn a relationship between the number of PCAT modes, k , injected signal amplitude, A , and the resulting recovered S/N S at some location. The objective function we chose to maximize is

$$f(A, S) = -S \log_{10} A.$$

This depends only on the recoverability of the injected signal, and does not attempt to optimize the S/N from reducing the no-injection data set. A separate forthcoming publication (J. D. Long et al. 2023, in preparation) will discuss the particulars of our approach to hyperparameter optimization and the distributed computing infrastructure developed for it. The optimization process was run for 200 iterations of selecting parameters, injecting a signal, reducing the data with said parameters, and measuring the recovered signal. Figure 5 shows the rapid convergence of the optimizer, with only marginal contrast gains after 50 iterations.

In recognition of the difference in noise statistics between points in the speckle-dominated region close to the star, and background-limited points further out, we did compute contrast limits with varying amounts of coadding (shown in Figure 6). Retaining every frame in full temporal sampling produced better contrast limits at small separations, while performance at




Figure 6. Median $S/N = 5$ contrast vs. separation for three different levels of coadding. Within $4 \lambda/D$, the median contrast degrades with temporal downsampling by coadding. This relationship turns over, and by $5 \lambda/D$ the 10:1 downsampled data cube actually provides better contrast.




Figure 7. Behavior of optimized median $S/N = 5$ contrast at each radius value with increasing number of optimization iterations, comparing PCAT to our implementation of KLIP. Both algorithms were optimized using the same Bayesian optimization procedure, with the median $S/N = 5$ contrast shown after 10 iterations as a dotted line, after 100 as a dashed line, and after 200 as a solid line.

larger separations benefited somewhat from coadding temporally adjacent frames.

We validated the performance of PCAT by comparison with our implementation of KLIP, with the same optimization scheme. We optimized the 10:1 coadded reduction with KLIP following the same procedure as PCAT, and show the results in Figure 7.

To compute the final contrast surface, we collected all combinations of parameters evaluated for a particular focal plane location where the injected signal was recovered with at least an $S/N \geq 8$. This cutoff was chosen arbitrarily to limit the analysis to points with relatively trustworthy estimates of the $S/N = 5$ contrast level $A_{S/N=5} = 5A/8$. The parameter combination with the smallest $A_{S/N=5}$ was saved, and those values define the surface in Figure 8. The k modes, injected amplitude A , and temporal downsampling factor n were allowed to vary, though we only evaluated fixed steps of $n \in \{1, 10, 50\}$ due to computing resource constraints. The spatial distributions of these parameters are shown in Appendix.

4. Detection Limits

Using the procedure described above, we calibrated the minimum brightness of an $S/N = 5$ signal at a set of focal plane locations arranged in rings around the center of the host-




Figure 8. Contrast map of the faintest signal detectable at $S/N = 5$, expressed as a ratio to the host-star flux.




Figure 9. The pattern of saturated pixels in each complementary half of the gvAPP-180 PSF, combined with the field rotation, leads to a varying number of frames from the original science data contributing to the value of each final pixel. This picture shows the result of rotating the data masks (set to 1.0 where data are kept) and summing along the time axis to visualize this effect. The overlaid points represent locations at which we injected and recovered a planet signal to calibrate our contrast limits.

star PSF (Figure 9). To obtain the detection map in Figure 10, we took those parameter combinations that produced the best $A_{S/N=5}$, and applied them to reduce the pristine data (with no injected signals). Our gvAPP-180 coronagraph allows for remarkably deep contrast at close separations, as shown in Figure 11. When scaled to the wavelength and diameter of the telescope used, the smaller IWA of the gvAPP-180 is apparent (Figure 11, right).

It is important to note that the radially averaged contrast limit curve does a poor job capturing the variation in coverage and achieved contrast floor with PA, shown in Figure 8. Due to the geometry of the coronagraph, masked region, and sky rotation, the final coverage in terms of the number of science frames contributing to each pixel varies as a function of PA—especially at the innermost separations probed (Figure 9).




