NRC·CNRC

NRC Publications Archive
Archives des publications du CNRC

Deep Orbital search for additional planets in the HR 8799 system

Thompson, William; Marois, Christian; Do Ó, Clarissa R.; Konopacky, Quinn; Ruffio, Jean-Baptiste; Wang, Jason; Skemer, Andy J.; De Rosa, Robert J.; Macintosh, Bruce

This publication could be one of several versions: author's original, accepted manuscript or the publisher's version. / La version de cette publication peut être l'une des suivantes : la version prépublication de l'auteur, la version acceptée du manuscrit ou la version de l'éditeur.
For the publisher's version, please access the DOI link below./ Pour consulter la version de l'éditeur, utilisez le lien DOI ci-dessous.

Publisher's version / Version de l'éditeur:
https://doi.org/10.3847/1538-3881/aca1af
The Astronomical Journal, 165, 1, pp. 1-20, 2022-12-23

NRC Publications Archive Record / Notice des Archives des publications du CNRC :
https://nrc-publications.canada.ca/eng/view/object/?id=77a2170a-b6b5-4b51-abcb-648fefb0f0d2
https://publications-cnrc.canada.ca/fra/voir/objet/?id=77a2170a-b6b5-4b51-abcb-648fefb0f0d2

Access and use of this website and the material on it are subject to the Terms and Conditions set forth at
https://nrc-publications.canada.ca/eng/copyright
READ THESE TERMS AND CONDITIONS CAREFULLY BEFORE USING THIS WEBSITE.

L'accès à ce site Web et l'utilisation de son contenu sont assujettis aux conditions présentées dans le site
https://publications-cnrc.canada.ca/fra/droits
LISEZ CES CONDITIONS ATTENTIVEMENT AVANT D'UTILISER CE SITE WEB.

Questions? Contact the NRC Publications Archive team at PublicationsArchive-ArchivesPublications@nrc-cnrc.gc.ca. If you wish to email the authors directly, please see the first page of the publication for their contact information.

Vous avez des questions? Nous pouvons vous aider. Pour communiquer directement avec un auteur, consultez la première page de la revue dans laquelle son article a été publié afin de trouver ses coordonnées. Si vous n'arrivez pas à les repérer, communiquez avec nous à PublicationsArchive-ArchivesPublications@nrc-cnrc.gc.ca.

National Research Council Canada | Conseil national de recherches Canada

Canada

The Astronomical Journal, 165:29 (20pp), 2023 January https://doi.org/10.3847/1538-3881/aca1af
© 2022. The Author(s). Published by the American Astronomical Society.
OPEN ACCESS

Deep Orbital Search for Additional Planets in the HR 8799 System

William Thompson[1], Christian Marois[2], Clarissa R. Do Ó[3], Quinn Konopacky[3], Jean-Baptiste Ruffio[4], Jason Wang[5], Andy J. Skemer[6], Robert J. De Rosa[7], and Bruce Macintosh[8]
[1] University of Victoria, Department of Physics and Astronomy, 3800 Finnerty Rd, Victoria, BC V8P 5C2, Canada; wthompson@uvic.ca
[2] National Research Council of Canada Herzberg, 5071 West Saanich Rd, Victoria, BC, V9E 2E7, Canada
[3] Center for Astrophysics and Space Sciences, University of California, San Diego, La Jolla, CA 92093, USA
[4] Department of Astronomy, California Institute of Technology, Pasadena, CA 91125, USA
[5] Department of Physics and Astronomy, Northwestern University, Evanston, IL 60208, USA
[6] Department of Astronomy and Astrophysics, University of California, Santa Cruz, Santa Cruz, CA 95064, USA
[7] European Southern Observatory, Alonso de Córdova 3107, Vitacura, Santiago, Chile
[8] Kavli Institute for Particle Astrophysics and Cosmology, Stanford University, Stanford, CA 94305, USA
Received 2022 July 15; revised 2022 October 14; accepted 2022 October 23; published 2022 December 23

Abstract

The HR 8799 system hosts four massive planets orbiting 15 and 80 au. Studies of the system's orbital stability and its outer debris disk open the possibility of additional planets, both interior to and exterior to the known system. Reaching a sufficient sensitivity to search for interior planets is very challenging due to the combination of bright quasi-static speckle noise close to the stellar diffraction core and relatively fast orbital motion. In this work, we present a deep *L*-band imaging campaign using NIRC2 at Keck comprising 14 observing sequences. We further re-reduce archival data for a total of 16.75 hr, one of the largest uniform data sets of a single direct imaging target. Using a Bayesian modeling technique for detecting planets in images while compensating for plausible orbital motion, we then present deep limits on the existence of additional planets in the HR 8799 system. The final combination shows a tentative candidate, consistent with 4–7 $M_{\rm jup}$ at 4–5 au, detected with an equivalent false-alarm probability better than 3σ. This analysis technique is widely applicable to archival data and to new observations from upcoming missions that revisit targets at multiple epochs.

Unified Astronomy Thesaurus concepts: Exoplanets (498); Direct imaging (387)

1. Introduction

The HR 8799 planetary system hosts four giant planets still glowing hot from their recent formation. It was the first multiplanetary system to be directly imaged (Marois et al. 2008b) in 2008. Since the subsequent detection of a fourth inner planet in 2010 (Marois et al. 2010), it became the benchmark system in direct imaging. Extensive follow-up observations of the four planets b, c, d, and e have characterized their orbits, compositions, (Currie et al. 2011; Skemer et al. 2012; Konopacky et al. 2013; Skemer et al. 2014; Currie et al. 2014; Barman et al. 2015; Wertz et al. 2017; Greenbaum et al. 2018; GRAVITY Collaboration et al. 2019; Wang et al. 2020; Wahhaj et al. 2021; Wang et al. 2021; Sepulveda & Bowler 2022) and now masses (Brandt et al. 2021).

Precision astrometric monitoring, orbit fitting, and dynamical modeling have found that the planets follow nearly coplanar orbits with relatively low eccentricity (0 to around 0.25) (Soummer et al. 2011; Bergfors et al. 2011; Currie et al. 2012; Sudol & Haghighipour 2012; Pueyo et al. 2015; Maire et al. 2015; Konopacky et al. 2016; Zurlo et al. 2016; Wertz et al. 2017; Wang et al. 2018; GRAVITY Collaboration et al. 2019). The orbits appear to form a near-resonant chain, with factor of two period multiples 1b:2 c:4d:8e (e.g., Maire et al. 2015; Wang et al. 2018; Goździewski & Migaszewski 2020). This is fortunate, because dynamical modeling shows that few configurations exist that are stable over millions of years besides these resonant chains.

Original content from this work may be used under the terms of the Creative Commons Attribution 4.0 licence. Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

These same orbital models show that the system's stability is increasingly tenuous if the inner three planets have masses much above 8 $M_{\rm jup}$, which matches the masses of approximately 7 $M_{\rm jup}$ derived through bolometric luminosity and evolution models. Recently, Brandt et al. (2021) combined previous sets of stable orbits from Wang et al. (2018) with careful modeling of the Hipparcos-GAIA proper motion anomaly to estimate the masses of c, d, and e as $9.6^{+1.9}_{-1.8}$ $M_{\rm jup}$. This dynamical mass measurement is in slight tension with the results from orbital stability and atmosphere modeling; however, any additional planets in the system would impact this dynamical mass measurement. In their orbit modeling, Goździewski & Migaszewski (2014) found some stable configurations that extend the resonant chain down to a fifth inner planet of up to roughly 6 $M_{\rm jup}$ near 7.5 au (1e:3f) or 9.5 au (1e:2f), and a larger "generally stable" region for test particles below ∼6–7 au. Brandt et al. do consider such a fifth planet, and place 3σ detectable mass limits at roughly 5.5 $M_{\rm jup}$ between 3 and 5 au, 6 $M_{\rm jup}$ between 5 and 7 au, and 7.5 $M_{\rm jup}$ near 9 au, but these limits do not consider how the space of stable orbits used to fit the mass may change by adding a fifth planet.

Many groups have undertaken extensive direct imaging searches for additional planets in the system (e.g., Currie et al. 2014; Maire et al. 2015; Wahhaj et al. 2021). The most sensitive constraints on the mass of an additional inner planet come from Wahhaj et al. (2021) in the YJH (IFS) and K bands (IRDIS). Using BT-Setl models (Spiegel & Burrows 2012) and assuming an age of 30 Myr, they place 5σ upper limits of 3.6 $M_{\rm jup}$ at 7.5 and 2.8 $M_{\rm jup}$ at 9.7 au.

These limits still leave room in the semimajor axis–mass parameter space where a fifth inner planet could hide; however, there are significant challenges to further improving our sensitivities. Current observations at such separations are

limited by quasi-static speckles (Marois et al. 2003). These speckles produce a non-Gaussian noise distribution that is highly correlated over time, and sensitivity improves poorly with increasing integration time. Observations thus use angular differential imaging (ADI; Marois et al. 2006a), at times in combination with spectral differential imaging (SDI; Racine et al. 1999; Walker et al. 1999; Marois et al. 2000) and reference star differential imaging (RDI; Wahhaj et al. 2021). These greatly improve sensitivity, but require observations to be scheduled near when the system transits the meridian or have a suitable reference star nearby. Compounding this issue is orbital motion. Planets at smaller semimajor axes have much shorter orbital periods according to Kepler's third law. At a separation of ∼5 au, a planet in this system would move fast enough that observations taken more than a few months apart would start to blur the planet. All told, this means that considerable integration time is required and that it is challenging to schedule time within the few-month window necessary to freeze orbital motion.

Combining images in the presence of orbital motion was previously considered in Males et al. (2013) and Nowak et al. (2018). These approaches consider "de-orbiting," in which images are transformed and stacked to counteract orbital motion. In doing so, they find that they are able to increase the S/N of a faint candidate despite orbital motion. One challenge with these approaches is that it becomes difficult to quantify the significance of such a detection in a way that includes uncertainty in the candidate's orbit. A second challenge is that the flux of a candidate can vary freely between epochs even if it is not consistent with later data. This latter point may be why Males et al. (2013) find an increasing false-positive rate with increasing orbital motion.

Separately, a large debris disk first noted in the star's spectral energy distribution (SED; Sadakane & Nishida 1986; Zuckerman & Song 2004; Rhee et al. 2007) and then described by Su et al. (2009) lies beyond the known planets and extends outward to perhaps as far as 1000 au (Matthews et al. 2014). The disk is only marginally resolved and models do not yet constrain the inner edge, with estimates varying from 104^{+8}_{-12} au (Wilner et al. 2018) to 145 au (Booth et al. 2016) or 170 ± 40 au (Faramaz et al. 2021). These works and the additional dynamical studies of Goździewski & Migaszewski (2018) and Geiler et al. (2019) consider several scenarios for what mechanism may have sculpted the inner edge of the disk, one of which is an additional outer planet between 0.1 and 3 $M_{\rm jup}$. According to Faramaz et al. (2021) the best limit on an additional outer planet in this regime is 1.25 $M_{\rm jup}$ by Maire et al. (2015); however, the contrast curves presented in that work end at 70 au of projected separation. Further out, Close & Males (2009) set a lower limit of ∼3 $M_{\rm jup}$ between ∼200 and 600 au, yet it appears no lower limits have been published on additional outer planets between b and the start of the outer debris disk, or ∼100–150 au. Though it will not access the inner region of HR 8799 with standard coronagraphic imaging, JWST is poised to place exquisite constraints on outer planets in this regime.

To search for these proposed additional inner and outer planets, we performed an extensive L' imaging campaign at Keck with NIRC2, reprocessed archival NIRC2 data using direct S/N optimization, and we used a joint Bayesian model of planet orbits and photometry to search for planets despite significant orbital motion. We present limits on the existence of any additional planets as well as a modest-S/N candidate at approximately 4–5 au worthy of further study.

2. Observations and Processing

2.1. Observations

To search for additional planets in the system, we conducted a campaign of deep L' imaging at Keck using the NIRC2 instrument (PI: K. Matthews) in 2017, 2020, and 2021. Our observations for this campaign totaled 14 quarter nights. We took observations in pupil tracking mode so that the field of view rotated during each sequence, but the speckle pattern remained fixed. This allowed us to process the data with angular differential imaging (ADI; Marois et al. 2006a) to suppress the halo of quasi-static speckles (Marois et al. 2003). Each observation was scheduled such that HR 8799 would transit the meridian roughly halfway through the sequence.

We chose L' imaging because it balances the favorable contrast of young planets at longer wavelengths with the need to access tight inner working angles, as well as limiting noise from the thermal background.

The observing strategy evolved over the course of the campaign. The 2017 epoch was captured with the 400 mas diameter Lyot coronagraph, differential atmospheric refraction (DAR) correction set to acquisition and track, and included dithering away from transit to improve background subtraction. For the 2020 epoch, we observed without a coronagraph and with limited dithering. This is because initial reductions showed that the quasi-static speckles near the edge of the mask were not as stable as those in non-coronagraphic data sets and because we found that the speckle pattern was not stable between dither positions. For example, see the ∼3.5× improvement in contrast between the otherwise very similar 2017 coronagraphic and 2020/2021 non-coronagraphic epochs near 200–300 mas separation (Figure 1 and Table 1). We did not switch to using the new optical vortex (PI: K. Matthews; Serabyn et al. 2017) installed in 2015, as it was available only in shared risk mode. Finally, for the 2021 data set, we also chose to set DAR (differential atmospheric refraction) to acquisition only and not acquisition and track to reduce the number of optics moving during our observations. Our 2021 data set shown in Figure 2 achieved the deepest contrast close to the star. We captured background images at the start and/or end of each sequence in order to reduce the thermal background. The exposure time of each sequence was adjusted to avoid saturating the first Airy ring of the stellar PSF. Individual exposures were co-added by the detector to create 30–80 s exposure images depending on conditions. Higher-cadence observations better capture the moment-to-moment variation in the stellar PSF; however, they come at the expense of considerable dead time after each image. For each sequence, we captured unsaturated non-coronagraphic images at the beginning and ends of each NIRC2 sequence to use as planet PSF templates and contrast calibrations.

2.2. Archival Data Selection

To this campaign, we added additional data from the Keck archive dating back to some of the first sequences taken of the system. We considered L′ sequences captured by members of our collaboration. We hoped this data would increase our sensitivity, in particular to a planet whose orbit might have appeared closer to the star during our main campaign. Combined with our dedicated observing campaign, this brought

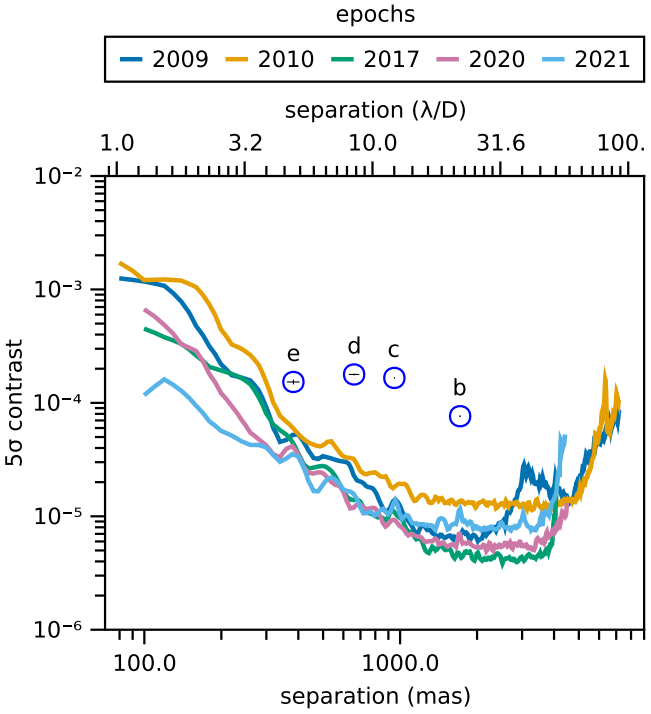

Figure 1. 5σ contrast limits at each of the five L' epochs on a log-log scale. Blue circles mark the average separations and photometry from our results. The contrast achieved close to the star has improved over time, so later epochs have more weight in our modeling.

the total integration time, not including overheads, bad frames, sky backgrounds, and other calibrations, to 16.75 hr of on-source data at L' . The sequences taken in 2009 include artifacts at wide separations due to nodding that we exclude from our models of the outer system.

