The Astronomical Journal, 165:57 (14pp), 2023 February https://doi.org/10.3847/1538-3881/aca60d

The Giant Accreting Protoplanet Survey (GAPlanetS): Optimization Techniques for Robust Detections of Protoplanets

Jéa I. Adams Redai[1,2], Katherine B. Follette[2], Jason Wang (王劲飞)[3,11], Clare Leonard[2,4], William Balmer[2,5,6], Laird M. Close[7], Beck Dacus[2], Jared R. Males[7], Katie M. Morzinski[8], Joseph Palmo[2], Laurent Pueyo[5,6], Elijah Spiro[2,9], Helena Treiber[2], Kimberly Ward-Duong[10], and Alex Watson[2]
[1] Center for Astrophysics |Harvard & Smithsonian, 60 Garden Street, Cambridge, MA 02138, USA; jea.adams@cfa.harvard.edu
[2] Amherst College Department of Physics and Astronomy, PO Box 5000, Amherst, MA 01002-5000, USA
[3] Department of Astronomy, California Institute of Technology, Pasadena, CA 91125, USA
[4] Epic Systems, 1979 Milky Way, Verona, WI 53593, USA
[5] Department of Physics & Astronomy, Johns Hopkins University, 3400 N. Charles Street, Baltimore, MD 21218, USA
[6] Space Telescope Science Institute, 3700 San Martin Drive, Baltimore MD 21218, USA
[7] Steward Observatory, University of Arizona, Tucson, AZ 85721, USA
[8] Steward Observatory, University of Arizona, 933 N. Cherry Avenue, Tucson, AZ 85721, USA
[9] NASA Kennedy Space Center, NASA Kennedy Space Center, Fl, 32899, USA
[10] Department of Astronomy, Smith College, Northampton, MA 01063 USA
Received 2022 June 28; revised 2022 November 17; accepted 2022 November 23; published 2023 January 18

Abstract

High-contrast imaging has afforded astronomers the opportunity to study light directly emitted by adolescent (tens of megayears) and "proto" (<10 Myr) planets still undergoing formation. Direct detection of these planets is enabled by empirical point-spread function (PSF) modeling and removal algorithms. The computational intensity of such algorithms, as well as their multiplicity of tunable input parameters, has led to the prevalence of ad hoc optimization approaches to high-contrast imaging results. In this work, we present a new, systematic approach to optimization vetted using data of the high-contrast stellar companion HD 142527 B from the Magellan Adaptive Optics Giant Accreting Protoplanet Survey (GAPlanetS). More specifically, we present a grid search technique designed to explore three influential parameters of the PSF subtraction algorithm `pyKLIP`: annuli, movement, and KL modes. We consider multiple metrics for postprocessed image quality in order to optimally recover at Hα (656 nm) synthetic planets injected into contemporaneous continuum (643 nm) images. These metrics include peak (single-pixel) signal-to-noise ratio (S/N), average (multipixel average) S/N, 5σ contrast, and false-positive fraction. We apply continuum-optimized KLIP reduction parameters to six Hα direct detections of the low-mass stellar companion HD 142527 B and recover the companion at a range of separations. Relative to a single-informed, nonoptimized set of KLIP parameters applied to all data sets uniformly, our multimetric grid search optimization led to improvements in companion S/N of up to 1.2σ, with an average improvement of 0.6σ. Since many direct imaging detections lie close to the canonical 5σ threshold, even such modest improvements may result in higher yields in future imaging surveys.

Unified Astronomy Thesaurus concepts: Direct imaging (387); Exoplanet detection methods (489); Astronomy data analysis (1858); Principal component analysis (1944)

1. Introduction

Over the past decade, high-contrast direct imaging has uncovered dozens of bound substellar companions to higher-mass stars (Bowler 2016; Currie et al. 2022). This technique is generally sensitive to faint companions at separations of >0.″1 and masses on the order of several Jupiter masses or larger. Imaging's ability to resolve the light emitted directly by exoplanet atmospheres makes it a powerful vehicle for planet characterization. Consequently, prospects for future work constraining planet composition, formation, and habitability are intertwined with refinement of the imaging techniques that will allow us to robustly isolate planetary signals (Seager & Deming 2010; Biller & Bonnefoy 2018).

Both current and future imaging campaigns are dependent on suppression of the stellar point-spread function (PSF). Raw data of imaged extrasolar systems are dominated by the diffraction-limited core of the stellar PSF, its broader seeing halo, and a field of uncorrected stellar "speckles." This leaves faint planets buried under the starlight in raw and conventionally combined data. The hardware (adaptive optics (AO), coronagraphs, apodizers, etc.) in some high-contrast imaging instruments, such as the Gemini Planet Imager (GPI; Macintosh et al. 2014) on the Gemini South telescope and the Spectro-Polarimetic High-contrast Exoplanet REsearch (SPHERE; Beuzit et al. 2019) on the Very Large Telescope (VLT), can suppress starlight and allow for raw planet/star contrasts (i.e., the 5σ noise level at planetary separations) of $\sim 10^{-4}$ to 10^{-5} (Bailey et al. 2016). However, detection of young planets in near-infrared thermal emission requires planet/star contrasts of at least 10^{-6}.

Postprocessing techniques, specifically PSF subtraction, can improve planet/star contrasts by a factor of 10–100 (see, e.g., Bailey et al. 2016). A range of algorithmic approaches are available to achieve this improved contrast, the most common

[11] 51 Pegasi b Fellow.

of which are the Karhunen–Loéve Image Processing (KLIP; Soummer et al. 2012; Pueyo 2016) and Locally Optimized Combinations of Images (LOCI; Lafrenière et al. 2007) techniques. Both techniques utilize angular differential imaging (ADI; Marois et al. 2006), in which the source is allowed to rotate in the image plane throughout the observation, while the instrumental PSF remains static. On-sky rotation of the source ensures that PSF features identified by the algorithms consist primarily of instrumental PSF features and exclude rotating high spatial frequency sources such as planets and narrow disk structures. In this work, we focus strictly on the principal component analysis (PCA) based technique, KLIP, implemented with the Python package `pyKLIP`[12] (Wang et al. 2015). Optimization of LOCI algorithms is discussed in Thompson & Marois (2021).

Despite their power to suppress the stellar PSF, postprocessing algorithms like KLIP are complex and highly tunable. Extracted photometry, astrometry, and spectroscopy of an exoplanet direct detection are greatly influenced by user-selected input parameters. The widely used KLIP algorithm `pyKLIP`, for example, utilizes 25 tunable input parameters that control features such as application of a high-pass filter, the complexity of the PSF model, the number and shape of regions for which model PSFs are constructed separately, and the size of the library of reference images.

Of particular concern regarding KLIP parameter choices are systems containing both planets and circumstellar material. Substructure is ubiquitous in planet-forming disks (Benisty et al. 2022), and disk features have the potential to appear planet-like in postprocessed images (e.g., K. B. Follette et al. 2017). A number of reported (proto)planet detections in these systems have been called into question when other techniques or data sets fail to reveal unambiguous planetary signals. These include the protoplanet candidates LkCa15 b and c (Kraus & Ireland 2012; Sallum et al. 2015), which have been contested by several papers highlighting their proximity to inner disk material (Thalmann et al. 2016; Currie et al. 2019). Similarly, the two planet candidates in the HD 100546 system (Quanz et al. 2013; Currie et al. 2015) were flagged as potential false positives because they were not recovered as orbiting point sources in other observations (Follette et al. 2017; Rameau et al. 2017).

We hypothesize that one cause of detection discrepancies among reported planet candidates is a lack of standardization in selecting PSF subtraction parameters. The de facto technique among the imaging community has been to make default choices for algorithmic parameters and to hand-tune those parameters once an apparent detection is made, or to optimize select parameters individually (Meshkat et al. 2014). These approaches are used in part because of the computational intensity of the KLIP algorithm, which does not lend itself well to optimization approaches, requiring thousands to millions of iterations.

In this work, we use Hα direct imagery of the 12–23 au separation $0.26^{+0.16}_{-0.14}\,M_{\odot}$ (Claudi et al. 2019) companion HD 142527 B to develop a `pyKLIP` optimization methodology. This well-characterized companion has been observed as part of the Giant Accreting Protoplanet Survey (K. B. Follette 2023, in preparation) over a long time baseline (2013–2018) and appears at a wide range of planet–star separations, making it more difficult to recover in some epochs than others.

This paper is organized as follows. We describe `pyKLIP` parameters of interest in Section 2. In Section 3, we detail our Magellan Adaptive Optics (MagAO) observations and basic data processing procedures. In Section 4, we outline our optimization approach, develop image quality metrics, and describe our selection of some fixed `pyKLIP` parameters. We present the final results of our optimizations in Section 5. Finally, we summarize our process and describe future steps in Section 6.

2. KLIP Parameters of Interest

All KLIP implementations rely on two fundamental elements of the algorithm: (1) the compilation of images to form a reference library, and (2) the construction of a custom PSF model for each image or region of an image from its PSF library, the complexity of which is controlled by the number of principal components ("KL modes"). "KL modes" are a set of orthogonal basis vectors oriented to describe the variance in the reference library images, with each additional mode describing a smaller proportion of the overall variance. They can be thought of as common patterns in the images. As the number of KL modes used to build a PSF model increases, the patterns identified appear in fewer of the reference images. Therefore, higher KL modes correspond to a more "aggressive" PSF subtraction. Of the 25 tunable user-input parameters to the `pyKLIP` algorithm, those that have a particularly marked effect on the quality of postprocessed images (even for bright companions, as shown in Figure 1) are as follows:

1. The `numbasis` parameter controls the number of "KL modes" that form the PSF model for subtraction, as described above. This parameter can be a single value or a list of values, in which case multiple PSF models will be created and subtracted to form multiple postprocessed images. The maximum number of KL modes that can be used to construct a PSF for a given image sequence is equal to the number of images in the reference image set (for ADI imagery, this is the number of images in the sequence minus any that were discarded through rotational masking; see below).
2. The `annuli` parameter controls the radial geometry of the optimization regions, specifically the number of annular zones within the image for which PSFs are constructed separately. In general, optimizing on smaller regions (a higher `annuli` parameter) means optimizing on fewer PSF features at once. These features include quasi-static speckles, the AO control radius, wind residuals, etc.
3. The `movement` parameter is effectively an ADI rotational mask. It sets the number of pixels that an astrophysical source (e.g., planet) located at the center of each annular zone is required to have rotated relative to the image for which the PSF is being constructed (the "target image") in order for another image in the sequence (a "reference image") to be included in the reference library. A smaller `movement` value corresponds to more "aggressive" subtraction, since more reference images with the planet located near the same position as in the target image are included in the reference library.