Figure 10. S/N map of the final derotated focal plane. To calculate the detection map, the optimized mode fraction k for a particular location (ρ, θ) is taken from the optimization process and applied to reducing the original data without an injected companion. The resulting S/N , measured in λ/D -spaced apertures and computed with the correction from Mawet et al. (2014), is shown in this figure. We found no $S/N \geq 4$ signals to report, and the lower S/N signals we inspected did not behave like a companion (moving with changing k , or disappearing) which points to a local maximum of the optimization process producing a spurious detection.

4.1. Interpretation as a Companion Mass Limit

Predicting the properties of planets that one would be sensitive to depends on choosing a model for planet evolution and spectra. We adopt the latest “Bobcat” models in the “Sonora” series of Marley et al. (2021), which span a range of $0.5 < M/M_J < 84$ in planet mass and $1 \text{ Myr} < t < 15 \text{ Gyr}$ in age. They model planets and substellar objects down to $T_{\text{eff}} = 200 \text{ K}$ with associated evolutionary tracks, high-resolution simulated spectra, and different metallicities. We perform our own synthetic photometry for this analysis using the published spectra.

The Bobcat models are published for $[M/H] \in \{-0.5, 0.0, +0.5\}$. Bond et al. (2017) models the evolution of Sirius A, reporting a slight metal deficiency relative to solar abundances and that an $[Fe/H]$ of -0.13 or -0.07 reproduces their observations. We adopt the $[M/H] = 0.0$ Bobcat models as a conservative option, because lower metallicity models have higher $[3.95]$ fluxes. As an example, an absolute $[3.95]$ mag of 14 corresponds to masses of $7.5, 8.5,$ and $8.8 M_J$ in the $[M/H] -0.5, 0.0,$ and $+0.5$ model suites, respectively.

4.1.1. Incorporating the Effects of Irradiation from Sirius A

Installation from Sirius A would be a major contributor to the effective temperature of a giant planet at a small separation from the star. Males et al. (2014) showed that the mass–separation–absolute-magnitude relationship depends only weakly on mass at small separations, effectively making lower-mass planets brighter and thus more accessible to current-generation instruments. To account for this in our mass limits, we computed the equilibrium temperature

$$T_{\text{eq}} = T_{\text{Sirius}} \sqrt{\frac{R_{\text{Sirius}}}{2a}} (1 - A_B)^{1/4}, \quad (1)$$




Figure 11. The line labeled “this work, median” indicates the median contrast level for which we would detect a source at $S/N = 5$, taken over the contrast levels for all PAs at that separation. The variation of contrast with PA is captured by the shaded gray region on either side of the line. Left: the contrast levels as a function of separation in arcseconds with the Thalmann et al. (2011), Vigan et al. (2015), and Pathak et al. (2021) reported curves for comparison. Right: the contrast levels scaled into λ/D units to compare the inner working angle across wavelengths and telescopes more effectively. When scaled by the telescope diameter and wavelength, the ability of the gvAPP-180 coronagraph to work at extremely small inner working angles is apparent.

where A_B is the Bond albedo and a is the separation between star and planet. Marley et al. (1999) found that cloud-free planets with incident flux from an A star have a mass- and temperature-dependent albedo ranging from $0.39 \leq A_B \leq 0.43$ in their simulations. That range extends to $0.39 < A_B < 0.93$ when considering different cloud models. For this analysis, we adopt $A_B = 0.5$ for consistency with Jupiter (Li et al. 2018).