Besides this data, the Keck archive contains on the order of a further 8 hr of L-band observations from other researchers. We did not include these sequences, due to the large manual effort required to reduce one-off observations captured with varying observing strategies—and in some cases, unsuitable choices of focal plane mask that obscure the inner planets. While not complete, our sample contains on the order of 65% of all L' data that have been recorded of HR 8799 by NIRC2. If we assume that the S/N of the combined observations grows with the square root of the total exposure time, then reducing all remaining observations could in theory increase the S/N by up to 25%.

Future work could additionally combine data from other wavelengths and observatories. For instance, the Large Binocular Telescope's (LBT) LMIRCam has similar capabilities to NIRC2 in the L band. Observations at other wavelengths, e.g., by GPI and SPHERE could also be added; however, using these observations would require us to assume a priori the color of the planets we search for.

For a full listing of the sequences used in this work, see Table 1.

2.3. ADI Reduction

Angular differential imaging (ADI) is a powerful technique for suppressing quasi-static speckles, but least-squares based algorithms (Lafrenière et al. 2007; Soummer et al. 2012; Marois et al. 2014) suffer from worsening contrast and planet self-subtraction at very small separations. We therefore developed a new technique for reducing differential imaging

Table 1
Observations

Date (UT)	Mask	Integration (min)	FoV Rotation (°)	Seeing: DIMM, MASS, WRF (")
2009-08-01	none	53	167	NA
2009-10-31	none	34	158	0.6, 0.2, 0.8
2009-11-01	none	72	162	0.8, 0.2, 1.4
2010-07-21	corona400	26	156	0.8–0.5—NA
2017-07-07	corona400	37	172	0.7, NA, 0.4
2017-07-11	corona400	57	179	NA, NA, 0.4
2017-07-12	corona400	90	179	NA, NA, 0.4
2017-07-13	corona400	73	178	0.5, 0.4, NA
2017-07-14	corona400	68	178	0.6, 0.4, NA
2020-08-23	none	62	169	0.6, 0.5, NA
2020-08-25	none	54	165	1.8, NA, NA
2020-08-27	none	43	170	0.8, 0.2, NA
2020-10-07	none	60	170	0.4, 0.2, NA
2020-11-17	none	53	178	0.8, 0.5, NA
2021-07-08	none	52	170	0.5, 0.2, NA
2021-07-09	none	65	173	0.5, 0.15, NA
2021-07-10	none	49	173	0.4, 0.2, NA
2021-07-11	none	57	172	0.7, 0.7, NA
Total 12 yr baseline		1005	3069	

Note. Observations grouped by year. The integration column gives the total science exposure time not including calibrations and overheads. Seeing information is summarized from the Maunakea Weather Center Archive where available. All sequences were captured with the L' filter using NIRC2 in narrow mode.

sequences called direct S/N optimization (Thompson & Marois 2021). This technique offers improved contrast close to stars by solving a system of quadratic equations maximizing S/N, rather than linear equations minimizing noise. Additionally, it optimizes stacks of multiple images in a sequence simultaneously to reduce correlated residual noise. See the above reverence for more information and a comparison of direct S/N optimization to LOCI on one of the 2020 sequences presented in this paper.

We used the Signal-to-Noise Analysis Pipeline (SNAP) to calibrate, align, and reduce each sequence. We pre-processed all images by subtracting darks, flat-fielding with Ks-band dome flats, and subtracting sky backgrounds where available. We applied distortion corrections using the solutions of Yelda et al. (2010) and Service et al. (2016) for data taken after the 2015 NIRC2 servicing. We then high-pass filtered the images to further suppress thermal background noise using a 25 pixel standard deviation. This step has the side effect of suppressing any diffuse emission in the system, though we do not expect to detect any emission from the inner or outer debris disks in these observations. We aligned the NIRC2 images using an iterative cross-correlation procedure because NIRC2 does not possess "satellite spots," off-axis faint copies of the stellar PSF used for astrometric and photometric calibration (Marois et al. 2006b; Sivaramakrishnan & Oppenheimer 2006). We aligned each image against the unsaturated PSF template and then stacked to create a master. We then cross-correlated each image against this master to improve their alignment, stacked them to create a new master, and then repeated the procedure a further two

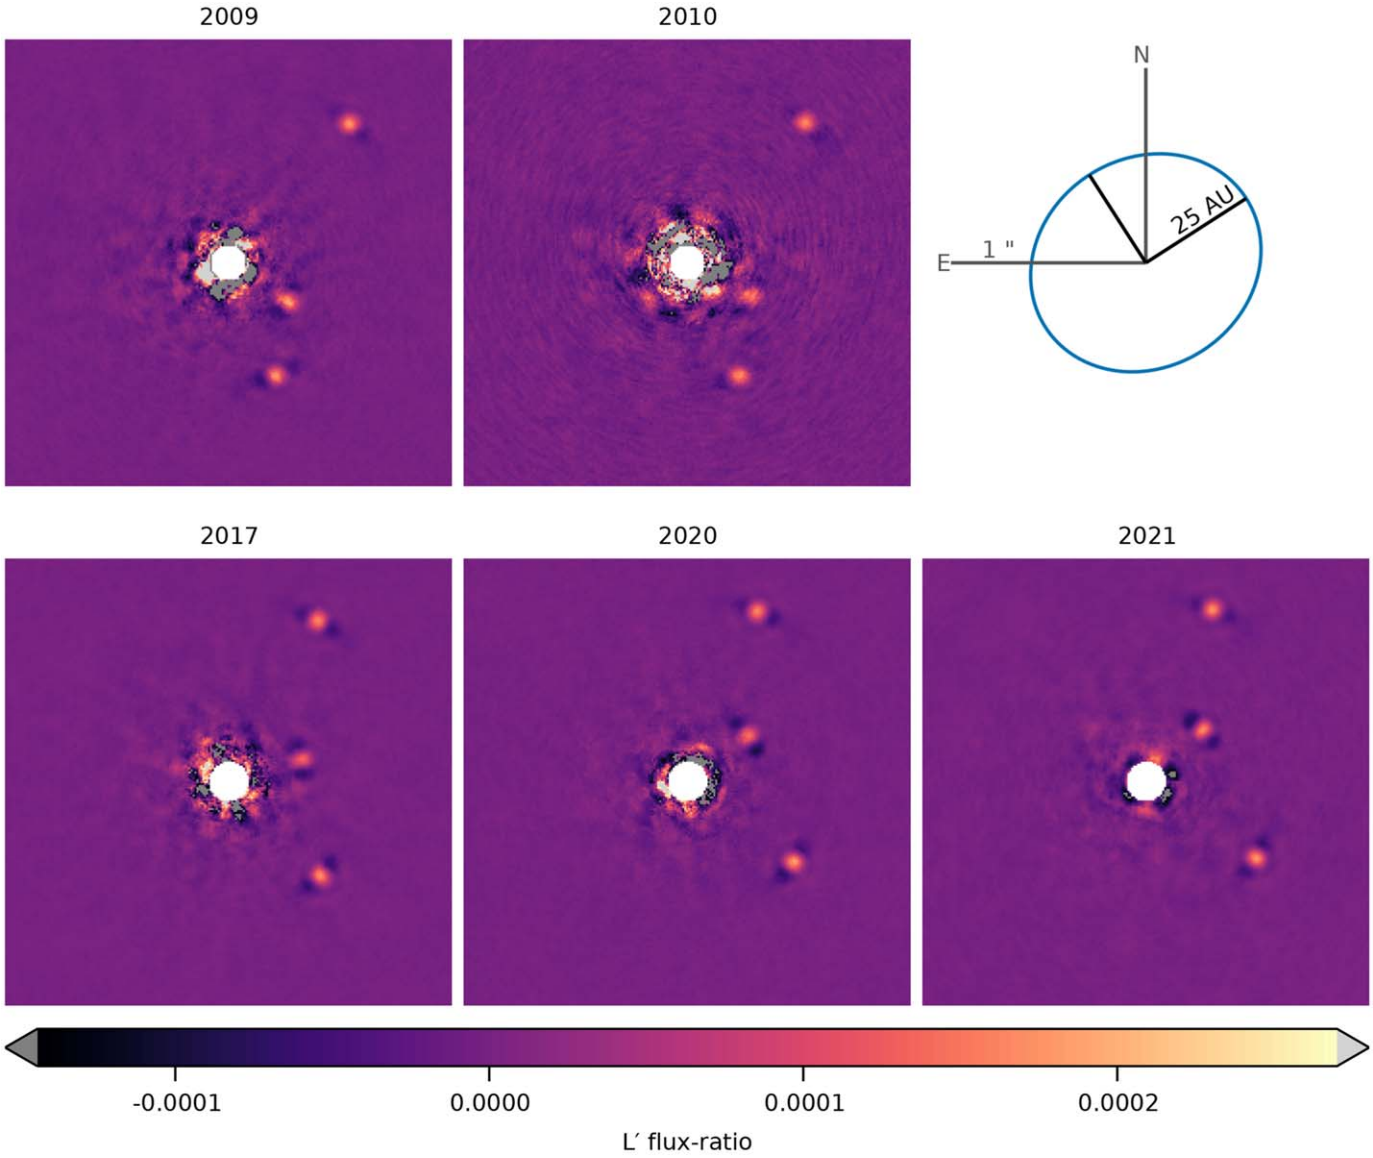

Figure 2. Combined images from each epoch, processed with SNAP and cropped to show the inner planets. The images are shown on a linear scale with correct throughput, normalized using the flux of planet d. Values outside the linear colorscale are shown in gray. Note that, when using SNAP, the forward-processed images should not be used for evaluating the noise. Close to the star, the PSFs of any potential planet can produce secondary positive and negative ghosts at other position angles. Instead, we calculate contrast curves using matching backward-rotated reductions. The top panels are comprised of archival data, while the bottom panels are the new campaign. Significant orbital motion is visible between epochs. The full images including planet b are listed in Appendix.

times. Finally, we rejected bad frames using a correlation threshold (usually 1–5 images per sequence).

We reduced the data using SNAP multitarget S/N optimization with batches of 10 images, and an optimized number of included reference images for each subtraction region. No parameters of the reduction were changed between sequences, to prevent human bias. The S/N optimization procedure does not include a rejection distance/exclusion zone or other adjustable aggressiveness parameter; all images are used in the optimization, including those in which the planet PSF overlaps.

As will be described in the following section, our models assume that the input images have well-calibrated planet throughput. The SNAP pipeline is throughput preserving for point sources as long as the instrument's PSF does not deviate significantly from the unsaturated templates captured before or after the sequence. The S/N optimization algorithm does not produce self-subtraction for the peaks of point sources, and oversubtraction is prevented by the use of optimization, buffer, and subtraction regions.

For each sequence, we also produced a matching, "backward-rotated" image in which the rotation direction of the ADI sequence was reversed. These backward-rotated images have the same noise distribution as the normal images, but do not contain any significant signal from the planets. This allows us to calculate a contrast curve for each sequence that is unbiased by planets. We verified that contrast curves of the backward-rotated sequences matched the regular sequences between planets b, c, d, and e, where any additional planets are very unlikely to orbit.

Finally, we combined the processed results of each sequence using a contrast-weighted median stack of images taken less than 3 months apart.[9] During this period, any planet on a

[9] SNAP processes groups of 10 frames into individual reduced images to reduce correlated noise. We stacked these individual processed images across nights within the three-month window.

Table 2
S/N Ratios in Combined Images

Year	b	c	d	e
2009	50	64	33	21
2010	28	48	27	13
2017	92	80	62	22
2020	56	90	75	27
2021	36	62	72	28

circular orbit with semimajor axis greater than or equal to 5 au would move less than $\frac{1}{3}\lambda/D$. Conveniently, this resulted in one combined image for each year HR 8799 was observed as shown in Figure 2 and Figure 16. The S/Ns of planets b, c, d, and e in each of these combined images are presented in Table 2.

2.4. Photometric Calibration

Because NIRC2 does not have satellite spots for astrometric and photometric calibration, we reduced the data in units of contrast relative to the star, as measured by unsaturated images taken before and/or after the sequence. However, variable weather conditions and AO performance during a sequence between the saturated and unsaturated images led to an additional photometric error on the order of 10% between sequences, which is greater than what we would expect from random variation at the recovered S/N. To compensate, we measured the flux of planets b, c, d, and e in each epoch and rescaled the images to the S/N-weighted average flux. This correction improved the photometric variation between epochs to the expected level, e.g., a 2% variation for planet c, which is close to what we would expect for a planet with S/N ratios of 60–90.

To calculate mass, we convert from convert from contrast to absolute magnitude using an apparent magnitude of 5.3 for HR 8799 in the L' filter (Marois et al. 2010).

3. Modeling

During the 12 yr baseline of our data set, any planet with a semimajor axis less than ≈3500 au would move more than $1\lambda/D$. This means that the data cannot be naïvely stacked to improve our sensitivity, as the signal of a planet would not be aligned between epochs. Because we wish to consider a very large parameter space of inclined and eccentric orbits, simple approaches like rotating and scaling the images would not be effective.

We therefore create a probabilistic model of the system and jointly model the orbits and photometry of the planets. This combines the process of detecting candidates with orbit fitting (e.g., Blunt et al. 2020) without the intervening step of extracting candidate astrometry at each epoch. Because astrometry is extracted from images, it follows that, if we can model images directly, then extracting astrometry and photometry as intermediate products is not necessary.

Crucially, this allows us to detect the signal of a planet that is too faint to see in a single epoch even if it has moved considerably between images. In fact, planets with arbitrarily low S/Ns per image can grow to detectable levels given a sufficient number of epochs. The sensitivity and limitations of this method will be expanded upon in a dedicated publication.

We consider nine parameters:

1. M, the total mass of the system;
2. Π, the parallax of the system;
3. a, the semimajor axis;
4. e, the eccentricity;
5. i, the orbital inclination;
6. Ω, the longitude of ascending node;
7. ω, the argument of periapsis;
8. τ, the time of periastron passage following the convention of Orbitize! (Blunt et al. 2020);
9. and L, the L' flux ratio.

The physical, geometric, and orbital parameters define a unique Keplerian orbit through our images and a position at each epoch. We assume that the planet has the same photometry over all observations. Following Ruffio et al. (2018), we consider the log-likelihood of a planet having those parameters as

$$\log \mathcal{L} \propto \sum_i \frac{1}{2\sigma_{x,i}^2}(L^2 - 2LL_{x,i}), \quad (1)$$

where i is the epoch, x is the computed position at epoch i, $L_{x,i}$ is the measured photometry extracted at position x from the image i, and $\sigma_{x,i}^2$ is the variance in that photometry. The schematic in Figure 3 illustrates this procedure.

To extract the photometry $L_{x,i}$ efficiently from the images, we perform a noise-weighted convolution by an Airy disk of $1\lambda/D$. Weighing the convolution locally at each pixel by the surrounding contrast is essential so that the peak S/N occurs at the location of the planet, rather than offset in the direction of lower noise. In regions with strongly sloped contrast curves (e.g., planet e), this correction prevents a position bias of up to 30 mas. We then look up the photometry at each coordinate using a bilinear interpolation. We estimate the variance $\sigma_{x,i}^2$ using the contrast at that separation in each image. We measured the contrast curves themselves using matching backward-rotated ADI reductions so that the signals of any planets do not bias the contrast.

