The Astronomical Journal, 163:253 (18pp), 2022 June https://doi.org/10.3847/1538-3881/ac64aa

The Perkins INfrared Exosatellite Survey (PINES) I. Survey Overview, Reduction Pipeline, and Early Results

Patrick Tamburo[1], Philip S. Muirhead[1], Allison M. McCarthy[1], Murdock Hart[1,2], David Gracia[1], Johanna M. Vos[3], Daniella C. Bardalez Gagliuffi[3], Jacqueline Faherty[3], Christopher Theissen[4,7], Eric Agol[5], Julie N. Skinner[1], and Sheila Sagear[6]
[1] Department of Astronomy & The Institute for Astrophysical Research, Boston University, 725 Commonwealth Avenue, Boston, MA 02215, USA
tamburop@bu.edu
[2] Perkins Telescope Observatory, 2780 Marshall Lake Road, Flagstaff, AZ 86005, USA
[3] American Museum of Natural History, 200 Central Park West, New York, NY 10024, USA
[4] Center for Astrophysics and Space Studies, University of California, San Diego, La Jolla, CA 91125, USA
[5] Department of Astronomy & Virtual Planetary Laboratory, University of Washington, Box 351580, U.W., Seattle, WA 98195, USA
[6] Department of Astronomy, University of Florida, Gainesville, FL 32611, USA
Received 2022 January 4; revised 2022 March 30; accepted 2022 April 4; published 2022 May 9

Abstract

We describe the Perkins INfrared Exosatellite Survey (PINES), a near-infrared photometric search for short-period transiting planets and moons around a sample of 393 spectroscopically confirmed L- and T-type dwarfs. PINES is performed with Boston University's 1.8 m Perkins Telescope Observatory, located on Anderson Mesa, Arizona. We discuss the observational strategy of the survey, which was designed to optimize the number of expected transit detections, and describe custom automated observing procedures for performing PINES observations. We detail the steps of the `PINES Analysis Toolkit` (`PAT`), software that is used to create light curves from PINES images. We assess the impact of second-order extinction due to changing precipitable water vapor on our observations and find that the magnitude of this effect is minimized in Mauna Kea Observatories *J* band. We demonstrate the validity of `PAT` through the recovery of a transit of WASP-2 b and known variable brown dwarfs, and use it to identify a new variable L/T transition object: the T2 dwarf WISE J045746.08-020719.2. We report on the measured photometric precision of the survey and use it to estimate our transit-detection sensitivity. We find that for our median brightness targets, assuming contributions from white noise only, we are sensitive to the detection of 2.5 $R_\oplus$ planets and larger. PINES will test whether the increase in sub-Neptune-sized planet occurrence with decreasing host mass continues into the L- and T-dwarf regime.

Key words: Surveys – Brown dwarfs – L dwarfs – T dwarfs – Transit photometry – Exoplanet astronomy – Infrared photometry

Supporting material: machine-readable table

1. Introduction

The L and T spectral types extend the traditional Harvard stellar classification system (Cannon & Pickering 1901) to include objects that are less massive, cooler, and spectroscopically distinct from the latest M dwarfs (Kirkpatrick et al. 1999; Burgasser et al. 2006). With effective temperatures ($T_{\rm eff}$) in the range from ∼700 to ∼2200 K (e.g., Nakajima et al. 2004; Vrba et al. 2004), L- and T-type dwarfs are faint at optical wavelengths, and only recently began to be detected in appreciable numbers, largely thanks to large-scale near-infrared (NIR) surveys like the DEep Near Infrared Survey of the Southern Sky (Epchtein et al. 1997), the Two Micron All Sky Survey (2MASS; Skrutskie et al. 2006), the UKIRT Infrared Deep Sky Survey (UKIDSS; Lawrence et al. 2007), and the space-based Wide-field Infrared Survey Explorer (WISE; Wright et al. 2010). Today, thousands of L- and T-type dwarfs are known. Studies of the initial mass function have found that though they are less common by number per unit volume than stars, L- and T-type dwarfs are still frequent outcomes of the process of star formation (Kroupa 2001; Chabrier 2005; Kirkpatrick et al. 2012, 2019).