The giant planets we might expect to find will glow with the heat of their formation and cool slowly over time to their equilibrium temperatures. We approximate the contribution to the equilibrium temperature T_{eq} from incident light by summing the planet luminosity

$$L_{eff} = L_{evol} + L_{eq}, \quad (2)$$

where

$$L_{evol} = \sigma_{sb} 4\pi R_{planet}^2 T_{evol}^4, \quad (3)$$

$$L_{eq} = \sigma_{sb} 4\pi R_{planet}^2 T_{eq}^4, \quad (4)$$

and T_{evol} is the evolutionary temperature from the Bobcat isochrones. Substituting the Stefan–Boltzmann law for L_{eff} gives us the following relationship for T_{eff} of the irradiated planet

$$T_{eff}^4 = T_{evol}^4 + T_{eq}^4, \quad (5)$$

This modified relationship in the specific case of Sirius and the [3.95] filter is shown in Figure 12. We do not claim to have




Figure 12. The masses detectable at a host-to-planet contrast of 10^{-5} and 5×10^{-6} are shown overlaid on a contour plot of planet absolute [3.95] mag predicted by the Bobcat models for $t = 242$ Myr and $T_{eff} = 9910$ K for Sirius.

captured any of the subtleties of the evolution of an irradiated planet, only to have chosen an appropriate T_{eff} subject to what we know about the minimum equilibrium temperature for an assumed separation and albedo value. Because of the way the Bobcat models were initialized, models with very high T_{eff} for very low masses were not available. Therefore, the effective minimum modeled mass for our analysis is $1 M_J$.

4.1.2. Magnitude-to-mass Relationship

To obtain the relationship between the modeled masses and their fluxes in the [3.95] filter, we assumed a system age of 242 Myr based on the determination of Sirius A’s age in Bond et al. (2017). For every modeled companion mass, T_{eff} and surface gravity were computed from the evolution tracks




Figure 13. Correspondence of [3.95] absolute magnitudes to masses assuming $A_B = 0.5$. Main panel: the relationship between magnitude and mass is plotted at an age of 242 Myr and several different equilibrium temperatures to illustrate its dependence on T_{eq} . The $T_{eq} = 685$ K line corresponds to the relationship at the smallest separation probed at the IWA. The dashed line indicates the median absolute-magnitude value of the contrast floor across all focal plane locations. The apparent discontinuity around $12 M_J$ is where the relationship becomes double valued as deuterium burning begins, and we take the lower value of the two. Lower panel: histogram of S/N = 5 detection levels reached at the probe locations shown in Figure 9 in 0.25 mag bins. Right panel: histogram of probe locations reaching particular mass sensitivity limits in $1 M_J$ bins.

subject to the modification described above. We then interpolated a spectrum from the Bobcat grid, and computed a [3.95] absolute Vega magnitude for such an object.

To compute the synthetic photometry, we adopted the latest Vega model from HST CALSPEC¹¹ (Bohlin et al. 2014) and an atmosphere with 2 mm of PWV and $\sec(z) = 1.15$. The modeled atmosphere retrieved from ESO SkyCalc (Noll et al. 2012; Jones et al. 2013) represents the La Silla observatory, quite close to Las Campanas where these data were taken, and its slightly lower altitude means its extinction is slightly pessimistic compared to the actual situation.

With a [3.95] mag for each mass computed, we inverted the relationship. This relationship is shown for a selection of T_{eq} values in Figure 13. Where the relationship is double valued for masses around the minimum mass for deuterium burning, we made it monotonic by taking the minimum mass value for that magnitude, resulting in the apparent discontinuity around $m_{[3.95]} = 11.5$.

4.1.3. Mass Limits

The minimum mass we would detect at S/N = 5 is shown as a function of separation in Figure 14. The overlaid scatter points and shaded region illustrate the variation in that mass limit as a function of angle at that separation, which is a

consequence of the rotation of the dark-hole region, changes in conditions over the course of the observation, etc.

To visualize the two-dimensional variation in the mass limit, when considering a point at angular coordinates (ρ, θ) in the focal plane, $a = d_{\text{Sirius}} \tan \rho$ is used as the separation for the purpose of computing a modified magnitude-to-mass relationship. The minimum detectable masses at each position are shown in Figure 15, adopting a Bond albedo of $A_B = 0.5$.