This likelihood function assumes that our convolved images are maximum likelihood estimates of planet photometry in the presence of Gaussian noise, and that contrast curves extracted from the backward-rotated noise maps provide good estimates of the variance in that estimate at each pixel.[10] Ruffio et al. (2018) provide a derivation of this likelihood function and show how non-Gaussianity does not significantly affect Bayesian upper limits by that definition. For detection thresholds, on the other hand, Section 4.5 of this paper discusses how we correct for mildly non-Gaussian noise.

This approach is similar to that of Mawet et al. (2019), in which direct images of ϵ Eridani are combined with radial velocity data; however, we do not include radial velocities (the planets of HR 8799 orbit almost face-on from our perspective) but instead combine images from multiple epochs as in Skemer & Close (2011). It also shares some similarities in concept with those of K-stacker (Le Coroller et al. 2015; Nowak et al. 2018) and for example the recent results of Le Coroller (2022).

Note also that we do not use a PCA (Soummer et al. 2012) or matched filter (Ruffio et al. 2017) based data processing as in Ruffio et al. (2018) and Mawet et al. (2019), neglect the effects

[10] One could compute the likelihood by injecting negative fake planets into the raw data; reperforming the post-processing for each epoch, position, and photometry; and examining the residuals; however, this would increase the compute time by a very large constant factor and is not computationally feasible at this time.

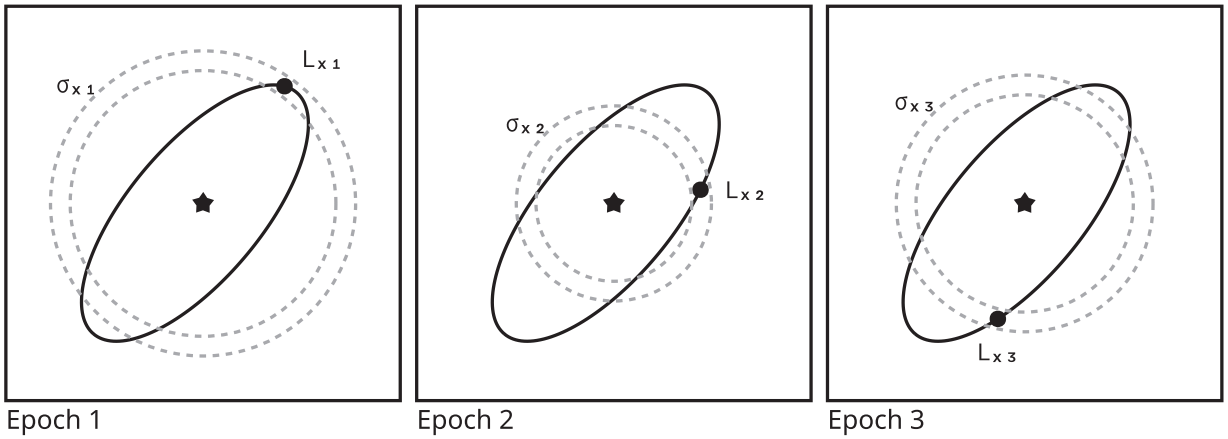

Figure 3. Schematic showing how we model planets across epochs with orbital motion. The solid black line shows the path of a hypothetical planet around a star, with a given set of parameters (a, i, e, Ω, ω, τ, M, and Π). We calculate the position x in each image i using Kepler's laws, and measure the flux at that location as $L_{x,i}$. We then calculate $\sigma_{x,i}$, the uncertainty in $L_{x,i}$, by measuring the contrast in an annulus at that separation from the matching backward-rotated image. $L_{x,i}$ is then compared to the model parameter L, which is the same for all epochs using the likelihood function given by equation N.

of distorted planet PSFs, and simply perform a noise-weighted convolution of our data by an Airy disk of $1\lambda/D$.

Because the purpose of this work is to detect or place limits on the mass of any additional planets rather than precise orbital characterization, we do not consider any systematic errors from the instrument pointing or north angle. Error in registration or north angle could bias the orbital parameters and reduce our ability to recover planets. Fortunately, the results of Yelda et al. (2010) and Service et al. (2016) indicate that the north angle and platescale of NIRC2 are very stable over time, varying less than 0.1 mas px^{-1}, respectively, between 2010 and the service in 2015.

We consider separate models for the four known planets, an additional outer planet, and an additional inner planet. Results pertaining to each planet are colored consistently across figures.

The purpose of the four known planet models is to confirm we can recover their photometry. These models will additionally confirm that there are no significant north angle offsets between epochs that could impact our ability to detect additional planets. We choose uniform priors for the angular parameters ω and τ for each planet and broad but informative priors based on previous work for the orbital plane of the planets, eccentricities, and the planets' L-band flux ratios. We adopt uniform priors on the planet semimajor axes but truncate them in order to separate the planet models and prevent them from each sampling all four modes of a single posterior. For the orbital planes, we chose wide Gaussian priors based on previous modeling of the outer planets' orbits and planetary radial velocities by Ruffio et al. (2019). This constrains the direction of motion along our line of sight, which roughly halves the size of the parameter space to explore.

We select priors on the flux ratio that require it to be greater than or equal to zero, but not less than zero. This is because ADI processing introduces dark wings around point sources. If a point source follows a Keplerian orbit through our images, the dark wings will nearly follow this same orbit, leading to spurious detections of negative planets and/or reducing the significance of a detection by introducing false uncertainty in its flux ratio. Besides enforcing positivity, we expect the exact shape of priors on the flux ratio to have little effect on the posterior, because we have more than enough data for the likelihood to overwhelm the prior.

For the inner and outer planet models, we again adopt broad but informative priors on the orbital plane of the system and a beta distribution to prefer low eccentricities. We adopt a uniform distribution for semimajor axis a between 1 and 14 au, which constrains our search to orbits closer in than planet e. Finally, for the flux ratio between the inner planet and star, we adopt both a uniform prior between 0 and $10\times^{-2}$ and a log normal prior centered on the expected pixel values. Given the amount of data, we expect the posterior to be relatively insensitive to this choice of prior (an assumption we will verify in Section 4.3). For the angular parameter ω, which has a uniform prior, we in fact sample from a pair of Gaussian distributions, ω_x and ω_y centered at zero, and calculate $\omega = \tan^{-1}(\omega_y, \omega_x)$. This is has the same distribution as a uniform prior on ω, but allows the sampler to easily wrap around past 0 and 2π. We do the same for τ, but restrict it to a domain of [0, 1) by diving by 2π.

Finally, for all models we adopted a Gaussian prior on host mass following Konopacky et al. (2016) and Wang et al. (2018) based on interferometric measurements of the stellar radius by Baines et al. (2012). For parallax Π of the system, we use a tight Gaussian prior from GAIA's EDR3 data release. We describe the priors further and sources for all parameters in Table 3.

3.1. Detection and Limits

To evaluate detections, we marginalize over all of the orbital parameters and inspect the flux-ratio (L) posterior. This histogram represents the posterior distribution of the planet's photometry regardless of the orbital parameters. Its central value is the best estimate of the planet's photometry, and the width of the distribution is the uncertainty in that estimate. A sharp peak that is well-separated from zero indicates a detection.

For ease of comparison with other methods, we summarize this posterior by calculating the S/N as the median divided by half the 84th–17th percentile distance. This is analogous to the standard S/N calculation used to evaluate point source detections in single images; however, it is marginalized over all plausible orbits making it a stricter measure. In a traditional S/N map, each point source is considered separately, even if there were, for example, many significant point sources with

Table 3
Model Priors

Parameter	Prior Distribution	Notes
M	$1.52 \pm 0.15\ M_\odot$	Baines et al. (2012), Wang et al. (2018), Konopacky et al. (2016)
Π	24.46 ± 0.05 mas	GAIA Collaboration (2021)
i		Ruffio et al. (2019)
Ω	$89° \pm 27°$	Ruffio et al. (2019)
ω	Uniform circular[a]	
τ	Uniform circular[a]	
b		
L_b	$1 \times 10^{-4} \pm 1 \times 10^{-4}$, $L_b \in (0, 1)$	Marois et al. (2008b)
a_b	Uniform(50, 180) au	Images masked outside of 130–200 px separation
e_b	Uniform	
c		
L_c	$2 \times 10^{-4} \pm 1 \times 10^{-4}$	Marois et al. (2008b)
a_c	Uniform(30, 55) au	Images masked outside of 80–120 px separation
e_c	Uniform(0, 1)	
d		
L_d	$1 \times 10^{-4} \pm 1 \times 10^{-4}$, $L_d \in (0, 1)$	Marois et al. (2008b)
a_d	Uniform(20, 80) au	Images masked outside of 56–80 px separation
e_d	Uniform(0, 1)	
e		
L_e	$1 \times 10^{-4} \pm 1 \times 10^{-4}$, $L_e \in (0, 1)$	Marois et al. (2010)
a_e	Uniform(8, 20) au	Images masked outside of 80–120 px separation
e_e	Uniform(0, 1)	
Outer		
L_f	Uniform(0, 10^{-5})	
a_f	Uniform(100, 160) au	Images masked outside of 180–500 px separation
e_f	Uniform(0, 1)	
Inner		
L_f	Uniform(0, 10^{-2})	
a_f	Uniform(1, 14) au	Images masked outside of 9–30 px separation
e_f	Beta(1.1, 5)	Low–moderate eccentricity

Note.
[a] Parameterized using the arctangent of two standard normal distributions.

varying brightness. Here, this S/N is testing the hypothesis that there is a *single* planet with consistent flux.

We can also use a fully Bayesian approach to assessing detections. We can proceed by evaluating the relative probabilities of two models: a model of a planet with a finite brightness (M_1) and a model where there is no light from the planet (M_0). The "Bayes factor" is then the ratio of the marginal likelihoods of the data given the models times a prior on which model is more likely:

$$B_{M_1:M_0} = \frac{P(M_1|D)}{P(M_0|D)} \frac{P(M_0)}{P(M_1)}.$$

We adopt the standard prior that both models are a priori equally likely, i.e., $P(M_1) = P(M_0)$. The Bayes factor between two models gives the relative probability of M_1 compared to M_0. For example, if the Bayes factor $B_{M_1:M_0} = 10$, then given this data, it is 10 times more likely that there is a planet than not.

Often Bayes factors are challenging to calculate numerically because MCMC-based methods only produce samples proportional to the posterior density. However, in our case, our two models are said to be "nested" because M_1 reduces to M_0 for $L = 0$. Because our prior on L is uncorrelated with the priors on the orbital parameters, we can calculate the Bayes factor between these nested models using the Savage-Dickey density ratio (Dickey 1971; Koop 2003). This allows us to calculate the Bayes factor $B_{M_0:M_1}$ as

$$B_{M_1:M_0} = \frac{P(L = 0)}{P(L = 0|D)}.$$

That is, the prior on L evaluated at (or near, for numerical purposes) $L = 0$ divided by the marginal posterior of L evaluated at that same location.

A benefit of this approach is that we can assess detections without assuming the marginal flux-ratio posterior is Gaussian. This could occur even with perfectly Gaussian noise in the images if there is a strong dependence of the flux ratio on one or more orbital parameters like a or e. Note, however, that the model M_1 itself assumes that the residual noise after post-processing is approximately Gaussian, an assumption we will verify in Section 4.5. Numerically, this approach requires posterior samples where the flux ratio is near zero, which for significant detections may be far in the tails of the distribution. This calculation would require an impractically large number of samples for very significant detections that are well-separated from zero. Of course, estimating the Bayes factor accurately in order to asses a highly robust detection is somewhat moot. In any case, the use of nested sampling may allow one to reliably calculate the Bayes factor between these two models for such significant detections. For an example of using nested sampling and a Bayes factor to evaluate the presence of an exoplanet from aperture masking interferometry data, see Blakely et al. (2022).

A limitation of both the S/N and Bayes factor approaches to evaluating detections is that they consider only a single planet. If there were, for example, two planets in the data with different orbits and/or flux ratios to the star, the overall S/N would suffer.

In these instances, the posterior must be examined more closely to disentangle the planets. For this situation, we attempt to reproduce a classical direct imaging S/N map that is not a direct output of this analysis method. To do so, we first select a given date. Natural choices could be the average date or date of the best input data set. At this date, each planet drawn from the posterior has a well-defined spatial position that we calculate. We then examine the marginal flux-ratio distribution of the posterior draws that fall in each given pixel on that date. Finally, we calculate the S/Ns of these samples in the same way as above. These maps are built using the posterior, so they only include the most a posteriori likely orbit and photometry parameters (no samples are available to perform this analysis along very unlikely orbits). We refer to these as "photometric accuracy" maps, to distinguish them from traditional S/N maps.

Finally, we present Bayesian upper limits following Ruffio et al. (2018) by calculating the 84th percentile of this same marginal flux-ratio posterior. In our results, we present the Bayes factor in addition to our analog of the classic S/N

whenever this calculation is feasible. The Bayes factor is arguably a more robust quantity, but it lacks a history and the established threshold conventions of the S/N in the context of direct imaging.

3.2. Sampling

Sampling from images is a difficult problem because the direct imaging likelihood function has strong modes at the locations of planets and speckles surrounded by large flat regions where the likelihood is negligible. Compared to fitting orbits to astrometry points, sampled orbits that fall far away from any modes do not have gradients that pull subsequent samples toward a mode. Additionally, fitting near face-on orbits is challenging due to degeneracies between Ω, ω, and τ in our chosen parameterization. These effects combine to require small step sizes to explore near the mode and many steps to adequately explore the tails of the posterior. Our numerical tests showed that simulated systems with companions injected with an overall S/N of approximately 3–6 are the most computationally demanding to sample, because the sampler must explore multiple peaks (S/Ns closer to 1 are associated with posteriors that are broad and relatively smooth, and above 6 the peak dominates).

Knowing this would be computationally demanding, we programmed our model in Julia (Bezanson et al. 2012). We used forward-mode automatic differentiation (Revels et al. 2016) to calculate the gradients of the log-posterior with respect to each parameter. This allowed us to use a higher-order sampler without manually deriving gradients for each model. We sampled from the posterior using the No U-Turn Sampler (Hoffman & Gelman 2014) variant of Hamiltonian Monte Carlo, as implemented in AdvancedHMC.jl (Xu et al. 2020). Internally, we used Bijectors.jl (Fjelde et al. 2020) to transform all constrained variables and priors to unconstrained distributions.

For each model, we used multiple independent chains with a maximum tree depth of 13 steps. We initialized each independent chain by drawing 50,000 samples from the priors and picking the sample with the highest posterior density. After adapting the step size and mass matrix for 30,000 iterations and discarding the first 150,000 iterations, we ran each chain in increments of 150,000 iterations until converged. We thinned each chain by its autocorrelation time and further discarded occasional chains that failed to adapt and remained at their initial parameters.

We checked for convergence by inspecting trace plots, ensuring all parameters had effective samples sizes (ESS) greater than 100, and that the Gelman, Rubin, and Brooks diagnostic was less than 1.3 (Gelman & Rubin 1992; Brooks & Gelman 1998).

The code used in this paper is available in the Julia packages PlanetOrbits.jl[11] and DirectDetections.jl.[12]

3.3. Stability

We further evaluate the results of our inner planet model by testing them for orbital stability using the Python REBOUND WHFast integrator (Rein & Tamayo 2015) to integrate sets of five-planet orbits for 100,000 yr. We determine whether the orbits are stable using the Mean Exponential Growth factor of Nearby Orbits (MEGNO; Cincotta et al. 2003) factor. The orbits that present a MEGNO of $\leqslant 2$ for 100,000 yr, which would indicate stability up until that time, are then integrated on a range of semimajor axis and eccentricity for 1 Myr to find possibly stable neighboring orbits.

We evaluate five-planet solutions in two ways. First, we start by sampling five-planet configurations from the posteriors for b, c, d, e, and the candidate planet f. Then, we also analyzed the stability of the candidate planet f from our posteriors with the b, c, d, and e planet parameters from the V_d model presented by Goździewski & Migaszewski (2014).