[12] https://bitbucket.org/pyKLIP/pyklip

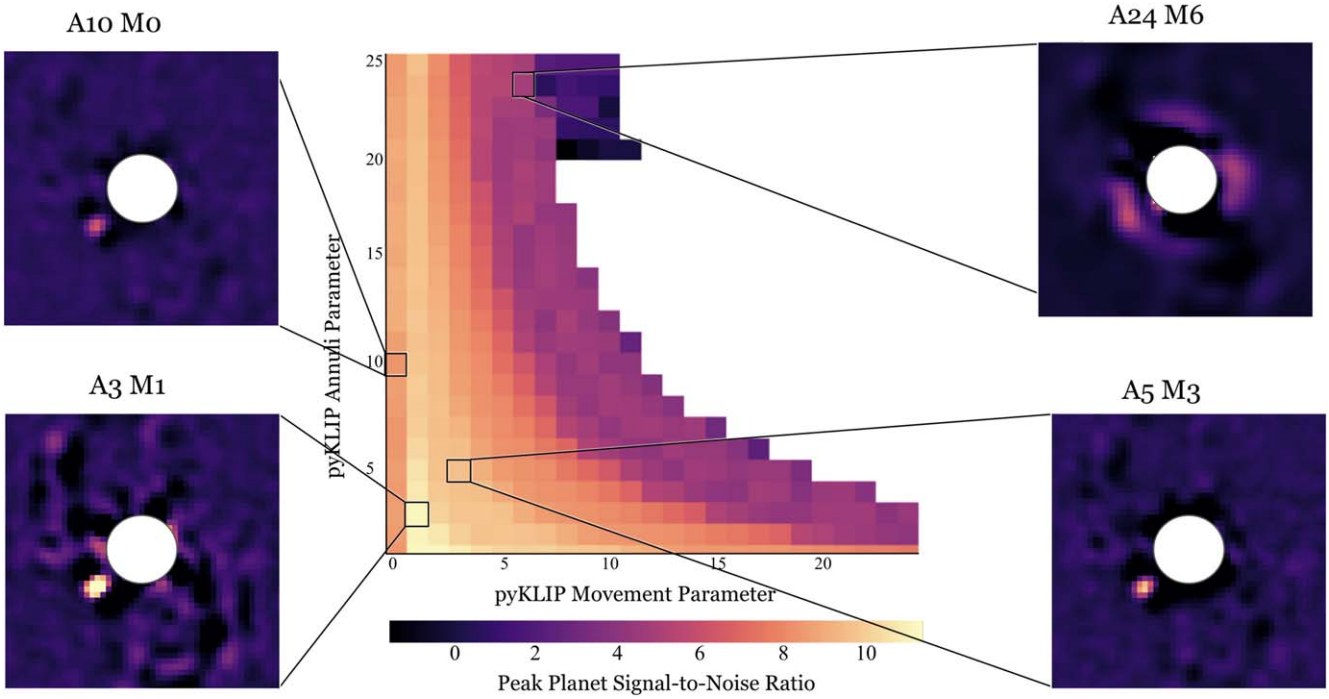

Figure 1. A depiction of `pyKLIP` output demonstrating the variation in the quality of recovery for a true companion (HD 142527 B) across a range of KLIP `movement` and `annuli` parameters. Each pixel in the central heat map represents a separate, independent KLIP reduction with the `movement` and `annuli` values depicted on the *x*- and *y*-axes, respectively. The color of that pixel reflects the highest single-pixel S/N value at the location of the companion in the S/N map. Inset images depict the S/N maps for a representative sample of individual reductions. While much of this parameter space yields recovery of the companion at $>5\sigma$, there is intense variation in the quality of the extraction, and even some parameters for which the signal from the bright companion is nearly absent. For this data set, an optimal combination of parameters to maximize the S/N of the companion for 10 KL modes is a `movement` value of 1 pixel and 3 `annuli`. An in-depth explanation of this grid search technique is provided in Section 4.2.

4. The `highpass` parameter improves starlight subtraction by attenuating low spatial frequency signals in an image before executing KLIP. `pyKLIP`'s Gaussian high-pass filter parameter controls the standard deviation of this filter, where a small standard deviation is considered aggressive. Very little low spatial frequency signal will survive an aggressive high-pass filter; however, planet light may also be attenuated in this process.

3. Data and Preprocessing

3.1. Observational Data

The data used in these analyses were taken with the Magellan Clay Telescope at Las Campanas Observatory using the MagAO (Close et al. 2013; Marchetti et al. 2014, 2016) system's visible light camera (VisAO; Males 2013; Males et al. 2014) in Hα simultaneous differential imaging (SDI) mode. The data were acquired between 2013 and 2018 as part of the Giant Accreting Protoplanet Survey (GAPlanetS K. B. Follette 2023, in preparation), a search for protoplanets inside of the gaps of transitional disk host stars. GAPlanetS data are processed with a custom IDL pipeline, as described in detail in Follette et al. (2017). In brief, they are dark-subtracted and divided by a flat-field image, which is generally acquired once per observing semester. This corrects primarily for attenuation of light by near-focus dust spots on the CCD window, as the VisAO CCD is otherwise flat to within 1%. The images are centered and aligned using Fourier cross-correlation, separated into Hα ($\lambda_c = 656$ nm, $\Delta\lambda = 6$ nm) and continuum ($\lambda_c = 642$ nm, $\Delta\lambda = 6$ nm) channels (acquired simultaneously on the detector), and cropped to a 451-pixel ($\sim 3\farcs5$) square. Images with cosmic rays within 50 pixels of the central star are removed from the image cube before analysis.

This work focuses on the transitional disk system HD 142527, which has a known $0.26^{+0.16}_{-0.14}\,M_\odot$ (Claudi et al. 2019) companion (Close et al. 2008; Biller et al. 2012, HD 142527 B) at separations ranging from 12.38 to 22.89 au (Balmer et al. 2022). In particular, we are interested in HD 142527 epochs where the companion is detectable with a range of KLIP parameters. These robust detections provide a stable "training" set of data for optimization. This was not the case for the GAPlanetS data collected on 2017 February 10, for which the companion was undetectable at signal-to-noise ratio (S/N) greater than 3 using any KLIP parameters. This is likely due to the limited rotation of this data set ($16\fdg1$) and the companion's tight separation (44.29 ± 2.57 mas; Balmer et al. 2022).

Seeing data indicated subarcsecond conditions for all data sets for which it was available; however, the site seeing was nonoperational during two of the epochs. The FWHM of the data sets ranges from 4 to 5.5 pixels (32–44 mas), with an average of 4.8 pixels or 38 mas, and on-sky rotation varies from $34\fdg8$ to $117\fdg4$, with an average of 74°. The properties of the data sets used are described in Table 1. Astrometry and photometry of the HD 142527 B companion from these data sets are discussed in detail in Balmer et al. (2022).

3.2. Data Set Selection

Multiwavelength imaging of targets with known companions is ideal for testing optimization of the `pyKLIP` algorithm. Wavelengths where the companion is fainter or not visible can be leveraged as "clean" images into which synthetic companions/planets can be inserted and optimized. Then, the efficacy

Table 1
GAPlanetS Observations of HD 142527 B Used in This Analysis

Date	Sep. (mas)	PA (deg)	Seeing	t_{int} (minutes)	r_{sat} (pixels)	Rot. (deg)	N_{ims}	FWHM (pixels)
11 April 13	81.08 ± 1.08	128.12 ± 0.49	0.56	74.2	6	65.3	1961	4.56
8 April 14	77.70 ± 1.68	117.01 ± 1.12	N/A	66.5	N/A	101.7	1758	4.00
15 May 15	70.16 ± 1.19	110.56 ± 0.80	0.55	90.3	N/A	117.4	2387	5.50
16 May 15	72.19 ± 2.02	107.84 ± 0.97	0.80	43.2	2	34.8	1143	5.01
18 May 15	70.00 ± 1.35	110.12 ± 0.72	0.66	79.5	9	76.8	159	5.24
27 April 18	44.34 ± 1.81	58.62 ± 1.67	N/A	48.3	3	49.2	580	4.37

Note. Orbital locations were derived from Balmer et al. (2022).

of a parameter optimization approach can be evaluated by applying it to target wavelengths.

In exploring data-driven approaches to KLIP optimization for GAPlanetS data, we focus on optimization of simulated planets injected at a range of separations into the continuum images for HD 142527 B data sets. We then apply this approach to the HD 142527 Hα data sets to test whether the approach results in robust single-epoch Hα recovery of the known accreting companion in these data sets. In Section 5, we discuss the S/N penalty of optimizing on injected planets injected into continuum images rather than directly on Hα images.

We note that we have focused our approach in this work on detection of planets at Hα alone, and have not applied Simultaneous Differential Imaging (SDI) as a part of our optimization approach i.e., We used continuum imagery to inject and optimize false planets, but we did not combine Hα and continuum images before or after PSF subtraction, treating them as wholly separate data sets. SDI results are reported in the GAPlanetS paper (K. B. Follette 2023, in preparation) released in conjunction with this work.

Although our analysis of this proposed approach to optimization relies on one object, the companion exists at a wide range of separations (~40–80 mas) across the 2013–2018 time baseline, and the data sets are of widely varying quality.

3.3. Pre-KLIP "Data Quality Cuts"

In order to reduce the computational scope of the optimization problem, we implemented "data quality cuts" prior to KLIP optimization as described in detail in K. B. Follette (2023, in preparation). In short, we fit the stellar PSF (or ghost in the case of saturated images) with a Moffat profile and extracted the peak value of this fit for every image in an image sequence. Using the peak value as a proxy for the "quality" of a given image, we then culled the data by discarding the images with the lowest peaks. We compared contrast curves (computed under a conservative set of KLIP parameters and with a high-pass filter of width of $0.5 \times$ FWHM) for the full image sequence to those with 5%, 10%, 20%, 30%, 40%, 50%, 60%, 70%, 80%, and 90% of the lowest-quality images discarded. We then chose a data quality cut by eye, balancing the overall contrast in the inner and outer regions of the image, generally adopting the cut with the highest achieved contrast (lowest contrast curve) in the inner 0.″1–0.″25 unless that cut was substantially worse in the outer regions of the image. The selected data quality cuts are listed in Table 2.