While this is sufficient to exclude a planet at (ρ, θ) of that mass (or greater) and that albedo at separation a , it cannot necessarily exclude a planet of that mass but at a greater separation that happens to fall on (ρ, θ) in the focal plane when viewed from Earth. A planet at that projected separation, but far enough away to experience negligible irradiation, is equivalent to a planet with an albedo of $A_B = 1.0$. The mass limits under this assumption are given in Figure 16.

Because this relationship depends on the projected separation (for the contrast limit) and the true separation (for the T_{eq}), it is more convenient and informative to fold these variables into the completeness analysis.

4.2. Completeness Limits

Our sensitivity is a function of the contrast floor (itself a function of position in the focal plane), companion mass, semimajor axis, and distance from the host star at the moment of observation. To capture this in our completeness analysis,

¹¹ alpha_lyr_mod_004.fits




Figure 14. The minimum mass detectable at $S/N = 5$ as predicted from our contrast levels using the Marley et al. (2021) models. The blue line and points are from this work. The contrast curves reported by Thalmann et al. (2011; Subaru/IRCS), Vigan et al. (2015; SPHERE), and Pathak et al. (2021; VISIR) have been reanalyzed using the Bobcat models for spectra and evolution, incorporating a T_{eq} correction as described in Section 4.1 with $A_B = 0.5$. The shaded region encompasses the minimum-to-maximum mass across all sampled points at that radius. Points at $1 M_J$ are likely overestimates, as that is the lower limit of the range of the model grid (shown as a cross-hatched region for out-of-bounds masses).




Figure 15. Visualization of the mass limit. This minimum mass is detectable when the separation is entirely in the plane of the sky; in other words, when the planet is experiencing the maximum irradiation for that projected separation. Since the irradiation depends on the actual separation, while the contrast limit depends on the projected separation, it is easier to interpret these results through completeness plots.

we drew 10^5 random orbital elements for a grid of masses and semimajor axis values. The eccentricities are drawn from a linearly descending prior given in Nielsen et al. (2019).

Projecting the position of the notional companion into the plane of the sky gives us the separation and PA we would observe, while the orbit determines the actual separation and hence the irradiation-affected effective temperature. The assumed age and metallicity, plus the companion mass and the projected and actual separations, then allow us to compute a [3.95] mag for the companion. In Figure 17, the fraction of the trial orbits that results in a [3.95] mag brighter than our $S/N = 5$ detection level gives the completeness at each mass and semimajor axis point.

Using the $S/N = 5$ contrast versus separation curves of each of the studies, as well as their dates of observation, we were able to assess whether a randomly drawn orbit would have been detected in at least one of the studies. We converted the other studies' contrast limits into mass limits following the same procedure as our own, assuming radial symmetry. For each point in (mass, semimajor axis) space, we drew 10^5 sets of orbital elements. Each was converted into a set of positions along the orbit using the observation dates of the studies.

We used the resulting true separations and projected separations to generate synthetic photometry in the bandpasses corresponding to each published contrast curve. As before, irradiation from the primary was taken into account. If any of the studies would have detected the planet at the corresponding position for their observation date, we counted it as detectable in the final combination of the completeness grid. The overall completeness from all four studies as a function of mass and semimajor axis is shown in Figure 18.




Figure 16. The minimum mass detectable at $S/N=5$ as predicted from our contrast levels using the Marley et al. (2021) models, without irradiation ($A_B = 1.0$). The blue line and points are from this work.




Figure 17. The fraction of planets in randomly drawn orbits around Sirius A that would be detectable in these observations, as a function of semimajor axis and mass. (Note that the separation axis is scaled differently from that in Figure 14.)

5. Considering a Second Epoch of Formation Triggered by Sirius B

The previous analysis assumed that any planets orbiting Sirius A would have formed contemporaneously with their host star. The less time that has elapsed since the time of formation, naturally the brighter a self-luminous giant planet would be, and thus easier for us to detect. If the red-giant stage of Sirius B altered the circumstellar environment of Sirius A to trigger a second epoch of planet formation, would we be sensitive to it?