4. Results

In this section, we begin by describing the results of our models of the four known planets. Then, we describe our results of applying the same approach to search for any additional outer and inner planets.

4.1. Recovery of Known Planets

The four known planets are easily recovered by our models at very high significance. Figure 4 shows orbital paths drawn from the posterior of the four-planet near-resonant model. Figure 5 shows the marginal posteriors of the photometry and selected orbital elements for each model. Despite not extracting astrometry points as an intermediate step, the orbital posteriors are consistent with previous studies (Wang et al. 2018).

The marginal L' histograms in Figure 6 show that the photometry posteriors are approximately Gaussian-distributed and are consistent with previously reported values (Marois et al. 2008b, 2010). The distributions are well-separated from zero, indicating robust detections. We find S/Ns of the planets b, c, d, and e from the combined observations of 122, 154, 125, and 32, respectively. Compared to the S/N measured at each individual epoch, these S/Ns are greater by roughly a factor of $\sqrt{5}$, which is the ideal improvement in S/N we would expect by combining five observations limited by Gaussian noise.

Compared to previous studies of the system's orbital configuration using, for example, GPI (Wang et al. 2018), SPHERE (Wertz et al. 2017), and GRAVITY (GRAVITY Collaboration et al. 2019), these longer-wavelength observations have reduced astrometric precision. We therefore present our orbital solutions here to show that the orbits derived by directly modeling the photometry in the images are consistent and to build confidence in this approach before applying it to search for additional unseen companions.

4.2. Limits on Additional Outer Planets

Now that we have established that our technique recovers the four known planets, we turn our search outward to search for any additional outer planets between 100 and 150 au. For this search, we ignored the 2009 epoch due to nodding artifacts beyond the orbit of planet b. The posterior of this model contained one peak that we identified as a bright artifact on the far top edge of the 2017 epoch. We dropped samples with orbits that intersected that artifact before proceeding with our analysis.

We find no evidence for a fifth outer planet above an 85th percentile L' contrast of 4.6×10^{-6}. The overall S/N from this posterior is 0.7, and the log Bayes factor for an additional outer planet given this data is -1.6. Figure 7 presents our sensitivity as a function of semimajor axis as well as histogram of the full

[11] https://sefffal.github.io/PlanetOrbits.jl/dev/
[12] https://sefffal.github.io/DirectDetections.jl/dev/

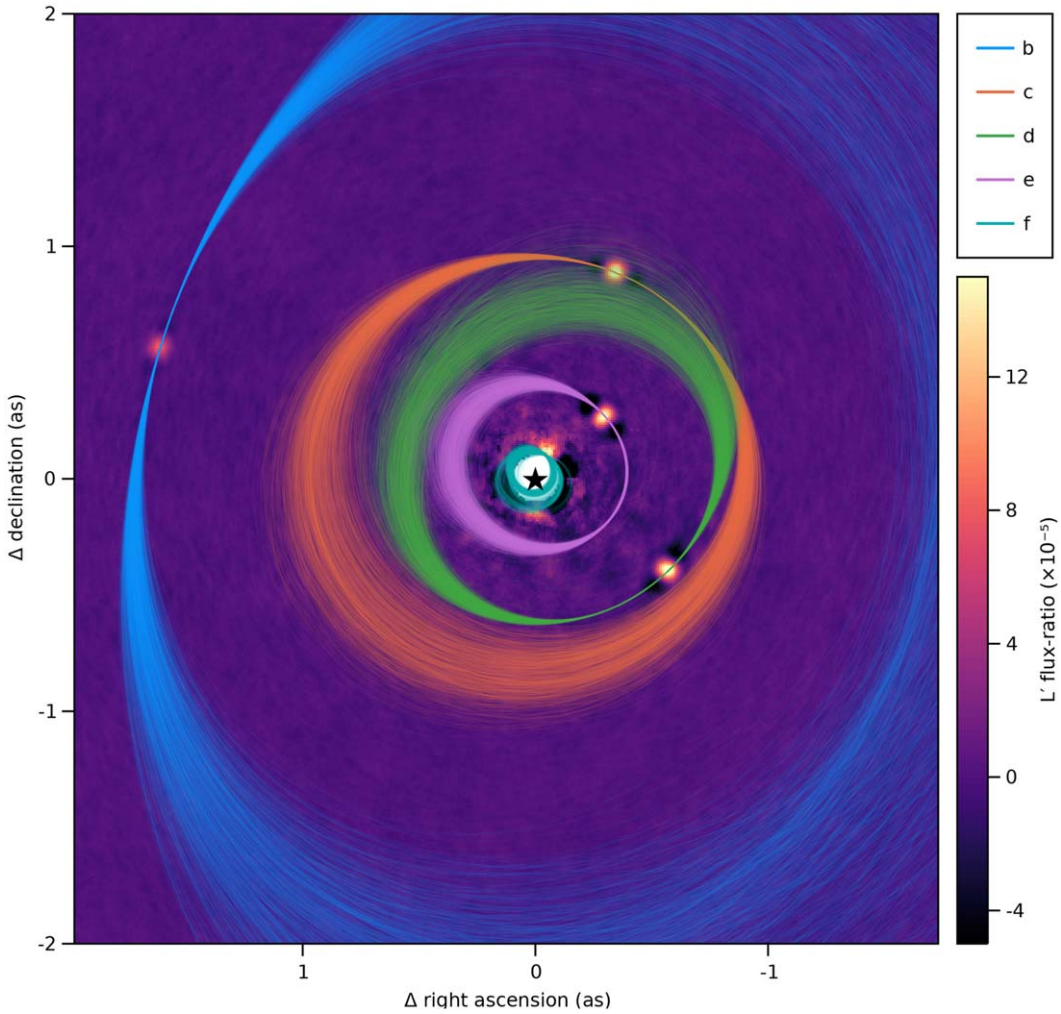

Figure 4. Visualization of 1500 orbit draws from the posteriors of the b, c, d, e, and inner planet models overplotted on the combined 2021 epoch. Directly modeling the photometry in our images allows us to simultaneously detect the known planets and produce orbital posteriors that agree with previous fits to extracted astrometry.

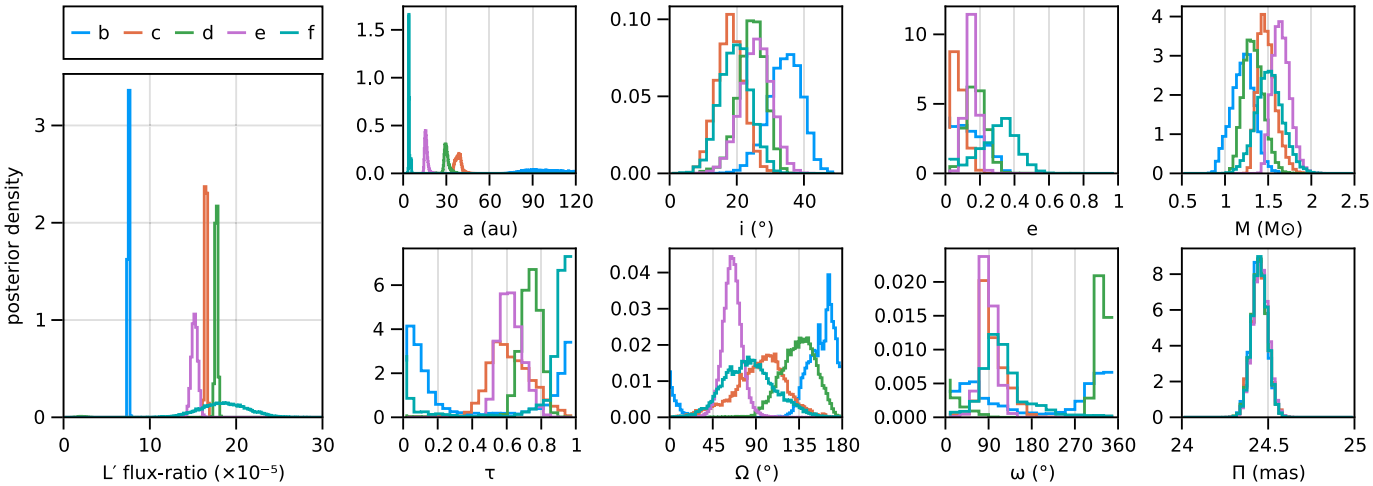

Figure 5. Marginal posteriors of photometry and orbital elements compared between the four known planets and the inner planet model.

marginal photometry posterior. Using the system ages of Sepulveda & Bowler (2022) and COND models of Baraffe et al. (2003), this corresponds to a 1σ mass limit of roughly 0.9 $M_{\rm jup}$.

4.3. Evidence for a Fifth Inner Planet

We now consider an additional planet interior to planet e. Against the full data set, the model finds a mode close to the star. Figure 8 shows the marginal photometry versus semimajor

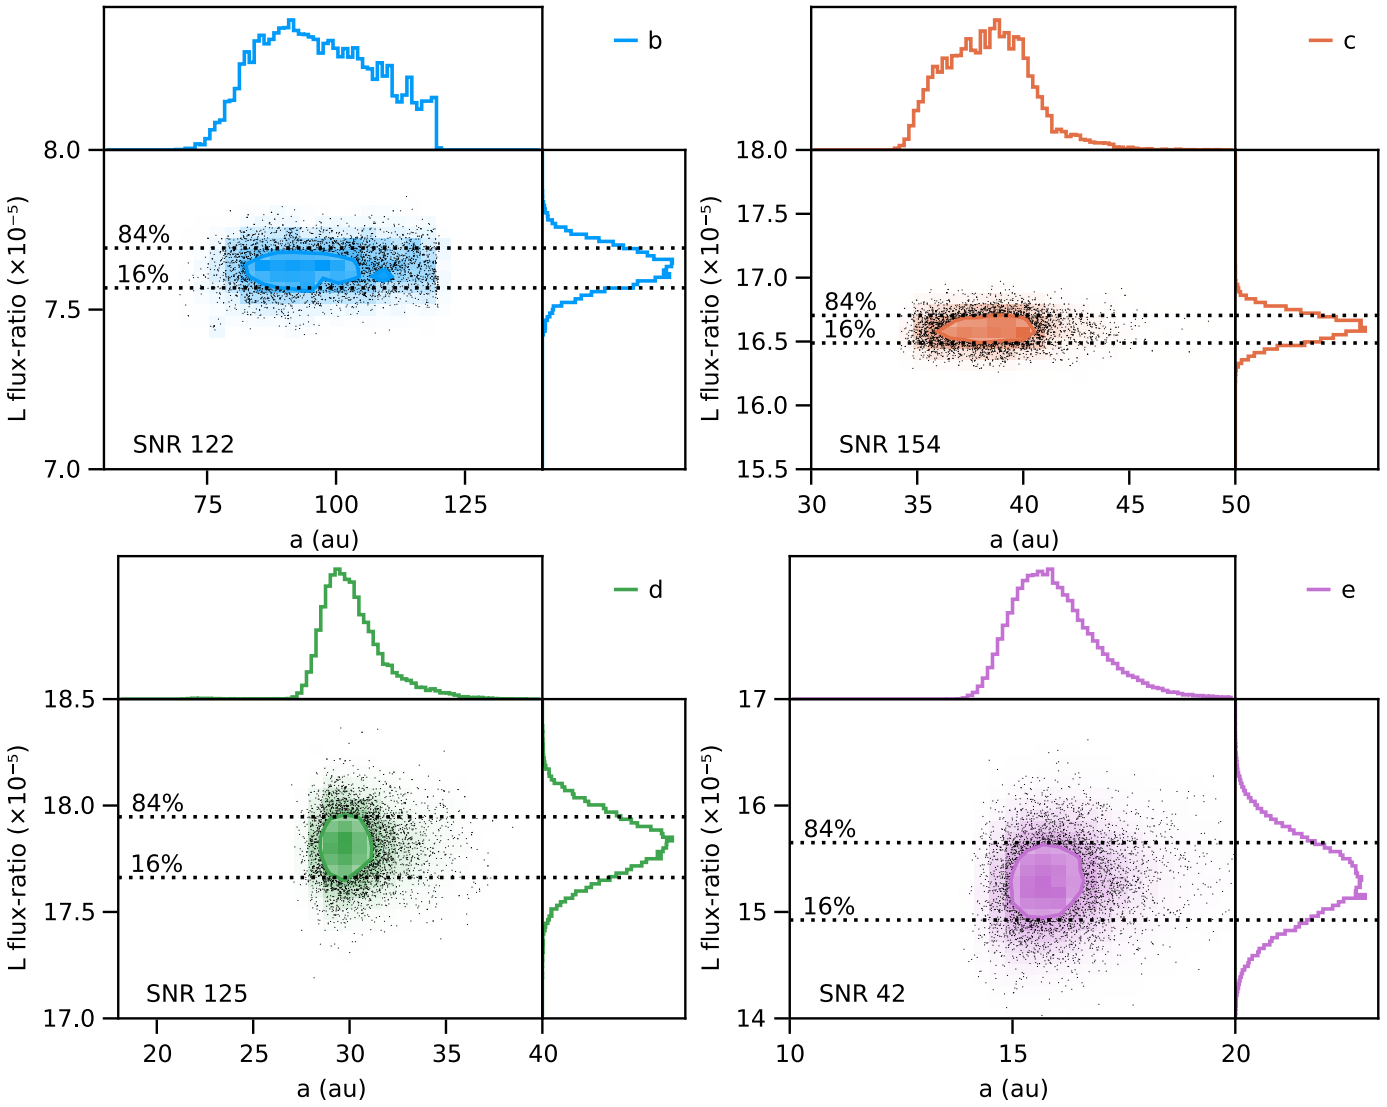

Figure 6. Marginal photometry vs. semimajor axis for planets b, c, d, and e. All four are detected at very high significance and S/Ns greater than in any individual epoch. The scales are different between each panel.

axis posteriors of our single inner planet model. The joint posterior of the inner planet model is multimodal, with 2–3 families of plausible orbits that all pass through the same locations in 2020 and 2021 (Figures 9 and 18). This multimodality is a result of weak photometric constraints in some epochs, leading the model to consider several plausible locations with consistent flux as the peaks found in higher quality epochs. Nonetheless, the marginal photometry posterior is roughly Gaussian and well-constrained.

The marginal photometry posterior is well-separated from zero, with a mean that is very similar to the photometry of planets c, d, and e. We find an S/N of 6.9 and a log Bayes factor of 18.

The marginal semimajor axis posterior is centered at 4.5 au but is cut off below ∼4 au due to detector saturation, so we cannot place a firm lower limit. The 85th percentile upper limit is 4.6 au. The chains for this posterior are available at 10.5281/zenodo.6823071.

Figure 9 shows images from our five epochs with positions calculated from orbits drawn from the posterior. The model places it NW of the star in 2021 and WNW in 2020. There are three plausible modes in 2017, and the location in earlier epochs is not well-constrained.

A full corner plot showing the values, uncertainties, and covariance between all nine parameters is available in Appendix, Figure 17.

In Figure 10, we draw orbits from the posterior and calculate their positions in 2021. We see that the model prefers a single location for the planet in 2021 north–north–west of the star. When we look at the median photometry of samples from the posterior that fall in this pixel, we find they are all roughly 2×10^{-4} in units relative contrast. Finally, when we look at the spread of the sample photometry, i.e., the "photometric accuracy," we again find a cluster of S/N 5–9. These maps support the posterior being consistent with a single object rather than two or more.