This technique, which builds on the principle of "Lucky Imaging" (Fried 1978), is a remarkably nuanced one. The "answer" for the optimal data quality cut for a given image set appears to vary with location within the image. We do not explore it in detail in this work, but we do note here that ultimately this parameter is likely also an important consideration for future optimization work and should in principle be optimized together with `KLIP` parameters.

Table 2
Injected Planet Parameters

Date	Cut	Fake Contrast	$N_{injected}$
2013 April 11	10	0.01	8
2014 April 8	0	0.01	5
2015 May 15	50	0.01	5
2015 May 16	80	0.05	5
2015 May 18	0	0.01	4
2018 April 27	0	0.05	8

Note. Fake contrast is the contrast of injected fake planets used to compute contrast curves (this same contrast is used to inject synthetic planets for pyKLIP-PE optimization). $N_{injected}$ is the number of injected planets between the IWA and control radius used to compute the optimal parameters.

3.4. Synthetic Planet Injection

The cornerstone of our proposed optimization approach is the assertion that KLIP optimization should be done on *simulated* companions to avoid cognitive biases in parameter selection. Therefore, another critical pre-KLIP step in our approach is injecting synthetic planets into the continuum images.

Among other considerations in planet injection is the region of the postprocessed image in which to optimize detections. In this work, we have opted to focus on the region inside of the AO system's control radius for planet insertion and recovery. This region (also known as the "dark hole"') is in most cases equivalent to or larger than the size of the cleared central cavities of GAPlanetS transitional disks, which is the region most likely to host detectable accreting protoplanets. In the case of HD 142527, however, the cavity is very large (~1″; Avenhaus et al. 2014). It is likely that optimization of companions in the outer portion of the HD 142527 cavity would yield different choices for optimal KLIP parameters than those reported here.

Injecting planets into raw data and recovering them at a range of separations and position angles (PAs) is a common method for quantifying azimuthal variation in recovered astrometry and photometry due to PSF asymmetries (e.g., Wagner et al. 2018). It also serves as a means to quantify the degree of flux attenuation ("throughput") introduced by the

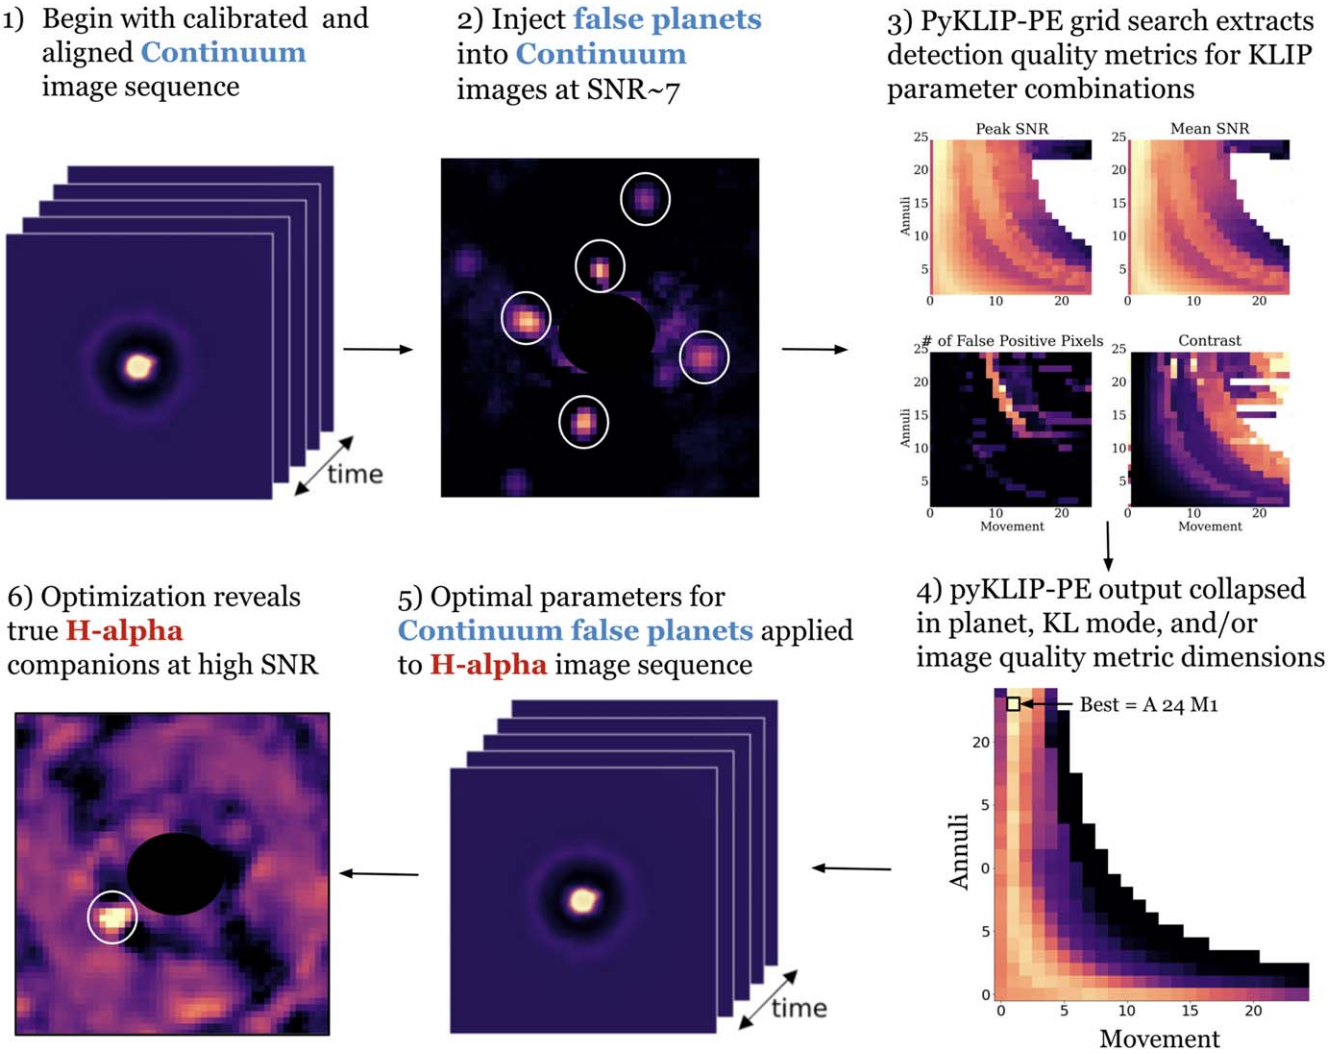

Figure 2. A schematic representation of the optimization process applied in this work, with images drawn from optimization of the HD 142527 2013 April 11 data set.

KLIP algorithm as a function of separation, which is needed to determine detection limits (Mawet et al. 2019).

This technique of inserting synthetic planet signals is also a potential tool for optimization. We hypothesize that if synthetic planets injected into continuum images are successful at mimicking real planets, their optimal KLIP parameters should result in high-quality recoveries of real companions in Hα data sets. This hypothesis is examined in detail in Section 5.2.

To this end, we injected planetary signals into the raw continuum image sequences (culled according to the chosen "data quality" cuts). The continuum images were chosen in order to reduce the influence of the companion on the results, though we note that it is recoverable at moderate S/N (S/N = 2–6; Balmer et al. 2022) in continuum wavelengths. In the case of HD 142527, we injected planets away from the companion to mitigate its effects, which are then limited to an inflation of the noise statistics at the companion's separation. In this way, the HD 142527 data mimic a case where optimization is being done on injected planets in the vicinity of real, previously unknown planets.

Synthetic planets are created by scaling the images of the central star (or ghost in the case of saturated data) to a desired contrast and injecting them into the raw images. The PSF of the injected planets is constructed from the stellar PSF image by image, so intrinsic variation in the stellar PSF is also captured in the injected planetary PSF (as would be the case with real planets). Contrasts of these injected planets are computed under a single set of KLIP parameters (`annuli` = 5, `movement` = 2) that ad hoc experimentation with MagAO data indicates will result in robust recovery of most companions. These injected planet contrasts were iterated upon until the average S/N across 5, 10, 20, and 50 KL modes was in the range of 6.5–7.5. We injected as many planets as would fit between the inner working angle (IWA; PSF FWHM) and the control radius separated by 1 FWHM radially. More specifically,

$$N_{\text{injected}} = \frac{\text{Control Radius} - \text{IWA}}{\text{FWHM}}. \quad (1)$$

The PAs of these fake injected planets were assigned such that the first planet would be placed at a PA of 0 and each subsequent planet would be advanced by a PA of 85°, probing different azimuthal regions of the PSF.

Step 2 in Figure 2 visualizes the stage in our reduction process where synthetic planet injection is done, and it shows the resulting fake planets. Table 2 shows the contrasts and number of fake injected planets in each data set.

4. Optimization Methods

The goal of this work is to develop a method for selecting KLIP parameter combinations that (1) is robust to false positives and (2) leads to high-quality recoveries of known objects. At the same time, we aim to develop a method that does not rely on the planetary signal itself for optimization. This approach is critical both in developing strategies for untargeted/uninformed exoplanet searches, where the location (s) of planet(s) are not known a priori, and in reducing cognitive biases in parameter optimization.

In other words, we should be able to apply these techniques to robustly recover planet signal without prior information about the planet's existence. In doing so, a first key question is what the metrics for a "good" planet recovery should entail. A second is how we should combine these metrics to refine an optimization approach. We apply several techniques to this problem, as outlined below.

4.1. Image Quality Metrics

In order to make decisions among possible `pyKLIP` parameter values, one or more metrics for the "quality" of the signal extraction are required. Due to the nuanced nature of postprocessed image quality, we chose to utilize three measurables in our approach to optimization: S/N maps, contrast curves, and false-positive thresholds.

4.1.1. Peak/Average S/N

A common metric for a high-quality recovery is a planet's S/N, i.e., its signal should be at least 3σ–5σ above the noise level, where the noise level is the standard deviation of background pixels at that separation.