A 101–126 Myr approximate nuclear lifetime is predicted by models of the Sirius B progenitor (Liebert et al. 2005). Since we have adopted 242 Myr for the system age, the more conservative estimate of a second epoch of planet formation would give us an age of 141 Myr. Calculating the synthetic




Figure 18. The fraction of planets in randomly drawn orbits around Sirius A that would be detectable in these observations and/or a past study of Sirius A, as a function of semimajor axis and mass. When combining the results of all four studies, at least one study would have detected a $9 M_J$ planet in the 2.5–7 au range for 99% of trials.




Figure 19. Minimum masses detectable at $S/N=5$ at our contrast levels using the Marley et al. (2021) models when assuming a second epoch of planet formation at $t = 141$ Myr. The blue line and points are from this work. The contrast curves from the literature have also been reanalyzed using the Bobcat models for spectra and evolution, incorporating a T_{eq} correction as described in Section 4.1. (Points at $1 M_J$ are likely overestimates, as that is the limit of the model grid.)

photometry for model planets of this age gives us the predicted mass limits in Figures 19 and 20. The completeness as a function of mass and separation increases accordingly, as shown in Figure 21.

6. Conclusions

We have demonstrated the use of a gvAPP-180 coronagraph to search for planets, allowing us to probe quite small separations from Sirius A. Expressed in λ/D , the inner working angle advantages of the gvAPP-180 are even more




Figure 20. Minimum masses detectable at $S/N = 5$ at our contrast levels using the Marley et al. (2021) models assuming a second epoch of planet formation, without irradiation. The blue line and points are from this work.




Figure 21. The fraction of planets formed in a hypothetical second period of formation at $t = 141$ Myr in randomly drawn orbits around Sirius A that would be detectable in these observations, as a function of semimajor axis and mass.

apparent. We have reanalyzed the reported contrast curves from previous high-contrast imaging searches with a common set of models, and report the overall completeness from all four studies.

Bond et al. (2017) report that stable planetary orbits around Sirius A exist up to a maximum period of 2.24 yr, citing a numerical stability analysis of planets in binary systems by Holman & Wiegert (1999). Assuming a circular orbit, this corresponds to a semimajor axis of 2.2 au. The only regions of Figure 18 with high completeness in our synthesis of previous direct-imaging studies are at larger semimajor axis values than this cutoff. This implies that we have barely begun to probe the region where a companion could exist in a stable orbit.

The albedo range reported in Marley et al. (1999) of $0.39 < A_B < 0.93$ lets us predict equilibrium temperatures for planets at $r = 2.2$ au of $217 \text{ K} < T_{\text{eff}} < 372 \text{ K}$. The low-mass end

of the Bobcat models is $0.5 M_J$, which will have cooled to 191 K by 242 Myr. Thus, 217 K as the pessimistic T_{eq} is still greater than the temperature predicted by the evolution model. Applying Equation (5) predicts $T_{\text{eff}} = 244 \text{ K}$ from the evolutionary and equilibrium temperatures. The model spectrum for a $0.5 M_J$ self-luminous planet divided by a reference Sirius spectrum predicts that the contrast will be lowest— 6×10^{-6} —at $4.6 \mu\text{m}$, when considering the $1\text{--}5 \mu\text{m}$ range of near-infrared detectors like Clio and JWST NIRC2. Aside from the inflection in contrast ratio around $4.6 \mu\text{m}$, the contrast ratio continues falling as wavelength increases, but achieving the required contrast and inner working angle becomes more challenging.

In our data reduction, we have shown that retaining the full data set rather than coadding can result in improved contrast limits. This creates challenges from the increased data volume, but we have shown that the time-domain PCA-based starlight-subtraction algorithm PCAT provides advantages in computational cost on large data sets.