The L spectral class spans the boundary between the stellar regime, where sustained hydrogen fusion maintains long main-sequence lifetimes, and the substellar regime of brown dwarfs (BDs) and planets, where the lack of sustained fusion causes objects to cool off over time, radiating away their heat of formation (Hayashi & Nakano 1963; Kumar 1963). It should be emphasized that the L and T taxonomy constitutes a spectral sequence, being defined by the evolution of observable spectral features across the classes; the L spectral class does not represent a demarcation between stars and BDs/planets. Rather, the spectral features that characterize early L dwarfs can equally be achieved in the atmospheres of stars ($\gtrsim$1 Gyr), BDs (∼0.01–1 Gyr), and planetary-mass objects ($\lesssim$0.01 Gyr; Burrows et al. 1997). Only by mid-L (∼L4) are objects unambiguously substellar.

Despite advances in the understanding of L and T dwarfs in recent decades, their short-period (P $\lesssim$200 days) planet population remains largely unconstrained. The occurrence rates of short-period planets around earlier spectral types have been studied in some detail with optical radial velocity (RV) and transit surveys (e.g., Johnson et al. 2010; Howard et al. 2012; Dressing & Charbonneau 2013, 2015; Mulders et al. 2015;

[7] NASA Sagan Fellow.

Hardegree-Ullman et al. 2019), but these surveys generally lack the sensitivity to detect planets around optically faint L and T dwarfs. The NIRSPEC Ultracool Dwarf Radial Velocity Survey (Blake et al. 2010) targeted a sample of 59 late-M/early L dwarfs at NIR wavelengths, but was designed to search for giant planets and BD companions. As a result, knowledge about sub-Neptune short-period L/T planet occurrence rates is currently limited to upper limits from Spitzer (He et al. 2017) and *K2* data (Sagear et al. 2020; Sestovic & Demory 2020). Individual *long*-period planets have been discovered around BDs through direct imaging (e.g., Chauvin et al. 2004; Close et al. 2007; Béjar et al. 2008; Luhman et al. 2009; Fontanive et al. 2020), gravitational microlensing (e.g., Han et al. 2013; Jung et al. 2018a, 2018b), and astrometry (Sahlmann et al. 2013), but those that have been confirmed have large planet-to-host mass ratios and orbital semimajor axes that are generally suggestive of formation through gravitational instability (e.g., Chauvin et al. 2005), rather than the core accretion mechanism that is thought to govern the in situ formation of short-period planets.

There is mixed evidence regarding the potential for planet formation around L and T dwarfs through core accretion. Analyses of Kepler data (Borucki et al. 2010) have revealed that short-period, sub-Neptune-sized planets occur more frequently with later host spectral types, with M-type stars hosting about three times as many sub-Neptunes as F-type stars (Howard et al. 2012; Mulders et al. 2015). M dwarfs are also frequently observed to host "compact multiple" systems, containing multiple sub-Neptune-sized planets on orbits with periods less than ∼10 days (Muirhead et al. 2015; Ballard & Johnson 2016). TRAPPIST-1, an M8 dwarf, is known to host seven transiting, Earth-sized exoplanets (Gillon et al. 2016, 2017). With a mass of 93 $M_{\rm Jup}$ (Van Grootel et al. 2018), TRAPPIST-1 is just slightly more massive than a field L or T dwarf and was one of only ∼50 systems targeted by the TRAPPIST Ultra-Cool Dwarfs Transit Survey, a prototype survey for SPECULOOS (Gillon et al. 2013a; Delrez et al. 2018). Such systems may also be common around L or T dwarfs, but a recent pebble accretion model from Mulders et al. (2021) predicts a turnover in the occurrence rate of super-Earths past the M spectral type, and that the formation of such planets may cease around objects less massive than 0.1 $M_{\odot}$.