In order to compute S/N maps, we mask a wedge-shaped region with a radial width of the PSF FWHM and an azimuthal width of 15° on either side of each planet, which masks the planet itself and all or most of the "self-subtraction lobes" that extend azimuthally from its location. We then estimate the noise at each separation by computing the standard deviation of the nonmasked (noise) pixels in 1 pixel `annuli` and apply a statistical correction following Mawet et al. (2014) to account for the small number of independent noise samples near the central star.

The classic measure of the quality of a high-contrast imaging detection is the maximum single-pixel value in an S/N map (which we will call "peak S/N" hereafter) of a recovered companion. However, we also extract the average S/N of the positive pixels under our planet masks as an alternative, potentially more robust, measure of the detection quality. It should be noted that, in our case, S/N is measured on postprocessed images that have been subject to a high-pass filter, and that the final postprocessed images were smoothed with a 1 pixel Gaussian kernel.

4.1.2. Contrast

Another common metric for the quality of a high-contrast imaging reduction is the achieved planet/star contrast limit. The optimal reduction by this metric should have the best planet/star contrast (lowest contrast value), allowing for the recovery of the faintest objects.

As is standard in the field, we compute contrast by (1) measuring the 5σ noise at a given location and (2) calibrating that noise to correct for the algorithm throughput. Throughput is computed as the ratio of a planet's true (injected) brightness to its postprocessed (recovered) brightness. The "raw" 5σ noise level divided by the throughput yields a limiting brightness for planets to be recovered at 5σ at a given location in the image.

4.1.3. False-positive Fraction

High-quality HCI reductions are also those in which all signals meeting the canonical S/N threshold are true signals and not false positives. On the assumption that there are no additional true signals in our images, we consider false-positive pixels to be those with values in the postprocessed maps with S/Ns above 5σ that are not at the location of the known or injected companion(s). We count the number of pixels between the IWA and control radius that meet this threshold, excluding those under the planetary mask(s). We note that more nuanced approaches to false-positive estimation are possible. Isolated single pixels with high values, for example, are less likely to be mistaken for a companion than clusters of pixels with high values. Previously unknown real companions that have not been masked and are visible at the optimization wavelength will also influence this value, as is the case in all of our data sets. Future work should implement a more nuanced version of this metric, perhaps also incorporating a forward-modeling approach (e.g., Ruffio et al. 2017).

4.1.4. Neighbor Quality

"Neighbor quality" metrics are created by smoothing the peak and average S/N metrics in movement/annuli space by a Gaussian with an FWHM of 3 pixels. They serve to create a measure of the quality of neighboring parameters in movement/annuli space. In other words, a high S/N value that is an outlier in a region of the annuli/movement heat map is penalized by the low S/Ns of neighboring values. Parameter regions that are stably high, such that small changes in movement/annuli values result in minimal changes in S/N, are given additional weight. A neighbor quality metric is created for both peak S/N and average S/N.

4.2. pyKLIP-PE

The final tool in our approach to optimization is the `pyKLIP` Parameter Explorer (`pyKLIP-PE`). This coarse grid search algorithm runs `pyKLIP` for every combination of specified `movement` and `annuli` parameter values for a specified range of `KL mode` values and returns all of the image quality metrics specified in Section 4.1.

The output for each data set is a 5D cube with the following dimensions: `annuli`, `movement`, `numbasis` (KL mode), planet, and image quality metric. In this work, we explore `annuli` values of 1–25, `movement` values of 0–24, and a KL basis set of `numbasis = [1,2,3,4,5,10,20,50,100]`. The maxima of the annuli and movement ranges are meant to limit the computational time of each run of the algorithm but can be expanded or reduced as necessary. KL mode values 1–5 were selected owing to the fact that they capture the most frequently occurring PSF modes. KL modes 10, 20, 50, and 100 were subsequently chosen in an attempt to capture a range of complexities gained by adding more KL modes.

To produce uniformity among data sets, we standardized each of the image quality metrics to take values between ∼0 and 1 by subtracting the minimum value and dividing by the

maximum–minimum value, so that 1 is the "best" combination of parameters and 0 is the "worst." In each case, this normalization is computed in `movement/annuli` space across KL modes so that for each planet, metric, and KL mode combination the "optimal" answer may vary across these parameters.

For contrast, we compute an image quality metric in log contrast space that ranges from 0 to 1 as follows:

$$\frac{\log_{10}(C) - \log_{10}(C_{\min})}{\log_{10}(C_{\max}) - \log_{10}(C_{\min})},$$

where C is the measured 5σ contrast for each injected planet, KL mode, annuli, and movement combination, and $C_{\min}$ and $C_{\max}$ are the minimum and maximum 5σ contrast values across all KL mode, annuli, and movement combinations for that planet.

We then convert the false-positive pixel count to a normalized metric. For each pixel exceeding the 5σ threshold inside the control radius for each KL mode, annuli, and movement combination, we (1) subtract its minimum value (generally 0), (2) divide by the difference between its maximum and minimum value across all KL modes, and (3) subtract this value from 1. This creates a metric for which the `movement` and `annuli` parameters with the highest number of false positives are assigned a value of 0 and those with the lowest number are assigned values near the maximum of 1.

A schematic depicting the various stages of the optimization process outlined in this section is shown in Figure 2.

Note that the nature of the 5D `pyKLIP-PE` output parameter space is such that it can be collapsed in many ways to select the "best" choice of `annuli`, `movement`, and `numbasis` parameters for a given data set. In Section 5 we explore several ways in which we chose to collapse these cubes to select a "best" parameter choice, though we note that there are many other possibilities that remain to be explored.

Once optimal parameters have been selected for synthetic planets injected into the continuum images, we apply those parameters to the unaltered (no simulated planets) Hα images to test the ability of this optimization method to robustly recover HD 142527 B at Hα (see Table 1 for details of these observations). At this stage, we apply our optimized parameters to the Hα images only and do not engage in SDI reductions.

4.3. `pyKLIP-PE` Structure

Because the movement parameter is applied at the center of each annular region and the widths of these regions vary with the value of the annuli parameter, certain combinations of annuli and movement lead to equivalent sets of reference images (equivalent rotational masks) and broadly similar reductions. For this reason, parameter exploration outputs frequently show diagonal structure from regions with a large number of annuli and low-movement values to fewer annuli with larger movement values.

A further feature of the parameter explorer heat maps are the jagged structures along the right-hand (high movement value) side. These occur when the synthetic planets are "passed" by the annular zone boundaries and move from being at the inner edge of an annulus whose center is at larger radii to the outer edge of an annulus whose center is at smaller radii. This shift to a new annulus allows for greater movement values to be applied before the algorithm runs out of reference images in the new zone.

Unphysical contrast values are also assigned `NaN` values, leading to a further source of white pixels in each parameter explorer.

4.4. Fixed Parameters

In order to make the computational time for GAPlanetS data sets tractable, we chose to limit the `pyKLIP-PE` grid search to only the `annuli`, `movement`, and `numbasis` KLIP parameters. We did, however, explore in a less systematic way the effect of the high-pass filter (`highpass`) and `IWA` values on postprocessed images and arrived at what we deemed to be reasonable fixed choices for the values of these additional influential KLIP input parameters.

4.4.1. Highpass

In the case of the Fourier high-pass filter width parameter `highpass`, we explored setting it to values of "False," "True" (filter size = image size/10), and the data FWHM for several HD 142527 data sets. We find that the application of an aggressive high-pass filter with a size near the stellar FWHM has a positive effect on the image quality when compared to images with no high-pass filter (`highpass` = False) or a conservative high-pass filter width (`highpass` = True). Example S/N maps showing various high-pass filter widths on the HD 142527 B 2013 April 11 data set are displayed in Figure 3. With no high-pass filtering, the detection has an S/N of 4.1. With high-pass filter set to "True," a width of 1/10 the image size, the S/N improves minimally to a value of 4.5. However, when we set the width to 1 FWHM, the S/N improved by nearly a factor of 2–8σ.

The effects of each filter size on image and detection quality were further explored with receiver operating characteristic (ROC) curves and contrast curves for the HD 142527 data sets. We explored 0.5, 1, or 1.5 times the FWHM but found no significant difference between these filter widths across data sets. Further analyses need to be done to understand the nuanced trade-offs of high-pass values in this regime.

4.4.2. Inner Working Angle

In the case of the IWA parameter `IWA`, MagAO data are noncoronagraphic, so this parameter can in principle be set to zero. The cleared central cavities of GAPlanetS targets are generally $<0\farcs1$ in radius, beneath the IWA of many coronagraphic HCI instruments, and the known companions in their gaps are very tightly separated: smaller IWAs leave more area for planet detection in this region. However, for initial `pyKLIP` parameters `annuli` = 5, `movement` = 2, `KL` = 10, we found that in five out of the six data sets, setting the IWA = FWHM instead of IWA = 0 moderately improved recovered S/N values, with an average improvement of 0.62σ.

We therefore elected to balance these considerations by setting `IWA` to a fixed value of either the median FWHM of the data set or, in the case of saturated data, the saturation radius. We chose these values because we would not be able to separate the light of companions from the starlight regardless of their intrinsic brightness at these separations. The range of IWA values for the data sets considered here spans 3–8 pixels.