The gvAPP-180 coronagraph on Clio, the ancestor of updated designs like those in ERIS and LBT (Kenworthy et al. 2018; Doelman et al. 2021), presents the same challenges in data reduction common to those coronagraphs: chiefly, dealing with the paired, nonradially symmetric PSF. We evaluated multiple methods of combining the data, whether performing starlight subtraction on both full PSFs, analyzing each dark-hole region separately, or—as we finally concluded was optimal—dividing out the AO-performance-related variation in the primary star to normalize the dark-hole regions into relative units, and then constructing a single-observation vector using both dark-hole regions.

The vAPP coronagraph makes the limitations of the one-dimensional contrast limit curve apparent when it comes to capturing the actual sensitivity of a set of observations. Other asymmetric coronagraphs like the PAPLC (Por 2020; Haffert et al. 2022) will also provide similarly complex contrast surfaces, and part of this work’s contribution is as an example of quantifying and visualizing that variation.

The appeal of direct-imaging studies of Sirius is no doubt in part due to its many convenient properties: brightness, proximity, and a wealth of existing work. However, successive IR direct-imaging studies have not revealed any new companions, and the contour of 95% completeness pushes ever inwards and downwards in mass sensitivity. Still, the absence of a detection with current facilities cannot be used to rule out the existence of any planet entirely. For example, a Jupiter-mass planet on a Jupiter-like orbit cannot be ruled out even after synthesizing the results of these previous studies. And, of course, reflected-light exoplanet imaging enabled by next-generation ground-based telescopes will probe a new mass–separation regime altogether.

Support for J.R.M. to conduct these observations was provided, in part, under contract with the California Institute of Technology (Caltech)/Jet Propulsion Laboratory (JPL) funded by NASA through the Sagan Fellowship Program executed by the NASA Exoplanet Science Institute. K.M.M.’s and L.M.C.’s work was supported in part by the NASA Exoplanets Research Program (XRP) by cooperative agreement NNX16AD44G.

This work has been supported by the Heising-Simons Foundation award #2020-1824 and NSF MRI Award #1625441 (MagAO-X). Parts of this research were done using

services provided by the OSG Consortium (Pordes et al. 2007; Sfiligoi et al. 2009), which is supported by the National Science Foundation awards #2030508 and #1836650. S.Y.H. is supported by NASA through the NASA Hubble Fellowship grant #HST-HF2-51436.001-A awarded by the Space Telescope Science Institute, which is operated by the Association of Universities for Research in Astronomy, Incorporated, under NASA contract NAS5-26555. This work also made use of High Performance Computing (HPC) resources supported by the University of Arizona.

The authors would also like to thank the anonymous reviewer for their comments which have improved this work.

Facilities: Magellan: Clay (MagAO, Clio).

Software: Astropy (Astropy Collaboration et al. 2018), Ray (Moritz et al. 2018), POPPY (Perrin et al. 2016), matplotlib (Hunter 2007).

Appendix Spatial Distributions of the Hyperparameters Found

Figure 22 shows the distribution across the focal plane of the hyperparameters tuned by the optimization process. The

maximum number of modes $k_{\max}$ is different at different locations in the focal plane due to the number of possible modes changing with the number of pixels masked out by the annular mask, as well as the total number of frames changing with coadding. For instance, when the number of frames is limiting the number of modes, a mode fraction of $k/k_{\max} = 0.5$ could be ≈ 5000 modes with no coadding or ≈ 100 modes with 50:1 coadding. Furthermore, the optimizer is not well suited to integer-valued hyperparameters so it works with $k/k_{\max}$, which is then converted into a number of modes to perform the starlight subtraction within the function to be optimized. For this reason, Figure 22(a) shows the mode fraction divided by the number of coadded frames to give a more accurate impression of the number of modes rather than the mode fraction.

The coadding fraction was grid searched rather than smoothly varied, so Figure 22(b) is shown without any attempt to interpolate between sample points smoothly. It is perhaps noteworthy that there are regions at $\theta = 0$ and 180° where the coverage was lowest (see Figure 9) that the optimizer favored the maximum amount of coadding.