Returning to the literature, various candidate point sources have previously been reported. The candidate reported by Maire et al. (2015) 3–4σ 0.″2 due south of the star in 2013 is not consistent with with our results. Likewise, no compatible point sources are visible in the shorter wavelength (YJH) IFS data of Wahhaj et al. (2021). Their IRDIS data in the K band (closer to L') do show a low-S/N point source northwest of the

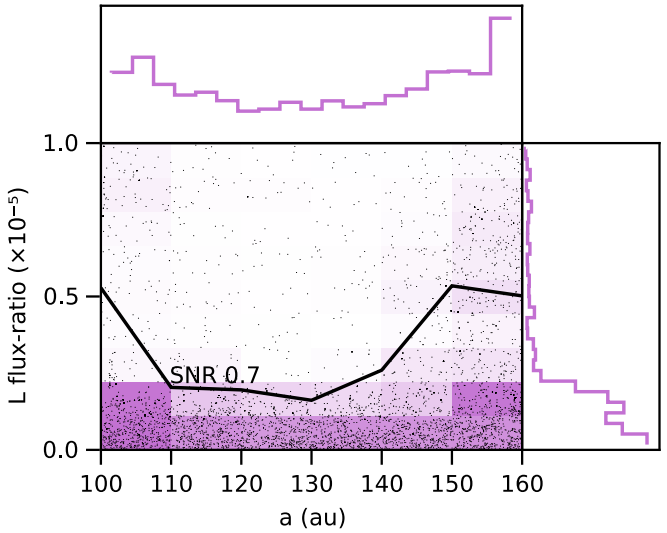

Figure 7. Flux-ratio and semimajor axis marginal posteriors for an additional outer planet between planet b and the start of the outer debris disk. The black line gives the 84th percentile flux as a function of semimajor axis. No planet is detected, and our sensitivity is not strongly dependent on any of the orbital parameters within this range of semimajor axis.

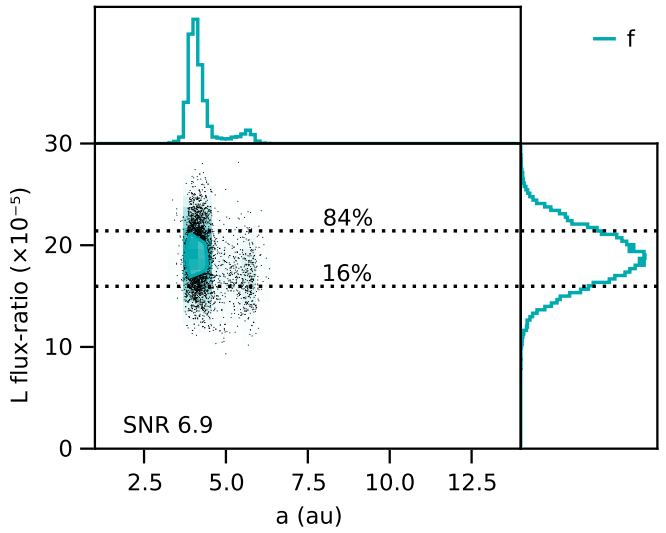

Figure 8. Flux-ratio and semimajor axis posteriors for an additional inner planet with a 1σ contour overplotted. The dashed lines show 16% and 84% percentile limits. The recovered photometry is consistent with that of planets c, d, and e.

star in 2019; however, again, its location may not be consistent with our orbital posterior to $>1\sigma$. On the other hand, the point source reported by Currie et al. (2014), 4σ north–north–west of the star in 2012, may have a roughly correct position angle if the candidate's semimajor axis is $\approx$4.7, though with a slightly greater separation. We did not include these sequences in our initial data selection (Section 2.2), so to add them afterward knowing they may or may not have compatible point sources could introduce confirmation bias.

This is an intriguing result, but given the novelty of this analysis technique, lack of confirmation at other wavelengths and instruments, and other points discussed later in this analysis, we do not yet consider this a robust detection.

4.4. Contribution of the 2021 Epoch

Figure 9 shows that the preferred location of model is the most localized in 2021. This was our best epoch and consists of observations taken over four nights. We now examine the impact of this epoch on the model in greater detail.

To begin, we injected a negative model planet into the 2021 sequences raw data prior to SNAP reduction. We placed the planet at the posterior expected position and photometry calculated from the full model of all five epochs. We then re-reduced the data with SNAP to produce Figure 11.

The panels in that figure show the image before and after injecting the negative planet model. The center panel giving the difference between these reductions shows the flux removed by the planet model. The model reproduces much of the irregular structure directly around the star, including a bright lobe opposite to the expected position. If the candidate is real, those effects can be understood as artifacts of the SNAP algorithm, tight separation, and rapid, near-180° field-of-view rotation as the system transits the meridian from Maunakea. In fact, the angle between the two bright spots is just under the average field-of-view rotation in the 2021 sequences. Encouragingly, the opposite bright lobe is not picked up by the inner planet model, meaning that a planet at that location and brightness in 2021 is not consistent with the other epochs. That said, the injected planet model is not a perfect match for the candidate. The negative sidelobes (artifacts of any ADI reduction) are somewhat mismatched with the model having a darker sidelobe to the east than and lighter to the west than the candidate. As a consequence, some flux remains to the northeast of the star after subtraction.

Next, we also run our model on the original unmodified data but fully exclude the 2021 epoch. In this case, we again find no detection though interestingly the earlier four epochs still predict a spot of high posterior density within $\sim$1 λ/D of the location found in the full model. These results do not mean that the earlier four epochs do not contribute to the S/N of the planet candidate. They still contribute by ruling out large swaths of the orbital parameter space and the bright area south of the star in 2021.

4.5. Noise Distribution and Sample Size

When evaluating a candidate in direct imaging, we should consider non-Gaussian noise and small sample statistics. Both act to increase the false-positive fraction (FPF) and reduce our confidence in a detection. We calculate penalty factors for both of these effects on 3σ and 5σ Gaussian-equivalent FPFs.

First, we consider non-Gaussian noise. Following Marois et al. (2008a), we compare the distribution of our data at the separation of the inner planet candidate to a Gaussian and estimate its effect on both detection thresholds. In the top panel of Figure 12, we plot the confidence levels of a Gaussian distribution and of standardized pixel data from the inner $1\lambda/D$ annulus of all five backward-rotated photometry maps. We extrapolate our data by fitting an exponential, and find that, to reach a confidence level equivalent to a Gaussian at 3σ and 5σ, we should penalize our S/N by factors of 1.07 and 1.35, respectively. The bottom panel of Figure 12 shows the log probability density function (PDF) of a Gaussian distribution and an empirical PDF (EPDF) of standardized pixel data from the inner $1\lambda/D$ annulus of the backward-rotated photometry maps. By expanding the Gaussian by these factors, it fully encompasses our data at a z-score of 3. The backward-rotated

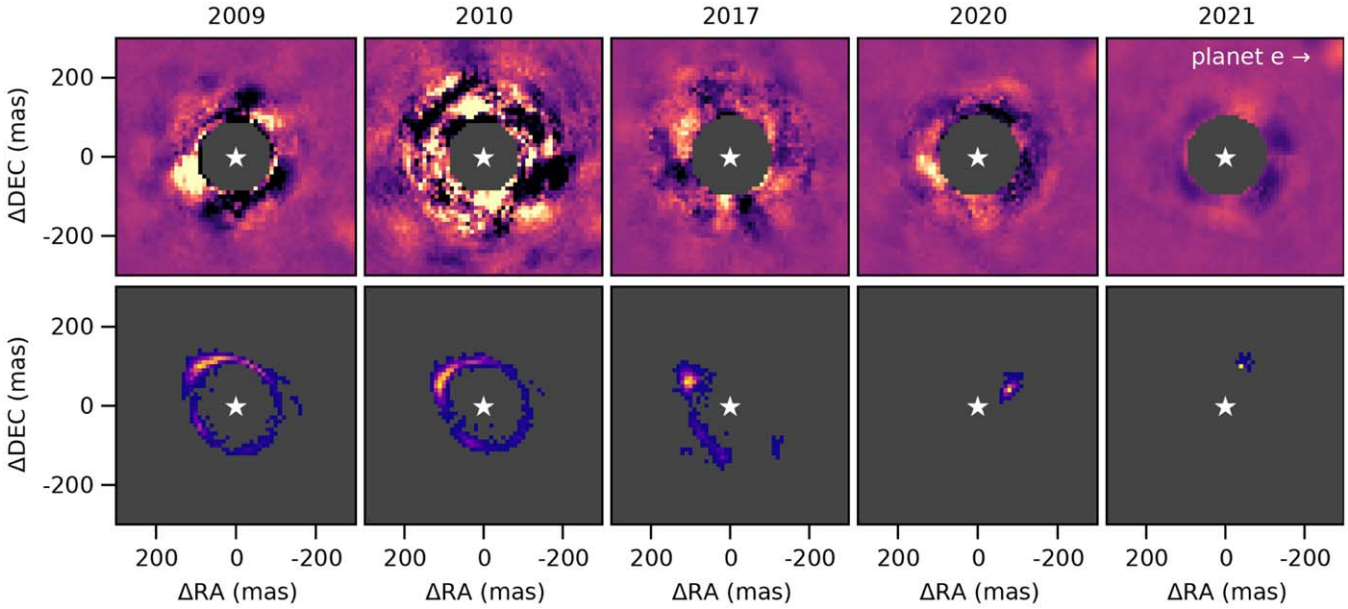

Figure 9. Top: images from each of the five epochs. Bottom: 2D marginal position posterior at each epoch. This shows the a posteriori most likely positions for an inner planet at each epoch. Note that the posterior is calculated jointly from all images; this figure merely visualizes it at individual epochs. Even epochs where the location is poorly constrained can still contribute to the S/N of the model as a whole by reducing the space of plausible orbits and planet photometry.

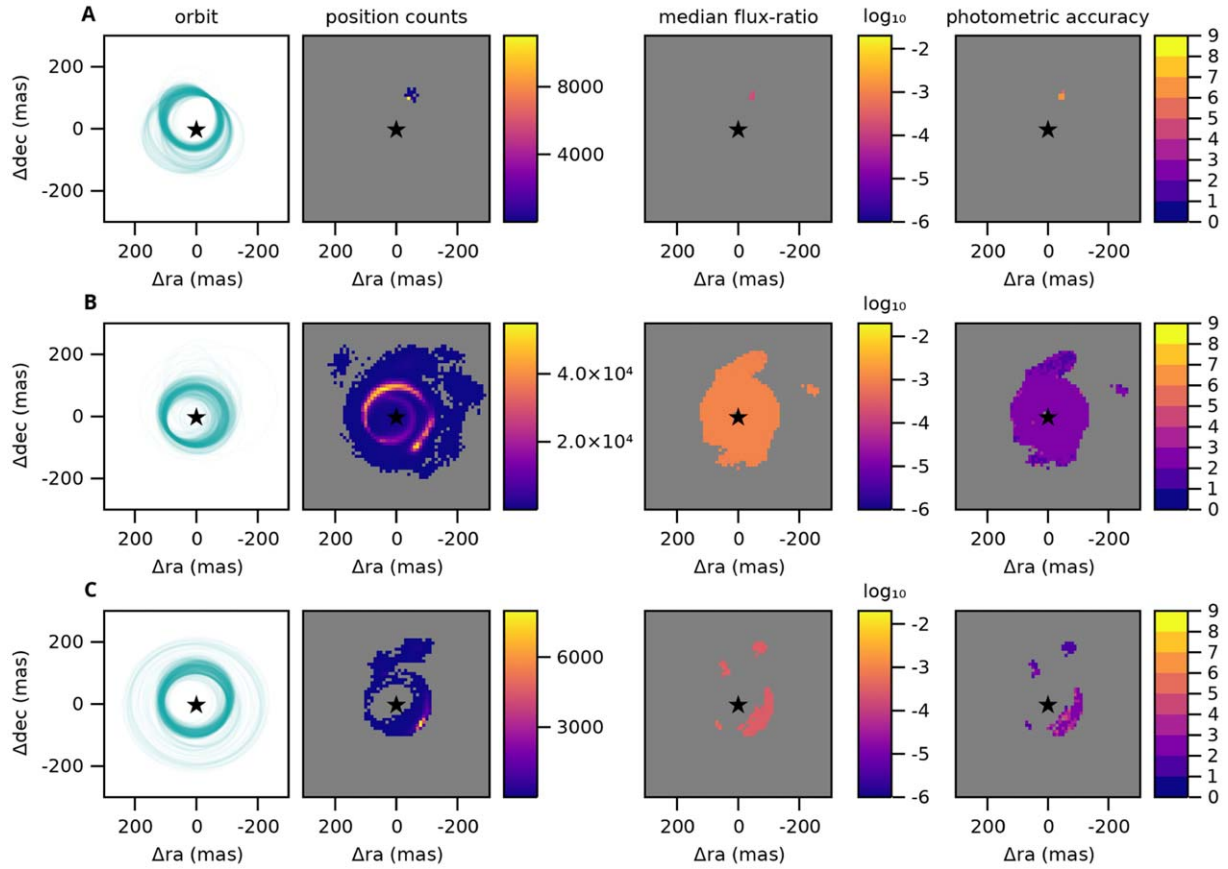

Figure 10. Comparison of the inner planet models against three different data sets. A. An inner planet model applied to the full data set. B. Same as A, but ignoring the best epoch (2021). C. Same as A, but after subtracting the best fitting model found by A from the 2021 epoch. The four panels show orbits drawn from the posterior, posterior density of planet position calculated at the 2021 epoch, median photometry if the planet were at that position in 2021, and photometric accuracy at those positions, i.e., the S/N of a planet if it were on an orbit that would pass through that pixel in 2021. Pixels are left blank where there are only a negligible number of samples.

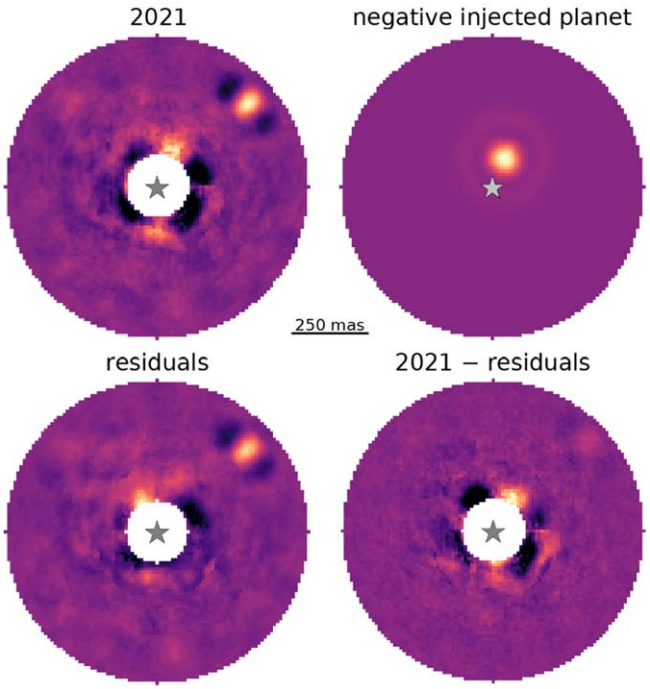

Figure 11. The 2021 epoch before and after subtracting a model planet injected at the expected location and photometry of the single inner planet model fit to all five epochs. The area interior to 100 mas is excluded from the processing, due to detector saturation. The bottom right image shows the flux removed by the negative planet injection. The structure is more complex than the typical dark wings from ADI processing, due to the high FoV rotation and tight separation.

noise maps do not contain any pixels at z-scores beyond 3.5 (the forward reduction does, of course, because it contains the signal of the candidate), but the extrapolation appears valid and conservative. These factors are relatively small, indicating that the residual noise in each epoch is close to Gaussian-distributed. This is not surprising, due to the central limit theorem, because each epoch besides 2010 is a stack of 3–5 sequences with uncorrelated noise.