We note that, because the KLIP `annuli` parameter divides the space between the IWA and outer working angle (OWA, in

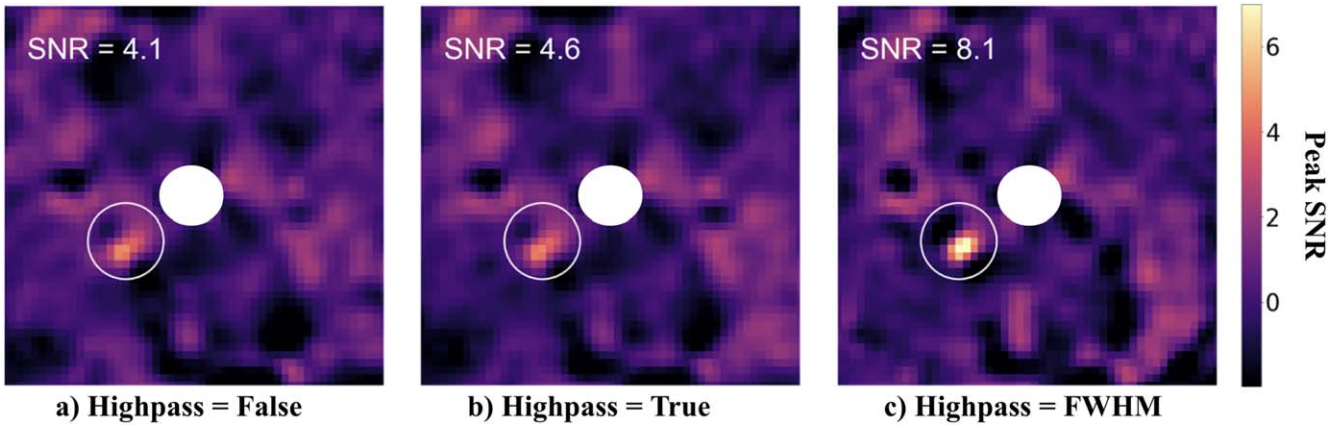

Figure 3. S/N maps showing the effect of high-pass filtering on the HD 142527 2013 April 11 data set. The left S/N map shows no high-pass filtering (`highpass` = False) and has a peak companion S/N of 4.1. The middle S/N map shows a default high-pass filter width of 0.1 × the image size (`highpass` = True) and has a slightly higher peak S/N of 4.6. The right S/N map shows a more aggressive high-pass filter width of 1 × FWHM and has a peak S/N of 8.1.

this case the image boundary) into evenly spaced annular zones, modifying the IWA from data set to data set results in a small variation in the "meaning" of the `annuli` parameter values, in the sense that the same number of annular zones may have slightly different annular widths in pixels if the OWA-IWA range is varied. However, this difference in annuli widths ranges from just 5 pixels for the lowest `annuli` value (1) to under 1 pixel for the highest (25). Similarly, because the `movement` parameter is defined relative to the center of each annulus, this results in a very small variation (<1°) in the size of the rotational mask corresponding to each `movement` value.

Although we chose to fix the `IWA` and `highpass` `pyKLIP` parameters in this investigation, with greater computational resources and/or for smaller single data sets, we strongly recommend optimizing them as well.

5. Optimization Analysis and Results

In an uninformed search for planets, we do not know a priori whether companions are present or at what separations and PAs they might appear in our data sets. For this reason, we have chosen to adopt an approach that averages over a number of planets injected at different separations and PAs in our continuum images, effectively optimizing over a region of interest rather than at a single azimuthal and radial location.

The HD 142527 data sets, with a known companion at a range of separations across the time baseline of our observations, serve as an ideal test case for this general methodology. If the continuum, multiplanet optimized, set of KLIP parameters is robust, then it should successfully recover true Hα companion signals at any location within the optimization region at reasonably high S/N. This S/N may not be the highest achievable, since the companion itself was not the target of the optimization. However, the method does not rely on the reality of an apparent companion signal and therefore is more robust to cognitive biases, particularly in the case of low-S/N detections. It also mitigates the risk of artificially inflating the S/N of a companion signal by optimizing on a single localized combination of companion signal and speckle field. In this section, we report the results of KLIP ADI optimization based on five detection metrics (see Section 4.1 for a description of these metrics) in Section 5.1. We then compare the results of these optimizations to real planet optimization in Section 5.2.

5.1. Generic versus Optimized Parameter S/N Maps

In many direct imaging surveys, KLIP reductions are done with the same set of parameters across all data sets. We mimicked this "survey" strategy by selecting a single, generic set of KLIP parameters and applying them to all HD 142527 data sets for comparison against our optimized reductions. Previous experience with these GAPlanetS data sets led us to choose `annuli` = 5, `movement` = 2, `KL Mode` = 10, which reveals most sources in our data at an S/N greater than 2. KLIP reductions with these initial parameters are shown in the top panel of Figure 4.

We then ran `pyKLIP-PE` on continuum planets injected into our six HD 142527 continuum data sets, with the synthetic planets injected as described in Section 3.4. From the `pyKLIP-PE` output, we first extracted optimum `movement`, `annuli`, and `KL mode` parameters for each of the three main image quality metrics individually, namely, peak S/N, average S/N, and contrast (see Section 4 for a description of how these metrics are computed). These are shown in panels 2, 3, and 4 of Figure 4. We then extracted the optimal parameters from the combination of these three metrics, as shown in panel 5 of Figure 4. Finally, we incorporated the false-positive fraction and neighbor quality metrics described in Section 4 and combined all six image quality metrics, shown in panel 6 of Figure 4. The peak S/N of each reduction is reported as is standard in direct imaging.

Figure 4 is therefore a gallery of KLIP postprocessed Hα images whose parameters were selected to maximize the recovery of continuum injected planets under each metric. Note that the false-positive pixel metric is not shown here individually because it is nearly degenerate and does not have a single optimal value. It is therefore most beneficial when combined with other metrics as a means of excluding bad parameter choices.

Of the 30 optimized companion detections depicted in Figure 4 (six data sets × five metric combinations), 22 produced a peak S/N improvement of up to 1.2σ relative to reductions with generic KLIP parameters (`annuli` = 5, `movement` = 2, `KL Mode` = 10). For five out of the six data

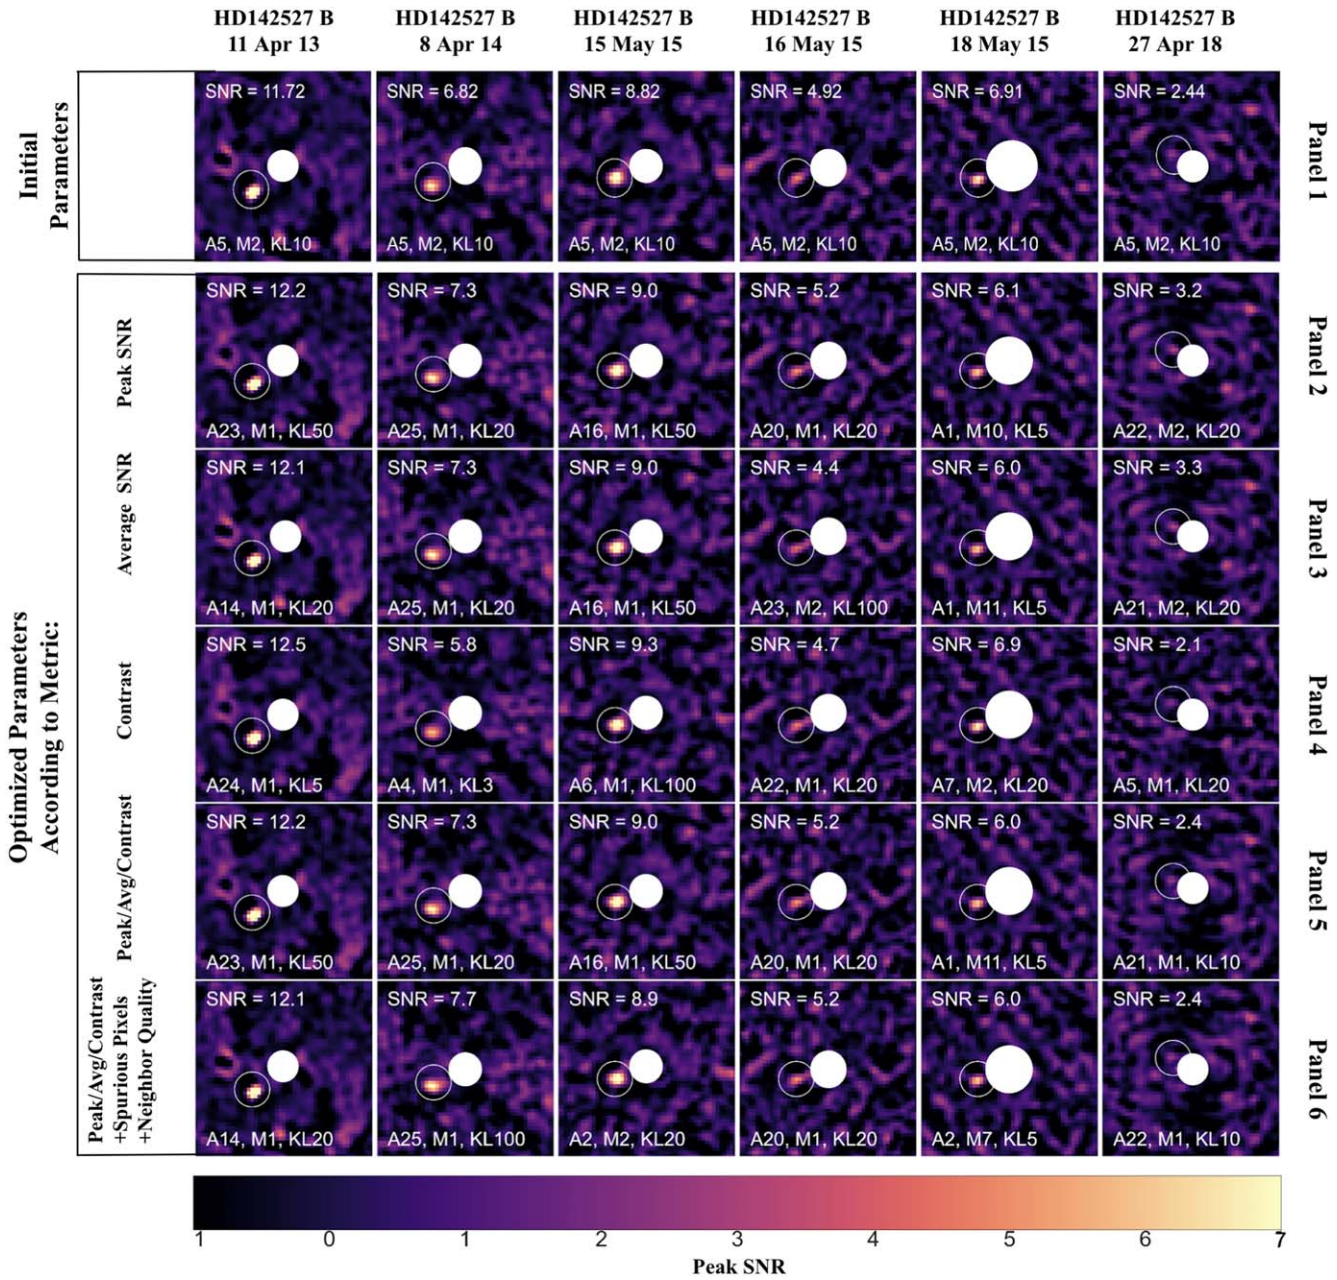

Figure 4. Optimization results for six HD 142527 data sets in which the companion is recovered. The top panel shows S/N maps of KLIP reductions using a single set of fixed KLIP parameters for all data sets: `annuli` = 5, `movement` = 2, `KL` = 10. Panels 2, 3, and 4 show S/N maps of Hα KLIP reductions with optimal parameters identified for injected continuum planets under the `pyKLIP-PE` image quality metrics peak S/N (panel 2), average S/N (panel 3), and contrast (panel 4) *individually*. Panel 5 shows the optimal parameters selected by the *combination* of these three metrics. Panel 6 shows the peak S/N, average S/N, contrast combination, along with the false-positive pixel and neighbor quality metrics. Note that the best parameters based on false-positive pixels are not shown individually because they are nearly degenerate for most data sets. In each of the six epochs, `pyKLIP-PE` was able to recover the real planet at a higher S/N than the generic reductions under each individual metric in 22/30 cases. A Gaussian smoothing was applied to these data with a standard deviation of 1 before computing the S/N maps.