Disks of gas and dust, the reservoirs from which planets form in the core accretion model, have been observed to be prevalent around BDs, occurring with similar frequencies and dissipation timescales as those around stellar-mass objects (Luhman et al. 2005; Luhman 2007; Monin et al. 2010; Dawson et al. 2013; Liu et al. 2015). In some cases, the masses of BD protoplanetary disks have been measured, with dust mass estimates that typically range from a fraction of an Earth mass to a few Earth masses (e.g., Klein et al. 2003; Scholz et al. 2006; Harvey et al. 2012; Broekhoven-Fiene et al. 2014; van der Plas et al. 2016; Rilinger et al. 2019). Rilinger & Espaillat (2021) measured masses for the largest sample of BD protoplanetary disks to date using spectral energy distribution (SED) modeling, and only 8 out of 49 disks in their sample possessed masses greater than 1 $M_{\rm Jup}$. Planet formation simulations suggest that these measured masses are generally insufficient to form Earth-mass planets through core accretion. Payne & Lodato (2007) performed core accretion modeling for planet formation around BDs and found that disks with masses less than 1 $M_{\rm Jup}$ rarely form planets more massive than 0.3 $M_{\oplus}$. Miguel et al. (2020) adapted a code built to study the formation of the Galilean satellites to planet formation in BD disks, finding that disks with masses greater than ∼10 $M_{\rm Jup}$ were required to form planets more massive than 0.1 $M_{\oplus}$. However, they note that their modeling fails to explain some of the larger known exoplanets around M dwarfs.

The tension between the competing lines of evidence presented above can only be resolved by searching for short-period planets around L and T dwarfs. With current technology, these planets are most amenable to detection using the transit method, as the diminutive radii of their hosts result in transit depths on the order of 1% for Earth-sized planets, signals that are accessible with 1–2 m class telescopes equipped with NIR photometers (Tamburo & Muirhead 2019). However, transit searches will also have to contend with short-duration transit events (down to 15 minutes, Delrez et al. 2018) and percent-level photometric variability on hours-long timescales that is frequently exhibited by L and T dwarfs (Radigan et al. 2014; Wilson et al. 2014; Metchev et al. 2015; Vos et al. 2019; Tannock et al. 2021).

One such transit survey, the Search for habitable Planets EClipsing ULtra-cOOl Stars (SPECULOOS), began in 2019 January and is looking for transiting planets around a sample of ∼1200 ultracool dwarfs (UCDs), of which ∼13.4% are L and T dwarfs (Delrez et al. 2018; Murray et al. 2020; Sebastian et al. 2021). SPECULOOS consists of a worldwide array of robotic telescopes (most of which are 1 m), which perform nightly photometric observations in a custom $I+z'$ filter.

In this work, we describe a new effort to search for transiting companions around L and T dwarfs at wavelengths longer than 1.0 μm: the Perkins INfrared Exosatellite[8] Survey (PINES). In Section 2, we provide an overview of the survey, describing the observing facility, target sample, and observing strategy. Section 3 details the custom photometric pipeline for analyzing PINES data. In Section 4, we quantify the survey's photometric performance and use this to estimate our transit-detection sensitivity. We also validate our pipeline's performance using observations of a transiting exoplanet and known variable L/T transition objects and identify a previously unknown variable.

2. The PINES Survey

2.1. Observing Facility and Instrument

PINES is conducted using Boston University's 1.8 m Perkins Telescope Observatory (PTO), located on Anderson Mesa, Arizona, at an altitude of 2206 m. The telescope is a Cassegrain reflector with a 0.4 m diameter secondary mirror. It is mounted on an English cross-axis equatorial mount originally built in the 1920s (Crump 1929).

In the survey, we perform photometric observations of L and T dwarfs using Mimir (Clemens et al. 2007), a NIR polarimeter, spectrometer, and imager that is permanently housed at the PTO. We use the instrument in its broadband *J*- and *H*-band wide-field imaging modes, which provide $10' \times 10'$ field-of-view (FOV) images. The FOV is large enough to provide several similarly bright reference stars in each target field, which are needed to measure and remove the effects from changing atmospheric conditions (Tamburo & Muirhead 2019). The *J*- and *H*-band filters are from the Mauna

[8] Because L and T dwarfs can themselves be planetary mass, the term "exosatellite" is used as a more general substitute for "exoplanet" in the survey title.