Next, we consider the effect of small sample statistics near the star. The contrast curves underlying our model at each epoch are calculated from as few as five independent samples at a separation of roughly $2\lambda/D$. If each epoch contributed to the S/N of our model in equal measure, then we might consider the noise sample to be five times larger, reducing the effects of small sample statistics. However, the 2021 epoch contributes significantly to the overall figure. It is not yet clear how to correct this model for small sample statistics, but we can take a conservative approach by considering the final S/N to come only from a single image (instead of 18 sequences) and applying the correction factors of Mawet et al. (2014), Table 1, at $2\lambda/D$. This gives penalty factors of 1.35 and 2.2 for 3σ and 5σ FPFs, respectively.

Combining these two sets of penalty factors, we should in fact apply 4.3σ and 14.9σ thresholds to reach a large sample size with Gaussian-equivalent FPFs for 3σ and 5σ, respectively. The candidate does easily meet the 3σ threshold, but it clearly falls well below a 5σ threshold for detection. This correction is perhaps overly pessimistic, but it reflects our goal to communicate that this signal is a candidate worthy of additional study rather than an unambiguous detection.

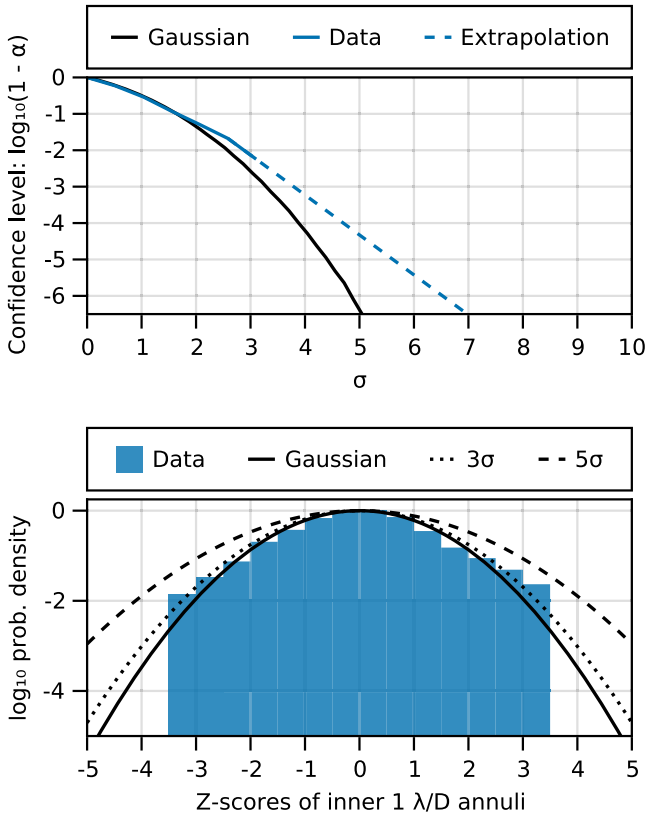

Figure 12. Top: Confidence levels of a Gaussian distribution and empirical confidence level from the inner $1\lambda/D$ annulus of our images extrapolated to $1–10^{-6}$ using an exponential function. We increase our S/N thresholds to account for this slight deviation from Gaussian noise. Bottom: Empirical log PDF of standardized pixel intensities compared to a Gaussian model, and Gaussians scaled to cover the tails of the distribution.

4.6. Mass and Proper Motion Anomaly

The photometry of the inner planet candidate is consistent with planets c, d, and e, and brighter than b. Without photometry at other wavelengths, we therefore assume that it would have a similar mass to the inner planets c, d, and e.

The masses of the HR8799 planets have been estimated using several approaches. From the beginning, luminosity modeling of the planets has suggested masses of approximately $7\,M_{\rm jup}$ for the inner planets (Marois et al. 2008b). Orbital stability analyses by Wang et al. (2018) support the inner planets having masses up to approximately $7\,M_{\rm jup}$, though small islands of stability may exist for high masses. Sepulveda & Bowler (2022), on the other hand, model the orbits of the planets in order to constrain the dynamical mass of the star. By combining this stellar mass with stellar and planet evolution models, they find that the inner planets c, d, and e likely have masses in the range of 4.1–7.0 $M_{\rm jup}$, where age is the dominant contributor to the uncertainty. In contrast to these estimates, Brandt et al. (2021) find somewhat higher masses of $9.6^{+1.9}_{-1.8}\,M_{\rm jup}$ by combining the stable orbits found by Wang et al. (2018) with the Hipparcos-Gaia Catalog of Accelerations (HGCA, Brandt 2021), which calibrates the Hipparcos catalog against Gaia EDR3 (GAIA Collaboration 2021). The addition of a fifth massive planet would alter the solution space for all three methods incorporating orbital dynamics; however, the method based on proper motion anomaly would be the most

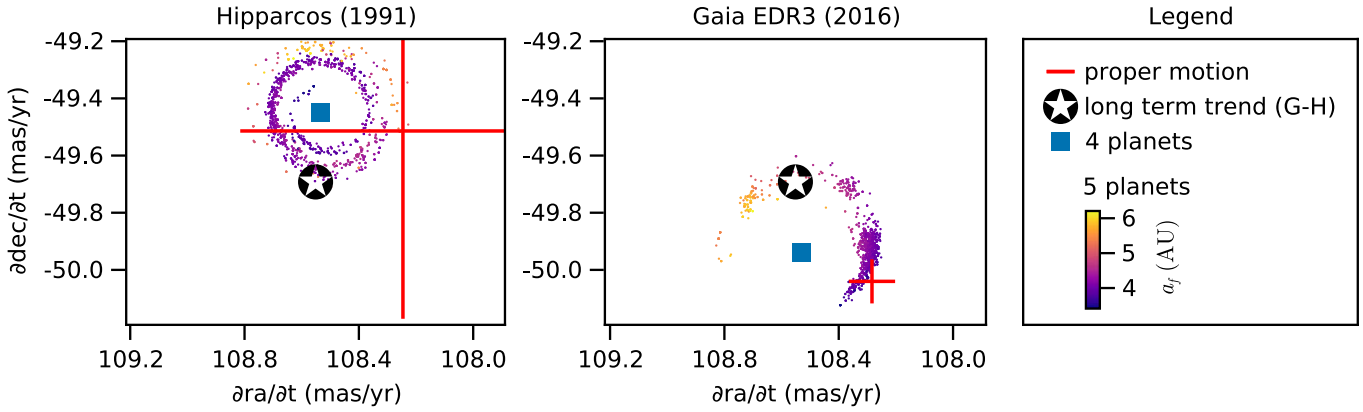

Figure 13. Stellar astrometric motion predicted at the Hipparcos and GAIA epochs for the stable four-planet V_d solution of Goździewski & Migaszewski (2014) (blue square) and with the addition of the fifth inner candidate (colored by the semimajor axis of the candidate). The red error bars show the stellar astrometric motion at the Hipparcos and GAIA epochs as calculated by the HGCA (Brandt 2021). The black marker shows the long-term proper motion from the HGCA calculated from the difference in position between both epochs. The planets c, d, e, and the candidate f are assumed to have masses of 7 $M_{\rm jup}$, while b is assumed to have a mass of 5 $M_{\rm jup}$.

Table 4
Stable Orbits

Planet	Mass (M_J)	a (au)	e	i (°)	Ω (°)	ω (°)	τ (°)
b	6.527020	68.597249	0.017425	27.502	63.953507	37.631852	0.872678
c	11.868008	39.486207	0.054102	27.502	63.953507	89.946387	0.403048
d	7.178005	25.705213	0.137796	27.502	63.953507	33.186128	0.139618
e	6.298260	15.660910	0.168239	27.502	63.953507	110.074917	0.902362
coplanar "f"	3.75000	4.325000	0.068600	27.502	63.953507	145.767982	0.833691
non-coplanar "f"	3.75000	4.510300	0.036997	14.986	82.304917	119.041162	0.837036

Note. Stable five-planet solutions. The non-coplanar solution is drawn from our orbital posterior and is stable for 0.75 Myr. The coplanar solution was found using using a search grid near the orbital parameters of the posterior of the inner planet model. The stellar mass is 1.716162 $M_\odot$. The parameters of planets b, c, d, and e are from the Goździewski & Migaszewski (2014) V_d model.

impacted. We now examine if this potential addition could account for the slight tension between these mass estimates.

Taking a more basic approach than that of Brandt et al. (2021), we model the proper motion anomaly of the star by assuming that the GAIA and Hipparcos missions each independently measure the position and instantaneous proper motion 25 times spaced equally throughout their respective missions. We then compare these quantities against observations for a four-planet and a five-planet model. For the four-planet model, we use the stable orbital parameters of Goździewski & Migaszewski (2014) and fixed masses of 5 $M_{\rm jup}$ and 7 $M_{\rm jup}$ for b, and c, d, e, respectively. For the five-planet model, we use the same parameters but add a fifth inner planet from our orbital posterior with a mass of 7 $M_{\rm jup}$.

Figure 13 shows the result of this comparison. Proper motion anomaly at the Hipparcos epoch is consistent with either the four- or five-planet model, but agreement at the GAIA epoch is significantly improved by the addition of the candidate at the most likely semimajor axis.

Though far from conclusive, this analysis demonstrates that the inclusion of the inner candidate could fully account for the observed proper motion anomaly if the candidate and inner planets c, d, and e have masses close to 7 $M_{\rm jup}$. This is slightly at odds with the Brandt et al. (2021) result, because they use the HGCA data to conclude that any additional inner planets more massive than $\approx$6 $M_{\rm jup}$ between 3 and 8 au are unlikely. The discrepancy could be due to their more sophisticated modeling or the lower masses used in our experiment for the known planets.

4.7. Stability

We now consider how the addition of an inner planet would impact the stability of the system. For these tests, we adopt a smaller than realistic mass for the candidate inner planet of 3.75$M_{\rm jup}$ in order to ease the search for stable orbits. For our five-planet models obtained entirely from the posteriors presented on Figure 5, we found that no configuration is stable for the age of the system of 10–23 Myr (Sepulveda & Bowler 2022). When using the V_d Model parameters presented on Goździewski & Migaszewski (2014) for planets b, c, d and e, we find one unconstrained orbit for "f" that remains stable for 0.4 Myr. A grid search around this sample's semimajor axis and eccentricity revealed similar orbits that are stable for up to 0.75 Myr. A grid search over coplanar orbits found regions that remain stable for up to 1.5 Myr. Most of the other configurations found from the posteriors become unstable within the first 0.5 Myr in our N-body simulations. These configurations of stable orbits are presented in Table 4 and MEGNO stability maps of the coplanar orbits are presented in Figure 14. These simulations do yet not explore changing the masses or orbital parameters of planets b, c, d, and e. It is likely that regions of greater stability could be found if these parameters were also explored in a future analysis.

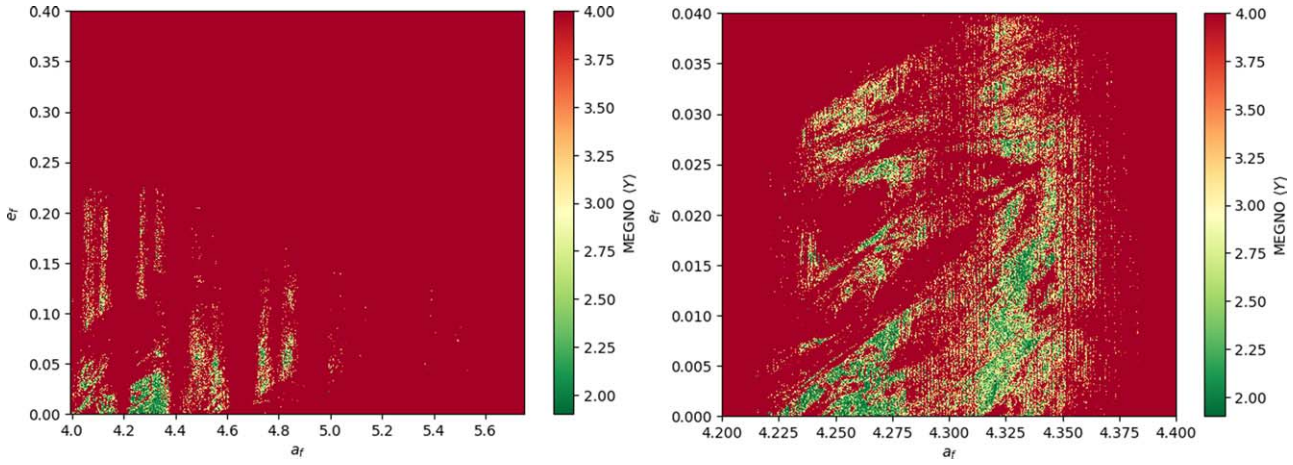

Figure 14. The MEGNO stability map for the coplanar "f" candidate using the parameters from Table 4, for an integration time of 0.5 Myr (left) and 1.0 Myr (right). Most of the configurations become unstable (MEGNO > 2) within this integration period.

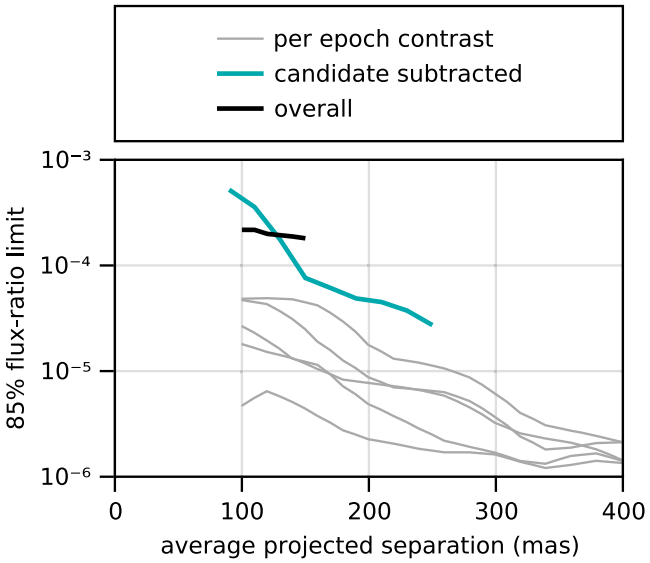

Figure 15. The 85th percentile upper limit of the marginal *L*-band flux-ratio posterior for an additional fifth inner planet (black). The turquoise lines show the same, but after subtracting a planet model from the raw data of 2021 epoch (C in Figure 10). The gray lines show the traditional 1σ projected contrast of individual epochs from backward-rotated reductions.

4.8. Sensitivity to Planets besides the 4–5 au Candidate

Figure 15 presents our sensitivity to additional inner planets besides the candidate near 4–5 au. We show the 85th percentile of the marginal L flux-ratio posterior, conditioned on different ranges of semimajor axis. We show the sensitivity for both the regular data set and the data set in which we subtracted the candidate from the 2021 epoch. Ignoring the candidate presented above, assuming the system ages of Sepulveda & Bowler (2022) (10–23 Myr), and extrapolating the COND models of Baraffe et al. (2003), this translates to 1σ upper mass limits of 4.3–3.0 $M_{\rm jup}$ for planets with orbits that have time-averaged projected separations of 150 and 250 mas, respectively. Note that this sensitivity versus average projected separation is not directly equivalent to contrast versus separation at a given epoch, as is usually quoted in the literature. In general, the sensitivity is a function of all orbital parameters. For instance, it improves with higher eccentricity at small semimajor axes, because such a planet would spend more time away from the star.

5. Conclusion

In this paper, we presented a deep targeted search in the HR 8799 system for additional planets using data at 3.8μm.

1. We observed HR 8799 for 14 quarter nights with NIRC2 between 2017 and 2021
2. We further gathered a further 4 quarter nights from the Keck archive for a 12 yr baseline.
3. We processed the data using direct S/N optimization (Thompson & Marois 2021) to improve contrasts at very small separations
4. We used the Hamiltonian Monte Carlo method to explore both the space of possible orbits and flux from planets, allowing us to search for planets in the presence of orbital motion. This approach differs from the K-Stacker (Nowak et al. 2018) technique in that it requires the flux to be consistent across epochs, gives the covariance between the planet's flux and orbital parameters, and ultimately gives a detection that can be evaluated against uncertainties in both flux and orbital parameters.