sets, there is at least one metric choice that improved detection S/N over the generic with an average improvement of 0.6. In the remaining data set, the maximum detection S/N selected by the metrics is equivalent to that of the generic parameters. However, even in cases where the optimal false planet parameters for a particular image quality metric lowered the S/N relative to the generic parameter choices, there was still a clear detection of the companion using all metrics. In the case of the most difficult data set (2018 April 27), where the companion is separated by only 5 pixels from the central star, the generic reduction could be classified as a nondetection, while the `pyKLIP-PE` reductions maximizing peak S/N and average S/N reveal the object at an S/N of $\sim 3.2\sigma$. Therefore, we find that in an uninformed or untargeted planet search of the HD 142527 system, `pyKLIP-PE` would recover the companion at an S/N greater than 3 in all epochs using a number of metric combinations.

The peak S/N metric optimization produced the best detections in five out of the six data sets analyzed. However, its optimizations were not significantly better than any other single metric. We also find that for a given data set, multiple metrics tend to converge on the same parameters (e.g., in 2014

April 8 four out of our five metrics were maximized at annuli = 25, movement = 1). In Section 5.2, we explore the effects of combining multiple metrics, as well as the effect of averaging across KL modes.

5.2. Comparison of False and True Companion Optimization

In the previous section, we show a gallery of KLIP Hα reductions with parameters chosen to maximize various combinations of the image quality metrics: peak S/N, average S/N, contrast, false-positive pixels, and neighbor quality. In each case, we ran the pyKLIP-PE optimizer on fake planets injected into the continuum data and applied the selected optimal parameters to Hα images to reveal the real companion. However, in a case such as this—where the existence of the companion is well established—a direct Hα optimization result could in principle be used in lieu of the continuum injected planet output to maximize the recovered signal of the planet. We expect that parameters selected in this manner will better reveal the companion than parameters chosen via fake planet injection. Therefore, an important test of the pyKLIP-PE's efficacy is its level of agreement between fake and real planet optimization.

In order to estimate the level of agreement between injected and true planet optimization, we first reduced the dimensionality of the pyKLIP-PE output. This output consists of 25 movement values, 25 annuli values, 9 KL modes (1, 2, 3, 4, 5, 10, 20, 50, 100), 6 parameter quality metrics (peak S/N, peak S/N neighbor quality, average S/N, average S/N neighbor quality, spurious pixel count, and contrast), and 4–8 injected planets for each data set. We begin by averaging across the continuum planets injected at various PAs and separations between the IWA and control radius (the "optimization region"), which allows for direct comparison with the single companion in Hα. We then collapsed the pyKLIP-PE output for each data set according to every possible combination of image quality metrics and KL modes. The result is an aggregate data quality (ADQ) map for each possible metric collapse scenario. Since these aggregate maps are combinations of the normalized image quality metrics, they have a range of possible maximum values equal to the number of image quality metrics that have been combined. We renormalize them by subtracting the 10th percentile value and dividing by the 90th percentile, creating a parameter quality map where most values lie between 0 and 1. We do not explore weighting of the image quality metrics in this work, and we combine them as a simple linear combination of the individual metrics, but the application of weighting coefficients is likely a fruitful avenue for future exploration.

In this work we use these Hα optimization maps only to test the structural consistency between continuum injected planet ADQ maps under each metric collapse scenario and those of the true Hα companion. To quantify the consistency between the two-parameter quality maps (continuum injected planets and true Hα companion), we subtract the map for the true Hα companion from that of the continuum injected planets. This creates a visualization of the difference in the *structure* of the parameter qualities (see Figure 5). We use summary statistics of these difference maps to analyze the relative merits of different collapse methods and metric combinations, as well as qualitative metrics such as the appearance of the final difference maps in the low-movement regions, where we expect the results to be most stable. This includes searching for the combinations that minimize the differences in the structure of pyKLIP-PE parameter space, which are the values where the true and false companion results are most closely aligned.

One pattern present in the Hα ADQ maps (shown in the left panels of Figure 5) is that they often exhibit two distinct peaks. There is commonly a region of stably optimal parameter space at low annuli and another at high annuli values. Movement space shows fewer clear patterns in the false versus true companion residual maps. However, movement values of zero, equivalent to no rotational mask, are generally unstable in most image metrics, so we have chosen to exclude them in this analysis.

We initially sought to identify the metric or sum of metrics and KL modes that minimized the difference between the real and injected planet optimizations. In principle, a collapse method that mimics the structure of the true companion parameter explorer will have a low standard deviation in the difference map, while maintaining a sum, median, and difference at the continuum peak near zero. Therefore, we utilized distributions of the following values from the difference maps across all possible metric and KL combinations for all six HD 142527 data sets: the median, standard deviation, and sum of the difference maps (1–3), and the value of the difference map at the location of the continuum injected planet peak (4), which quantifies the relative penalty of optimizing on the combination of continuum planets in lieu of the single known companion.

We isolated all metric combinations for which the standard deviation of the false-true difference map was among the lowest *X*%, and where the sum, median, and difference at the peak of that map were among the *X*% closest to the mean, where *X* is a value that we varied to get a sense for the patterns in these parameters under the assumption that, with the small number of data sets under consideration, patterns in "good" metrics would be more informative/universal than the single "best" metric.

The metrics chosen varied, though again there was a preference for more than one image quality metric. We explored parameter combinations within the top 10% and 12% "best" values for all four measures (the median, standard deviation, and sum of the difference maps, and the value of the difference map at the location of the continuum injected planet peak) and found that in the top 10%, 5 KL modes was chosen the most frequently, and in the top 12%, 20 KL modes was chosen then most frequently (Figure 6). Among all metric combinations that fell in the top 12%, in 25 out of 26, at least 1 KL mode 5 or lower and 1 KL model 10 or higher were selected. Therefore, we opted to use both 5 and 20.

Informed by this analysis, we opted for a final "best" collapse method of 5 and 20 KL modes and an equal weighting of all six image quality metrics.

We show the continuum injected planet and true Hα companion aggregate maps for this combination, as well as the difference map and a histogram of its values for this combination of image quality metrics and KL modes for all data sets in Figure 5.

The overall distributions of the standard deviation, sum, median, and difference at the peak of the difference maps for all collapse scenarios with the final choice marked in red are shown in Figure 7.