Figure 22. Visualization of the spatial distributions of hyperparameters found by the optimization process. (a) Distribution of the optimal modes fraction $k/k_{\max}$ at each position scaled by $1/n$, where n is the optimal number of frames to coadd from the $n \in \{1, 10, 50\}$ values evaluated. (b) Spatial distribution of the optimal number of frames to coadd $n \in \{1, 10, 50\}$. (c) Optimized amplitude of injected signal used to measure an S/N and compute the limiting amplitude where $S/N = 5$. Since injecting a signal with an amplitude much greater than the limit will cause us to overestimate the limit, we vary the amplitude in the optimization process as well as described in Section 3.6.

ORCID iDs

Joseph D. Long  <https://orcid.org/0000-0003-1905-9443>
 Jared R. Males  <https://orcid.org/0000-0002-2346-3441>
 Sebastiaan Y. Haffert  <https://orcid.org/0000-0001-5130-9153>
 Logan Pearce  <https://orcid.org/0000-0003-3904-7378>
 Mark S. Marley  <https://orcid.org/0000-0002-5251-2943>
 Katie M. Morzinski  <https://orcid.org/0000-0002-1384-0063>
 Laird M. Close  <https://orcid.org/0000-0002-2167-8246>
 Gilles P. P. L. Otten  <https://orcid.org/0000-0002-6717-1977>
 Frans Snik  <https://orcid.org/0000-0003-1946-7009>
 Matthew A. Kenworthy  <https://orcid.org/0000-0002-7064-8270>
 Christoph U. Keller  <https://orcid.org/0000-0002-1368-841X>
 Philip Hinz  <https://orcid.org/0000-0002-1954-4564>
 John D. Monnier  <https://orcid.org/0000-0002-3380-3307>
 Alycia Weinberger  <https://orcid.org/0000-0001-6654-7859>
 Volker Tolls  <https://orcid.org/0000-0003-1841-2241>