Using these methods, we presented the results of our search:

1. We demonstrated this technique by redetecting the four known planets b, c, d, and e at very high S/N
2. We presented limits on the flux of any additional outer planets between b and the start of the debris disk and did not find any additional outer planets above a 5σ significance threshold, or a mass limit of roughly 0.9 $M_{\rm jup}$.
3. We found a modest-S/N candidate interior to the orbit of planet e. This planet would have a semimajor axis of approximately 4–5 au and *L*-band photometry similar to the inner planets c, d, and e. According to the analysis Sepulveda & Bowler (2022), this would imply a mass of 4.1–7.0 $M_{\rm jup}$.

For this candidate planet, we considered several factors.

1. We verified that subtracting a planet model from the raw data at the appropriate position and intensity in the best epoch and rerunning the model removes the detection.

2. We adjusted our detection thresholds to account for slightly non-Gaussian noise and a conservative treatment of the impacts of small sample statistics.
3. We showed that the addition of a fifth inner planet could improve the agreement between the dynamical mass measurements of the planets and mass estimates/limits from atmosphere models and orbital stability modeling.
4. We performed rejection sampling with REBOUND using the inner candidate orbital posterior combined with a four-planet solution known to be stable. We were successful in finding five-planet orbital solutions that were stable for 0.75 Myr using a lowered planet mass. A grid search over nearby coplanar orbital parameters found small families of orbits that are stable for up to 1.5Myr.

Overall, we found that the inner candidate at S/N 6.9 easily met a 3σ equivalent FPF threshold (4.3σ) but did not meet a 5σ equivalent FPF threshold (14.9σ). This is primarily because of uncertainty in the contrast at each epoch, due to the limited sample size at small separations from the star. We consider this evidence intriguing, but caution that these results fall short of a conclusive detection.

Wahhaj et al. (2015) presented the most sensitive limits on the *K*-band flux of an additional inner planet. They found a 5σ upper limit of 3×10^{-5} relative to the star at 100 mas separation. Adjusting this figure for small sample statistics in the same way as our data gives an upper limit of 7×10^{-5}. As such, their nondetection of this candidate at *K*-band is a point against the candidate. On the other hand, if the candidate is confirmed, this would give the candidate a very red color compared to c, d, and e ($K - L > 2$). If real, such a red color might be caused by the different environment much closer to the star. One possible explanation might lie in photochemical hazes as proposed for 51 Eri b (Macintosh et al. 2015; Zahnle et al. 2016), due to such a planet receiving more than an order of magnitude more light from the star than planet e does.

Additional follow-up observations would be necessary to confirm this candidate; however, assuming an ideal $\sqrt{N}$ growth in S/N and a decreasing penalty for small sample statistics, this would require a further 12 quarter nights with NIRC2 of similar quality to 2021. Instead, the best chance at confirming or rejecting this candidate may come from upcoming instruments with improved contrast at 100–250 mas. Follow-up observations from GRAVITY (GRAVITY Collaboration et al. 2017) might be possible, but they would be challenging given the remaining uncertainty in the candidate's orbit. Another avenue that may be worth exploring is searching for the candidate with a fiber-fed spectrograph like the Keck Planet Imager and Characterizer (Delorme et al. 2021), though this again requires a well-determined orbit.

Regardless of whether this candidate is confirmed, we have demonstrated the utility of searching for planets in direct images by combining orbit modeling and planet detection. This approach could considerably loosen scheduling requirements when searching for rapidly moving targets like planets around Alpha Cent, in addition to making the best use of direct imaging archives.

The authors thank Dori Blakely, Thayne Curry, and the anonymous referee for providing feedback that greatly improved this work. W.T. acknowledges the support of the Natural Sciences and Engineering Research Council of Canada (NSERC), 466479467. This research used the facilities of the Canadian Astronomy Data Centre operated by the National Research Council of Canada with the support of the Canadian Space Agency. The data presented herein were obtained at the W. M. Keck Observatory, which is operated as a scientific partnership among the California Institute of Technology, the University of California and the National Aeronautics and Space Administration. The Observatory was made possible by the generous financial support of the W. M. Keck Foundation. This research has made use of the Keck Observatory Archive (KOA), which is operated by the W. M. Keck Observatory and the NASA Exoplanet Science Institute (NExScI), under contract with the National Aeronautics and Space Administration. This work has made use of data from the European Space Agency (ESA) mission Gaia (https://www.cosmos.esa.int/Gaia), processed by the Gaia Data Processing and Analysis Consortium (DPAC, https://www.cosmos.esa.int/web/Gaia/dpac/consortium). Funding for the DPAC has been provided by national institutions, in particular the institutions participating in the Gaia Multilateral Agreement.

The authors wish to recognize and acknowledge the very significant cultural role and reverence that the summit of Maunakea has always had within the indigenous Hawaiian community. We are most fortunate to have the opportunity to conduct observations from this mountain.

Software: Julia (Bezanson et al. 2012), AdvancedHMC.jl (Xu et al. 2020), REBOUND (Rein & Tamayo 2015), Makie.jl (Danisch & Krumbiegel 2021), Matplotlib (Hunter 2007), GR (Heinen 1985), PairPlots.jl.

Appendix
Additional Figures

This section includes additional information and context to aid comparisons with new or archival observations not considered in this paper. Figure 16 provides the full images used by these models.[13] The corner plot in Figure 17 provides error bounds and covariance between all model parameters for the inner planet model. Figure 18 the most likely positions of the candidate inner planet at every year from 2009 to 2028.

[13] Also available in machine readable format at DOI:10.5281/zenodo.6823071.

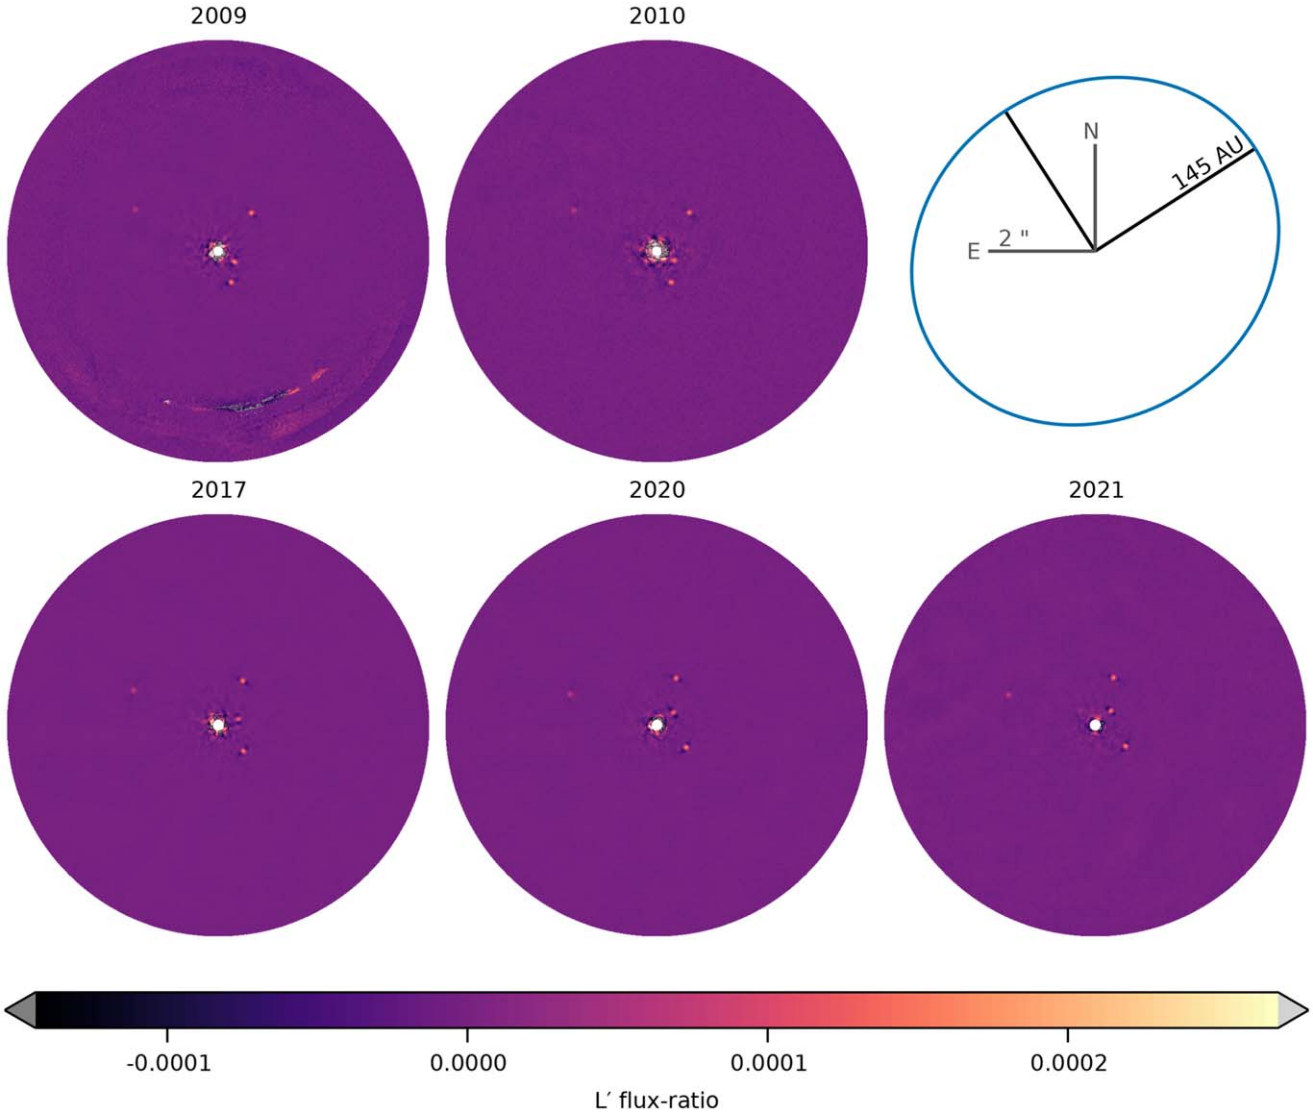

Figure 16. Full combined image from each epoch. The top two panels consist of archival data, while the bottom three are from the new campaign. The artifacts south of the star in the 2009 image are from a chopping strategy used to subtract the thermal background. The approximate location of the inner edge of the outer debris disk is outlined in blue. These processed images are available at 10.5281/zenodo.6823071.

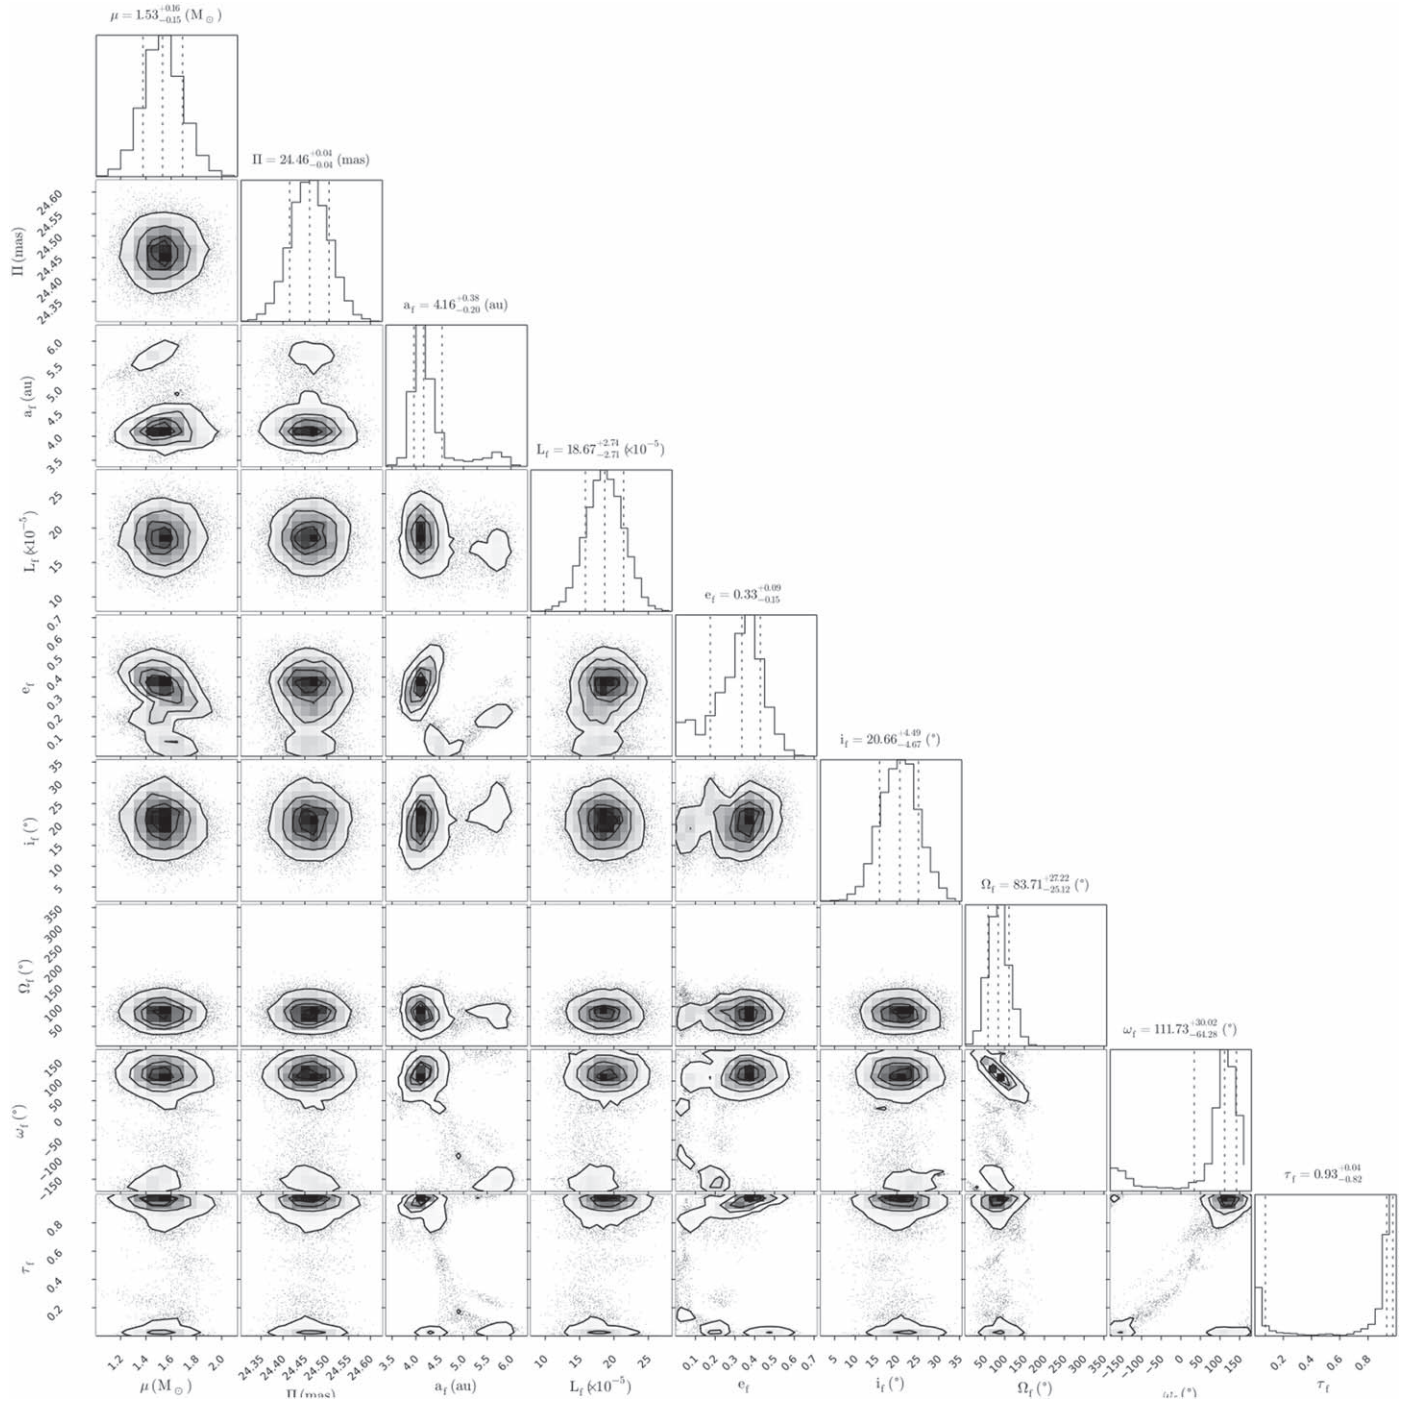

Figure 17. Corner plot showing the posterior of the inner planet model applied to all epochs. Note that the angular parameters i, Ω, ω, and τ are periodic. Chains are available at 10.5281/zenodo.6823071.