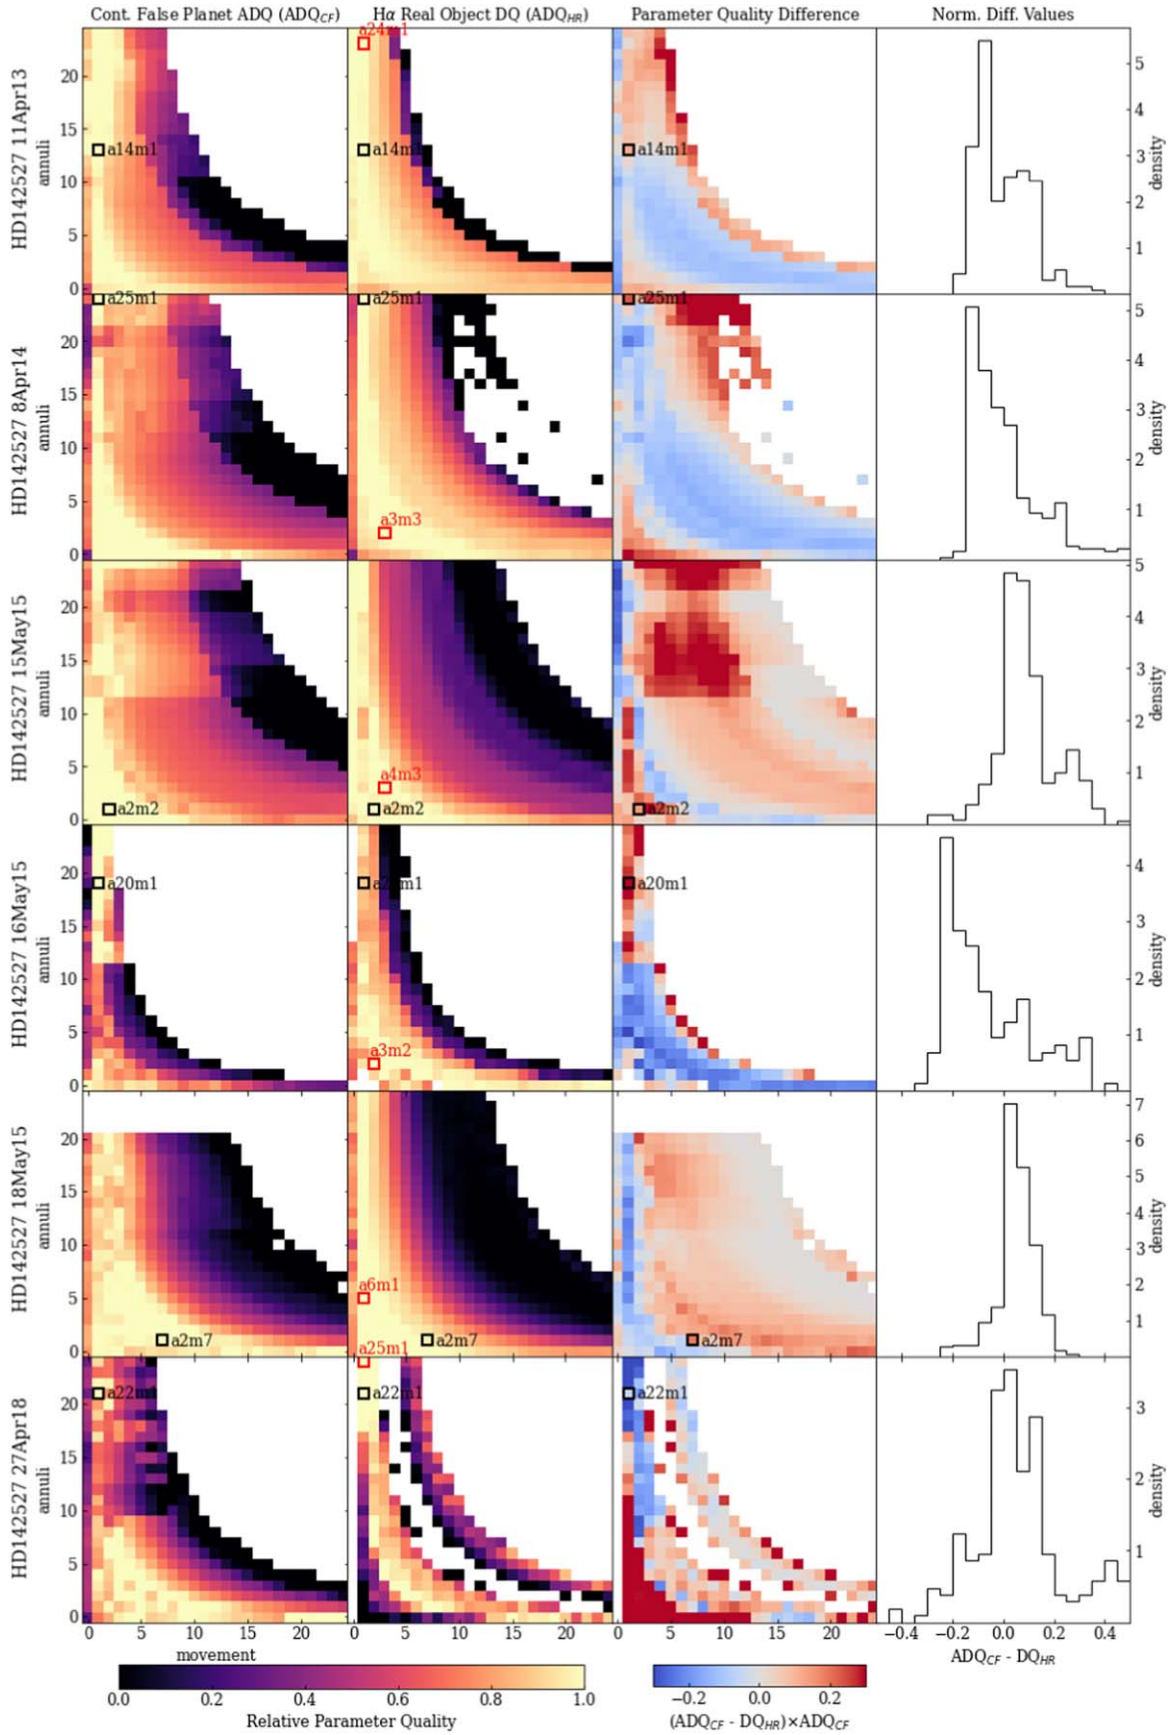

Figure 5. Maps of `pyKLIP-PE` output for continuum injected (CF) planets (first column) and the Hα real (HR) companion (second column), for each HD 142527 B data set, normalized so that the 10th–90th percentiles in the ADQ metric span the range 0–1. The difference between the two normalized maps is shown in the third column, where red values indicate parameter combinations that are more strongly favored for false planets and blue indicates the opposite. Shown in the rightmost column are histograms of the difference maps. The highest relative parameter quality annuli/movement combinations are highlighted in red in the two leftmost columns.

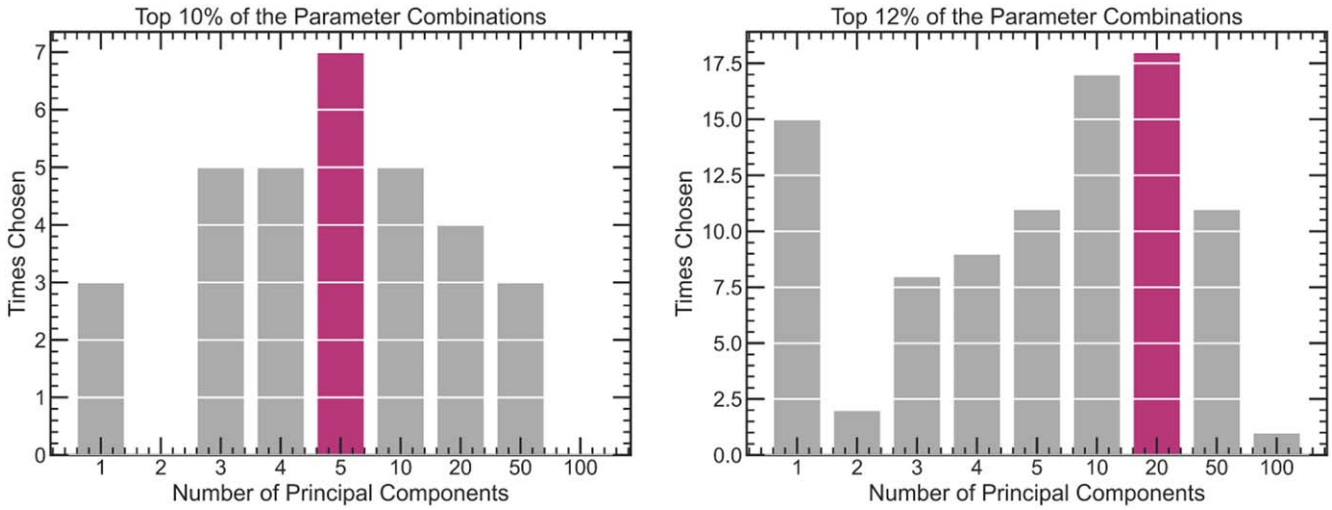

Figure 6. Histograms showing the distribution of KL modes for all parameter combinations (choice of weights for image quality metrics and choice of KL modes for combination) where the four quantitative measures of similarity between parameter explorer heat map structures for true and injected companions (described in detail in the text) were among the top 10% (left) and top 12% (right) of values according to all four measures. In the top 10% of reductions, 5 KL modes was chosen most frequently. In the top 12% of reductions, 20 KL modes was chosen most frequently.

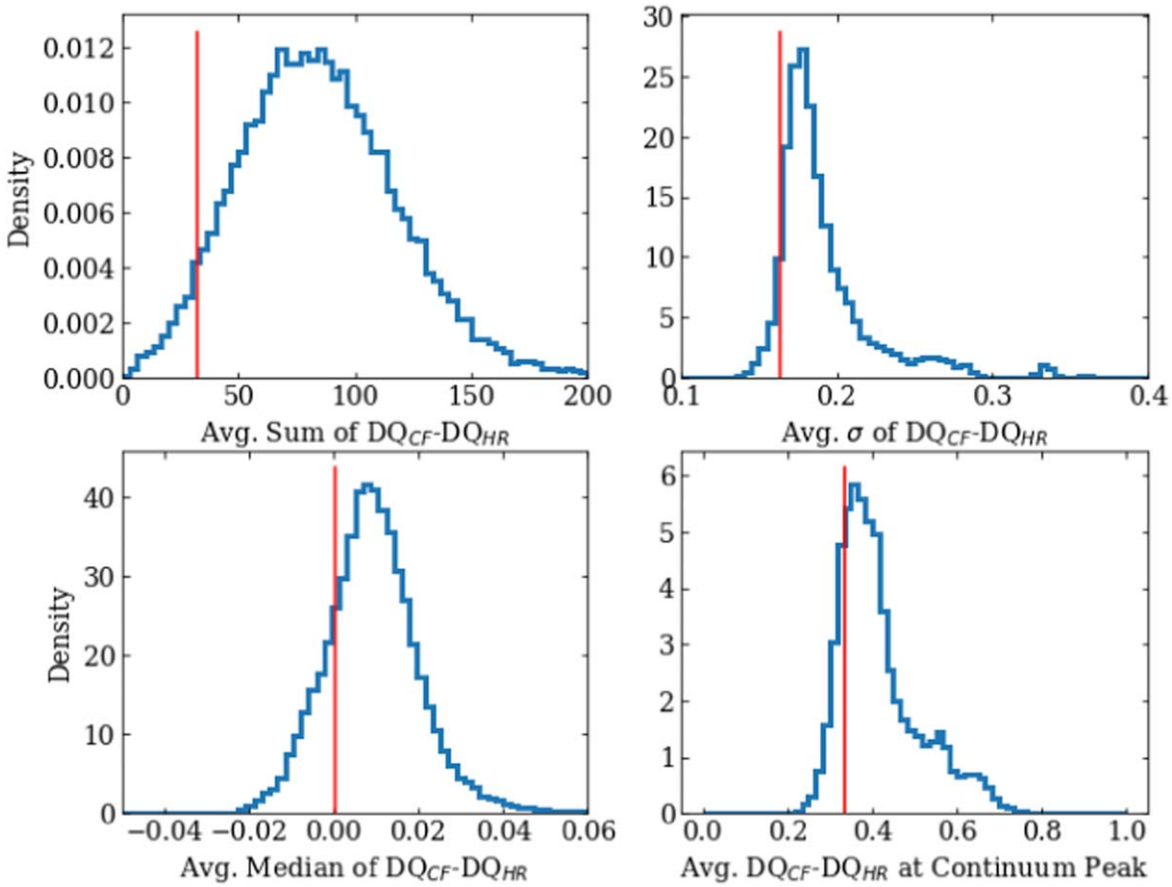

Figure 7. Histograms of summary statistics for continuum injected planet—real Hα companion difference maps. These maps are generated for each of 32,768 possible combinations of image quality metric and KL mode from the pyKLIP-PE output. For each map, the sum of the difference map (top left), standard deviation of the difference map (top right), median of the difference map (bottom left), and difference in the ADQ metric score between the injected and true companions at the location of the peak (bottom right) is computed. These quantities are averaged across the six HD 142527 B detections, and these averaged quantities are depicted in the histograms. The red vertical lines indicate the value of these statistics for our choice of "best" collapse method.