References

- Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, *AJ*, **156**, 123
- Bailer-Jones, C. A. L., Rybizki, J., Fouesneau, M., Demleitner, M., & Andrae, R. 2021, *AJ*, **161**, 147
- Benest, D., & Duvent, J. L. 1995, *A&A*, **299**, 621
- Bessel, P. 1844, *MNRAS*, **6**, 136
- Bohlin, R. C., Deustua, S. E., & de Rosa, G. 2019, *AJ*, **158**, 211
- Bohlin, R. C., Gordon, K. D., & Tremblay, P.-E. 2014, *PASP*, **126**, 711
- Bond, G. 1862, *AN*, **57**, 131
- Bond, H. E., Schaefer, G. H., Gilliland, R. L., et al. 2017, *ApJ*, **840**, 70
- Close, L. M., Males, J. R., Morzinski, K., et al. 2013, *ApJ*, **774**, 94
- Doelman, D. S., Snik, F., Por, E. H., et al. 2021, *ApOpt*, **60**, D52
- Ghezzi, L., Montet, B. T., & Johnson, J. A. 2018, *ApJ*, **860**, 109
- Haffert, S. Y., Males, J., Van Gorkom, K., et al. 2022, *Proc. SPIE*, **12185**, 1218581
- Holman, M. J., & Wiegert, P. A. 1999, *AJ*, **117**, 621
- Hunter, J. D. 2007, *CSE*, **9**, 90
- Hunziker, S., Quanz, S. P., Amara, A., & Meyer, M. R. 2018, *A&A*, **611**, A23
- Johnson, J. A., Fischer, D. A., Marcy, G. W., et al. 2007, *ApJ*, **665**, 785
- Jones, A., Noll, S., Kausch, W., Szyszka, C., & Kimeswenger, S. 2013, *A&A*, **560**, A91
- Kenworthy, M. A., Codona, J. L., Hinz, P. M., et al. 2007, *ApJ*, **660**, 762
- Kenworthy, M. A., Snik, F., Keller, C. U., et al. 2018, *Proc. SPIE*, **10702**, 1070246
- Kervella, P., Thévenin, F., Morel, P., Bordé, P., & Di Folco, E. 2003, *A&A*, **408**, 681
- Krige, D. G. 1951, PhD thesis, Univ. the Witwatersrand
- Li, L., Jiang, X., West, R. A., et al. 2018, *NatCo*, **9**, 3709
- Liebert, J., Young, P. A., Arnett, D., Holberg, J. B., & Williams, K. A. 2005, *ApJ*, **630**, L69
- Long, J. D., & Males, J. R. 2021, *AJ*, **161**, 166
- Long, J. D., Males, J. R., Morzinski, K. M., et al. 2018, *Proc. SPIE*, **10703**, 107032T
- Males, J. R., Close, L. M., Guyon, O., et al. 2014, *Proc. SPIE*, **9148**, 914820
- Males, J. R., Fitzgerald, M. P., Belikov, R., & Guyon, O. 2021, *PASP*, **133**, 104504
- Marley, M. S., Gelino, C., Stephens, D., Lunine, J. I., & Freedman, R. 1999, *ApJ*, **513**, 879
- Marley, M. S., Saumon, D., Visscher, C., et al. 2021, *ApJ*, **920**, 85
- Marois, C., Lafreniere, D., Doyon, R., Macintosh, B., & Nadeau, D. 2006, *ApJ*, **641**, 556
- Mawet, D., Milli, J., Wahhaj, Z., et al. 2014, *ApJ*, **792**, 97
- Moritz, P., Nishihara, R., Wang, S., et al. 2018, in *Proc. 13th USENIX Conference on Operating Systems Design and Implementation* (Berkeley, CA: USENIX Association), 561
- Morzinski, K. M., Males, J. R., Skemer, A. J., et al. 2015, *ApJ*, **815**, 108
- Nielsen, E. L., De Rosa, R. J., Macintosh, B., et al. 2019, *AJ*, **158**, 13
- Noll, S., Kausch, W., Barden, M., et al. 2012, *A&A*, **543**, A92
- Otten, G. P. P. L., Snik, F., Kenworthy, M. A., et al. 2017, *ApJ*, **834**, 175
- Pathak, P., Petit dit de la Roche, D. J. M., Kasper, M., et al. 2021, *A&A*, **652**, A121
- Perrin, M., Long, J., Douglas, E., Sivaramakrishnan, A., & Slocum, C. 2016, *POPPY: Physical Optics Propagation in PYthon, Astrophysics Source Code Library*, ascl:1602.018
- Por, E. H. 2020, *ApJ*, **888**, 127
- Pordes, R., Petravick, D., Kramer, B., et al. 2007, *JPhCS*, **78**, 012057
- Samland, M., Bouwman, J., Hogg, D. W., et al. 2021, *A&A*, **646**, A24
- Schroeder, D. J., Golimowski, D. A., Bruckardt, R. A., et al. 2000, *AJ*, **119**, 906
- Sfiligoi, I., Bradley, D. C., Holzman, B., et al. 2009, in *2009 WRI World Congress on Computer Science and Information Engineering* (New York: IEEE), 428
- Snik, F., Otten, G., Kenworthy, M., et al. 2012, *Proc. SPIE*, **8450**, 84500M
- Snoek, J., Larochelle, H., & Adams, R. P. 2012, in *Proc. 25th Int. Conf. on Neural Information Processing Systems*, Vol. 2 (Red Hook, NY: Curran Associates Inc.), 2951
- Soummer, R., Pueyo, L., & Larkin, J. 2012, *ApJ*, **755**, L28
- Sutcliffe, B. J., Bohn, A. J., Birkby, J. L., et al. 2021, *MNRAS*, **506**, 3224
- Thalmann, C., Usuda, T., Kenworthy, M., et al. 2011, *ApJ*, **732**, L34
- Vigan, A., Gry, C., Salter, G., et al. 2015, *MNRAS*, **454**, 129
- Yang, L., & Shami, A. 2020, *Neurocomputing*, **415**, 295