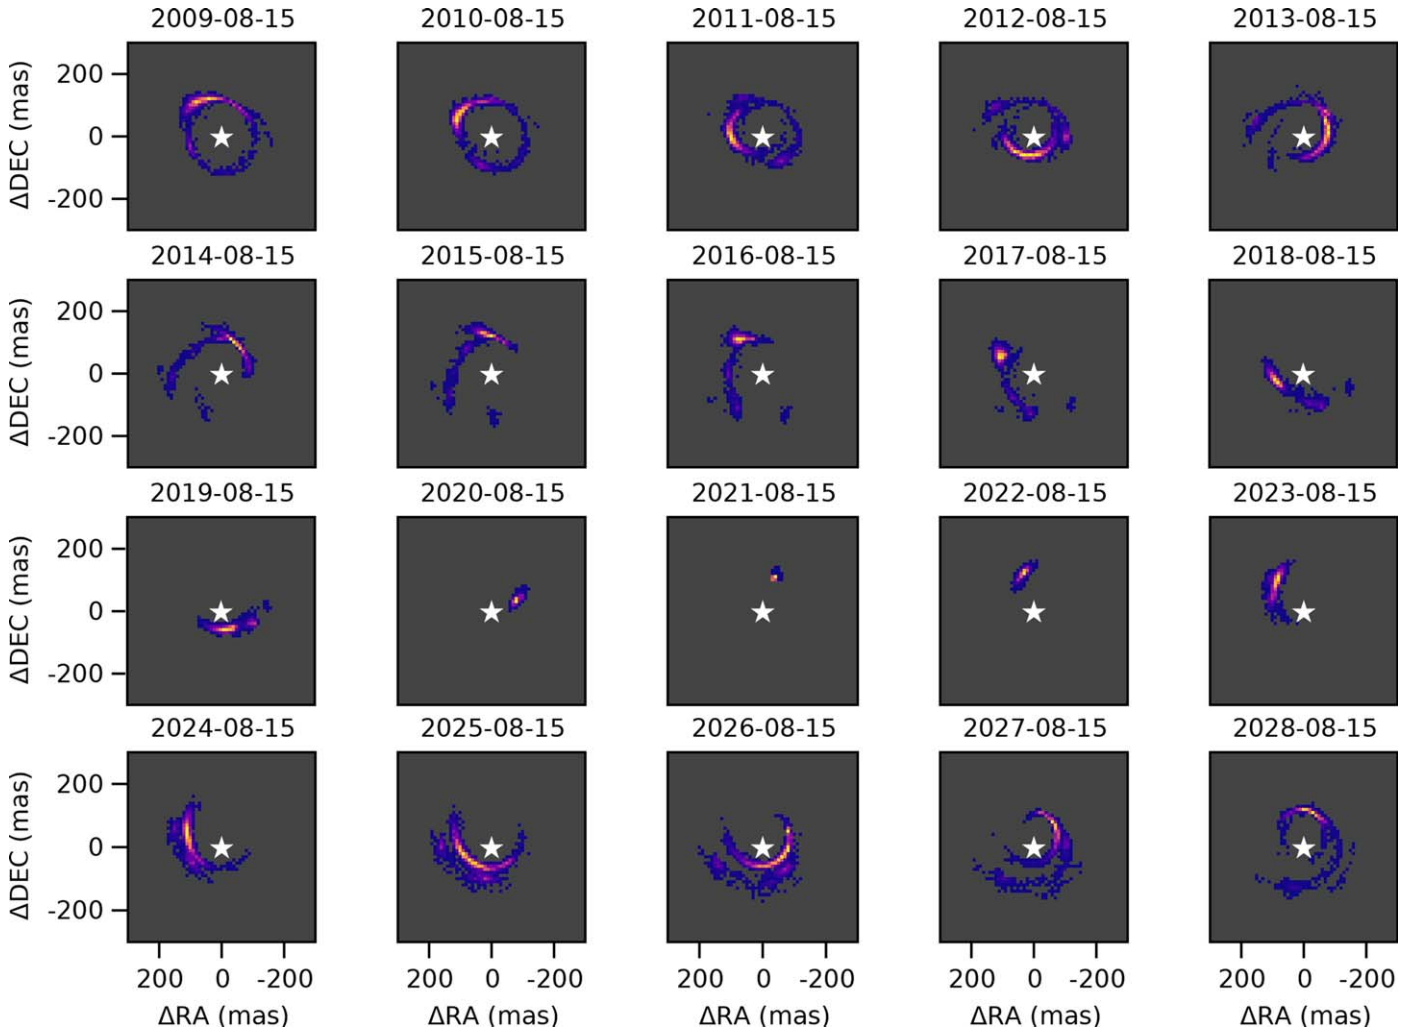

Figure 18. Position posterior density for the inner planet model at different epochs. The position is shown on August 15th of each year between 2009 and 2024.

ORCID iDs

William Thompson https://orcid.org/0000-0001-5684-4593
Christian Marois https://orcid.org/0000-0002-4164-4182
Clarissa R. Do Ó https://orcid.org/0000-0001-5173-2947
Quinn Konopacky https://orcid.org/0000-0002-9936-6285
Jean-Baptiste Ruffio https://orcid.org/0000-0003-2233-4821
Jason Wang https://orcid.org/0000-0003-0774-6502
Andy J. Skemer https://orcid.org/0000-0001-6098-3924
Robert J. De Rosa https://orcid.org/0000-0002-4918-0247
Bruce Macintosh https://orcid.org/0000-0003-1212-7538

References

Baines, E. K., White, R. J., Huber, D., et al. 2012, ApJ, 761, 57
Baraffe, I., Chabrier, G., Barman, T. S., Allard, F., & Hauschildt, P. H. 2003, A&A, 402, 701
Barman, T. S., Konopacky, Q. M., Macintosh, B., & Marois, C. 2015, ApJ, 804, 61
Bergfors, C., Brandner, W., Janson, M., Kohler, R., & Henning, T. 2011, A&A, 528, A134
Bezanson, J., Karpinski, S., Shah, V. B., & Edelman, A. 2012, arXiv:1209.5145
Blakely, D., Francis, L., Johnstone, D., et al. 2022, ApJ, 931, 3
Blunt, S., Wang, J. J., Angelo, I., et al. 2020, AJ, 159, 89
Booth, M., Jordan, A., Casassus, S., et al. 2016, MNRAS, 460, L10
Brandt, G. M., Brandt, T. D., Dupuy, T. J., Michalik, D., & Marleau, G.-D. 2021, ApJL, 915, L16
Brandt, T. D. 2021, ApJS, 254, 42
Brooks, S. P., & Gelman, A. 1998, J. Comput. Graph. Stat., 7, 434
Cincotta, P. M., Giordano, C. M., & Simó, C. 2003, PhyD, 182, 151
Close, L. M., & Males, J. R. 2009, ApJ, 709, 342
Currie, T., Burrows, A., Girard, J. H., et al. 2014, ApJ, 795, 133
Currie, T., Burrows, A., Itoh, Y., et al. 2011, ApJ, 729, 128
Currie, T., Fukagawa, M., Thalmann, C., Matsumura, S., & Plavchan, P. 2012, ApJL, 755, L34
Danisch, S., & Krumbiegel, J. 2021, JOSS, 6, 3349
Delorme, J.-R., Jovanovic, N., Echeverri, D., et al. 2021, JATIS, 7, 035006
Dickey, J. M. 1971, Ann. Math. Stat., 42, 204
Faramaz, V., Marino, S., Booth, M., et al. 2021, AJ, 161, 271
Fjelde, T. E., Xu, K., Tarek, M., Yalburgi, S., & Ge, H. 2020, in Proc. of the 2nd Symp. 118: Advances in Approximate Bayesian Inference (London: PMLR), 1
Gaia Collaboration 2021, A&A, 649, A1
Geiler, F., Krivov, A. V., Booth, M., & Lohne, T. 2019, MNRAS, 483, 332
Gelman, A., & Rubin, D. B. 1992, StaSc, 7, 7
Goździewski, K., & Migaszewski, C. 2014, MNRAS, 440, 3140
Goździewski, K., & Migaszewski, C. 2018, ApJS, 238, 21
Goździewski, K., & Migaszewski, C. 2020, ApJL, 902, L40
GRAVITY Collaboration, Abuter, R., Accardo, M., et al. 2017, A&A, 602, A94
GRAVITY Collaboration, Lacour, S., Nowak, M., et al. 2019, A&A, 623, L11
Greenbaum, A. Z., Pueyo, L., Ruffio, J.-B., et al. 2018, AJ, 155, 226
Heinen, J. 1985, GR Framework
Hoffman, M. D., & Gelman, A. 2014, JMLR, 15, 15
Hunter, J. D. 2007, CSE, 9, 90

Konopacky, Q. M., Barman, T. S., Macintosh, B. A., & Marois, C. 2013, Sci, 339, 1398
Konopacky, Q. M., Marois, C., Macintosh, B. A., et al. 2016, AJ, 152, 28
Koop, G. 2003, Bayesian Econometrics (Hoboken, NJ: Wiley)
Lafrenière, D., Marois, C., Doyon, R., Nadeau, D., & Artigau, E. 2007, ApJ, 660, 770
Le Coroller, H., Nowak, M., Arnold, L., et al. 2015, in Twenty Years of Giant Exoplanets, ed. I. Boisse, O. Demangeon, F. Bouchy et al. (Saint-Michel-l'Observatoire: Observatoire de Haute-Provence), 59
Le Coroller, H., Nowak, M., Wagner, K., et al. 2022, A&A, 667, A142
Macintosh, B., Graham, J. R., Barman, T., et al. 2015, Sci, 350, 64
Maire, A.-L., Skemer, A. J., Hinz, P. M., et al. 2015, A&A, 576, A133
Males, J. R., Skemer, A. J., & Close, L. M. 2013, ApJ, 771, 14
Marois, C., Correia, C., Galicher, R., et al. 2014, Proc. SPIE, 9148, 91480U
Marois, C., Doyon, R., Nadeau, D., Racine, R., & Walker, G. A. H. 2003, EAS Publ. Ser., 8, 233
Marois, C., Doyon, R., Racine, R., & Nadeau, D. 2000, PASP, 112, 91
Marois, C., Lafrenière, D., Doyon, R., Macintosh, B., & Nadeau, D. 2006a, ApJ, 641, 556
Marois, C., Lafrenière, D., Macintosh, B., & Doyon, R. 2006b, ApJ, 647, 612
Marois, C., Lafrenière, D., Macintosh, B., & Doyon, R. 2008a, ApJ, 673, 647
Marois, C., Macintosh, B., Barman, T., et al. 2008b, Sci, 322, 1348
Marois, C., Zuckerman, B., Konopacky, Q. M., Macintosh, B., & Barman, T. 2010, Natur, 468, 1080
Matthews, B., Kennedy, G., Sibthorpe, B., et al. 2014, ApJ, 780, 97
Mawet, D., Hirsch, L., Lee, E. J., et al. 2019, AJ, 157, 20
Mawet, D., Milli, J., Wahhaj, Z., et al. 2014, ApJ, 792, 11
Nowak, M., Coroller, H. L., Arnold, L., et al. 2018, A&A, 615, A144
Pueyo, L., Soummer, R., Hoffmann, J., et al. 2015, ApJ, 803, 31
Racine, R., Walker, G. H., Nadeau, D., Doyon, R., & Marois, C. 1999, PASP, 111, 587
Rein, H., & Tamayo, D. 2015, MNRAS, 452, 376
Revels, J., Lubin, M., & Papamarkou, T. 2016, arXiv:1607.07892
Rhee, J. H., Song, I., Zuckerman, B., & McElwain, M. 2007, ApJ, 660, 1556
Ruffio, J.-B., Macintosh, B., Konopacky, Q. M., et al. 2019, AJ, 158, 21
Ruffio, J.-B., Macintosh, B., Wang, J. J., et al. 2017, ApJ, 842, 22
Ruffio, J.-B., Mawet, D., Czekala, I., et al. 2018, AJ, 156, 15
Sadakane, K., & Nishida, M. 1986, PASP, 98, 685
Sepulveda, A. G., & Bowler, B. P. 2022, AJ, 163, 163
Serabyn, E., Huby, E., Matthews, K., et al. 2017, AJ, 153, 7
Service, M., Lu, J. R., Campbell, R., et al. 2016, PASP, 128, 095004
Sivaramakrishnan, A., & Oppenheimer, B. R. 2006, ApJ, 647, 620
Skemer, A. J., & Close, L. M. 2011, ApJ, 730, 4
Skemer, A. J., Hinz, P. M., Esposito, S., et al. 2012, ApJ, 753, 14
Skemer, A. J., Marley, M. S., Hinz, P. M., et al. 2014, ApJ, 792, 17
Soummer, R., Hagan, J. B., Pueyo, L., et al. 2011, ApJ, 741, 55
Soummer, R., Pueyo, L., & Larkin, J. 2012, ApJ, 755, L28
Spiegel, D. S., & Burrows, A. 2012, ApJ, 745, 174
Su, K. Y. L., Rieke, G. H., Stapelfeldt, K. R., et al. 2009, ApJ, 705, 314
Sudol, J. J., & Haghighipour, N. 2012, ApJ, 755, 38
Thompson, W., & Marois, C. 2021, AJ, 161, 161
Wahhaj, Z., Cieza, L. A., Mawet, D., et al. 2015, A&A, 581, A24
Wahhaj, Z., Milli, Z., Romero, J., et al. 2021, A&A, 648, A26
Walker, G., Chapman, S., Mandushev, G., et al. 1999, ESOC, 56, 449
Wang, J., Wang, J. J., Ma, B., et al. 2020, AJ, 160, 160
Wang, J. J., Graham, J. R., Dawson, R., et al. 2018, AJ, 853, 156
Wang, J. J., Ruffio, J.-B., Morris, E., et al. 2021, AJ, 162, 162
Wertz, O., Absil, O., Gómez González, C. A., et al. 2017, A&A, 598, A83
Wilner, D. J., MacGregor, M. A., Andrews, S. M., et al. 2018, ApJ, 855, 56
Xu, K., Ge, H., Tebbutt, W., et al. 2020, in Proc. of The 2nd Symp. on Advances in Approximate Bayesian Inference, ed. C. Zhang, F. Ruiz, T. Bui et al. (London: PMLR), https://proceedings.mlr.press/v118/xu20a.html
Yelda, S., Lu, J. R., Ghez, A. M., et al. 2010, ApJ, 725, 331
Zahnle, K., Marley, M. S., Morley, C. V., & Moses, J. I. 2016, ApJ, 824, 137
Zuckerman, B., & Song, I. 2004, ApJ, 603, 738
Zurlo, A., Vigan, A., Galicher, R., et al. 2016, A&A, 587, A57