5.3. Computing Efficiency

While a grid search of KLIP parameters applied to injected planets can improve detection quality, we recognize that it may not always be the most time-efficient choice. Note that the `pyKLIP-PE` algorithm takes hours to weeks to run on a single GAPlanetS data set on an 8 core, 32 GB machine. Therefore, next steps in this exploration process will include investigating

more systematically whether there are regions of KLIP parameter space that yield consistently poor detections and can be discarded. This would reduce processing time by decreasing the number of values tested. We caution that these parameters may be peculiar to the nature of the wavelength regime and instruments used for a given data set.

5.4. Future Directions

Some parameters that were not explored in detail in the scope of this research may also be vital in optimizing detections. One example of this is the "high-pass" `pyKLIP` parameter. Preliminary explorations of the high-pass parameter show that it can change S/N by a factor of 3 or more. We find that using high-pass values close to that of the PSF's FWHM usually result in high-quality detections. However, this relationship should be further explored, and perhaps even incorporated into the parameter exploration grid.

These optimization techniques should also be tested on data from other telescopes. Our specific grid optimization technique is certainly biased toward GAPlanetS data, so it is vital to assess to what extent our conclusions are instrument or wavelength dependent.

In order to gain a full understanding of the detectability of accreting protoplanets embedded in disks, future analyses should also consider forward modeling of systems with a disk and planet combination.

6. Summary and Conclusion

In this paper, we demonstrate that a systematic approach to optimization of input parameters to the PSF subtraction algorithm `pyKLIP` results in equivalent or higher-S/N detections of companions in the GAPlanetS sample relative to using a single generic set of parameters across the survey data set. We begin with six data sets of the HD 142527 system taken over the course of 5 yr. We gauge the quality of our parameter selection method, which relies on optimization of planetary signals injected into the continuum images, based on its ability to recover these known companions in as many data sets and at as high an S/N as possible.

We introduce a grid search tool to optimize `pyKLIP` parameters using a number of postprocessed image quality metrics. More specifically, we explore the role of the `pyKLIP` user-input parameters `movement`, `annuli`, and `numbasis` (KL modes) on the quality of postprocessed images. To gauge image quality, we combine a number of metrics computed from the final S/N maps of the postprocessed images. These metrics are the peak (single-pixel) S/N of the recovered planet, the average S/N of the positive pixels under the planet mask ($r \sim 0.5$FWHM), the star/planet contrast achieved at the planet location, the number of false-positive pixels between the IWA and control radius of the S/N map, and the quality of nearby combinations of PyKLIP parameters.

The process of optimization utilized in this work is summarized as follows:

1. Inject two to eight synthetic planets into the *continuum* wavelength images for each GAPlanetS target in between the IWA and control radius.
2. Fix the KLIP parameters `highpass` and `IWA` at values that were found to be universally reasonable (1 FWHM in most cases).
3. Complete a coarse grid search of the `movement`, `annuli`, and `numbasis` parameters with the `pyKLIP-PE` algorithm.
4. Compute S/N maps for each postprocessed image and record various metrics for the "quality" of each injected planet detection, namely, peak S/N, peak S/N neighbor quality, average S/N, average S/N neighbor quality, contrast, and number of $>5\sigma$ false-positive pixels inside of the control radius.
5. Combine all six metrics and average among the 5 and 20 KL mode reductions to arrive at a "best" choice of KLIP parameters for a given set of continuum injected planets.
6. Apply the best injected planet parameters to the Hα data and record the peak S/N of the real companion detections. If the injected planet was able to effectively mimic a real signal, then its optimal parameters should be well approximated for a real planet in the same data set.

We demonstrate that, relative to reductions with a generic fixed choice of KLIP parameter values, this simple grid search technique is able to reveal the HD 142527 B companion in every epoch and improve the S/Ns of the detections by up to 1.2σ.

This simple parameter grid search can help shape our understanding of how to find planets in an uninformed planet search. By showing that synthetic planets injected into images at a wavelength where true sources are expected to be dim can be used as a reasonable proxy for true planets at neighboring wavelengths, we can start to conceive of parameter optimization via more advanced mechanisms (such as neural networks). Furthermore, we have established a reliable and systematic method to select KLIP parameters without relying on the reality of the planetary signal itself.

This paper provides one solution to minimizing false-positive protoplanet detections by developing a robust data-driven method for KLIP parameter optimization that does not rely on the reality of an apparent planetary signal. It introduces a new tool that will help high-contrast imaging surveys make the most of available data, and may even help reveal planets missed in previous explorations.

K.B.F., W.O.B., and J.A. acknowledge funding from NSF-AST-2009816. L.M.C.'s work was supported by NASA Exoplanets Research Program (XRP) grants 80NSSC18K0441 and 80NSSC21K0397. W.O.B. thanks the LSSTC Data Science Fellowship Program, which is funded by LSSTC, NSF Cybertraining grant No. 1829740, the Brinson Foundation, and the Moore Foundation; their participation in the program has benefited this work. K.M.M.'s work has been supported by the NASA XRP by cooperative agreement NNX16AD44G.

This paper includes data gathered with the 6.5 m Magellan Telescopes located at Las Campanas Observatory, Chile.

J.A. would like to thank Sarah Blunt for her help reviewing the extended thesis version of this work and her continued support and mentorship. J.A. especially thanks Sam and Captain Levi for their encouragement writing this paper and their help bringing it to completion. J.A. also thanks Prof. Thayne Currie for his thoughtful and constructive feedback on this paper.

We thank the anonymous reviewer for their considerate and instructive review that greatly improved the quality of this paper.

ORCID iDs

Jéa I. Adams Redai https://orcid.org/0000-0002-4489-3168
Katherine B. Follette https://orcid.org/0000-0002-7821-0695
Jason Wang (王劲飞) https://orcid.org/0000-0003-0774-6502
William Balmer https://orcid.org/0000-0001-6396-8439
Laird M. Close https://orcid.org/0000-0002-2167-8246
Jared R. Males https://orcid.org/0000-0002-2346-3441
Katie M. Morzinski https://orcid.org/0000-0002-1384-0063
Laurent Pueyo https://orcid.org/0000-0003-3818-408X
Helena Treiber https://orcid.org/0000-0003-0660-9776
Kimberly Ward-Duong https://orcid.org/0000-0002-4479-8291

References

Avenhaus, H., Quanz, S. P., Schmid, H. M., et al. 2014, ApJ, 781, 87
Bailey, V. P., Poyneer, L. A., Macintosh, B. A., et al. 2016, Proc. SPIE, 9909, 99090V
Balmer, W. O., Follette, K. B., Close, L. M., et al. 2022, AJ, 164, 29
Benisty, M., Dominik, C., Follette, K., et al. 2022, arXiv:2203.09991
Beuzit, J. L., Vigan, A., Mouillet, D., et al. 2019, A&A, 631, A155
Biller, B., Lacour, S., Juhasz, A., et al. 2012, ApJL, 753, L38
Biller, B. A., & Bonnefoy, M. 2018, Handbook of Exoplanets (Berlin: Springer), 2107
Bowler, B. P. 2016, PASP, 128, 102001
Claudi, R., Maire, A. L., Mesa, D., et al. 2019, A&A, 622, A96
Close, L. M., Hubin, N., Max, C. E., et al. 2008, Proc. SPIE, 7015, 70150Y
Close, L. M., Males, J. R., Morzinski, K., et al. 2013, ApJ, 774, 94
Currie, T., Biller, B., Lagrange, A.-M., et al. 2022, arXiv:2205.05696
Currie, T., Cloutier, R., Brittain, S., et al. 2015, ApJL, 814, L27
Currie, T., Marois, C., Cieza, L., et al. 2019, ApJL, 877, L3
Follette, K. B., Rameau, J., Dong, R., et al. 2017, AJ, 153, 264
Fried, D. L. 1978, JOSA, 68, 1651
Kraus, A. L., & Ireland, M. J. 2012, ApJ, 745, 5
Lafrenière, D., Marois, C., Doyon, R., Nadeau, D., & Artigau, E. 2007, ApJ, 660, 770
Macintosh, B., Graham, J. R., Ingraham, P., et al. 2014, PNAS, 111, 12661
Males, J. R. 2013, PhD thesis, The University of Arizona
Males, J. R., Close, L. M., Morzinski, K. M., et al. 2014, ApJ, 786, 32
Marchetti, K. M., Close, L. M., Males, J. R., et al. 2014, Proc. SPIE, 9148, 914804
Marchetti, K. M., Close, L. M., Males, J. R., et al. 2016, Proc. SPIE, 9909, 990901
Marois, C., Lafrenière, D., Doyon, R., Macintosh, B., & Nadeau, D. 2006, ApJ, 641, 556
Mawet, D., Hirsch, L., Lee, E. J., et al. 2019, AJ, 157, 33
Mawet, D., Milli, J., Wahhaj, Z., et al. 2014, ApJ, 792, 97
Meshkat, T., Kenworthy, M. A., Quanz, S. P., et al. 2014, ApJ, 780, 17
Pueyo, L. 2016, ApJ, 824, 117
Quanz, S. P., Amara, A., Meyer, M. R., et al. 2013, ApJL, 766, L1
Rameau, J., Follette, K. B., Pueyo, L., et al. 2017, AJ, 153, 244
Ruffio, J.-B., Macintosh, B., Wang, J. J., et al. 2017, ApJ, 842, 14
Sallum, S., Follette, K. B., Eisner, J. A., et al. 2015, Natur, 527, 342
Seager, S., & Deming, D. 2010, ARA&A, 48, 631
Soummer, R., Pueyo, L., & Larkin, J. 2012, ApJL, 755, L28
Thalmann, C., Janson, M., Garufi, A., et al. 2016, ApJL, 828, L17
Thompson, W., & Marois, C. 2021, AJ, 161, 236
Wagner, K., Follette, K. B., Close, L. M., et al. 2018, ApJL, 863, L8
Wang, J. J., Ruffio, J.-B., De Rosa, R. J., et al. 2015, pyKLIP: PSF Subtraction for Exoplanets and Disks, Astrophysics Source Code Library, ascl:1506.001