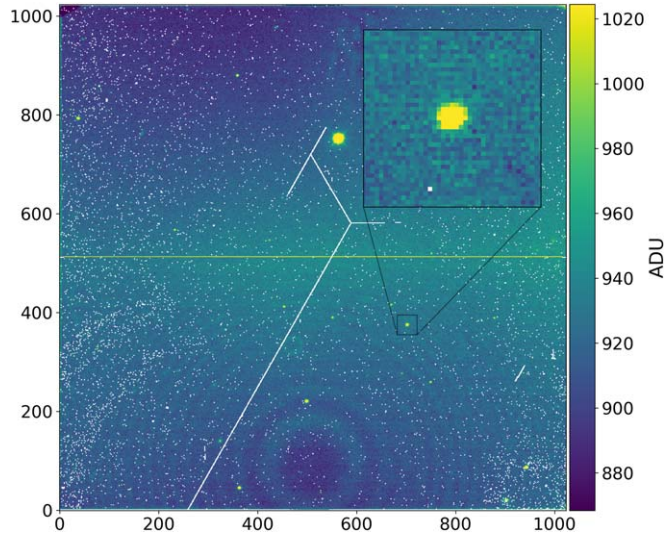

Figure 1. A typical reduced *J*-band Mimir image. Bad pixels in this image are shown in white. The inset shows a zoom-in on the target in this field, which is deliberately placed in a detector region that is relatively devoid of bad pixels.

Kea Observatories near-infrared (MKO-NIR) set, which was designed in part to avoid contamination from water vapor lines (Simons & Tokunaga 2002; Tokunaga et al. 2002).

The Mimir detector is an ALADDIN III InSb 1024 × 1024 array detector, which is operated at 33.5 K with the help of a Trillium 1050 two-stage, closed-cycle helium cooling system. A typical reduced Mimir image is shown in Figure 1. On average, we find that 2.7% of the pixels are either significantly hot, dark, or variable, and are thus unsuitable for precise photometry; these pixels are marked in white in Figure 1, and we detail how they were identified in Section 3.2. Bad pixels are spread more-or-less homogeneously across the chip, except for in a physical crack that runs through all four detector quadrants. Additionally, Figure 1 shows an interference pattern in the bottom two quadrants, likely due to sky background interference within one or several refractive optical elements. The resulting interference pattern causes a series of concentric rings of alternating brightness in the lower two quadrants. We modeled the ring pattern to remove it from the reduced images but found that doing so had minimal impact on the resulting light curve quality.

2.2. Target Sample

We assembled the PINES target sample using an online database of UCDs maintained by J. Gagné,[9] which incorporates targets listed in Mace (2014), Dupuy & Liu (2012), and the DwarfArchives.org catalog. All objects in the list have either an optical or an NIR spectral type of L0 or later. We selected targets from this list that reside between the decl. limits of the PTO ($-10° < \delta < +70°$) and have an apparent 2MASS magnitude of $m_J < 16.5$. Known unresolved binaries were also removed from the list. In total, the PINES sample consists of 393 spectroscopically confirmed L- and T-dwarf targets, the properties of which are given in Table 1.

Histograms of the magnitudes and spectral types of targets are shown in Figure 2. The full 393-target sample is shown as a dashed line, while the targets that have been observed as of 2021 October are filled in. Because the PINES sample is magnitude-limited, it is biased toward early-type L dwarfs, with 73% of our targets being L3 or earlier. The median 2MASS *J*-band apparent magnitude of the sample is $m_J = 15.1$, and the median spectral type is L1.

A sky map of the sample is shown in Figure 3. PINES targets are distributed uniformly across the sky, except near R.A. 90° and 300°, where the galactic plane inhibits the detection of faint L and T dwarfs. Targets in the sample were divided into groups using a clustering algorithm, which are represented by the solid lines in Figure 3. We designed these groups such that the targets in each were separated by less than 15° on the sky, to reduce average slew times between group members.

2.3. Observational Strategy

PINES uses the observing strategy detailed in Tamburo & Muirhead (2019), in which a group of four to seven targets is repeatedly observed throughout an observing night (see the inset in Figure 3). Each target is observed for an equal amount of time, and cycles are designed to take a maximum of one hour to capture short-duration transit events. This results in noncontinuous time coverage for PINES targets, but transits can be detected by significant decrements of individual *blocks* of data. Groups of targets are scheduled for five nights of observations, which was found to maximize the number of observed transits in Tamburo & Muirhead (2019). This duration heavily biases the survey toward detecting short-period planets, and the same simulation showed that the vast majority of transits will belong to planets with orbital periods of 10 days or less.

The strategy of cycling between multiple targets has been employed successfully in the search for transiting planets by other surveys (e.g., MEarth, Nutzman & Charbonneau 2008), but as noted in Delrez et al. (2018), a staring procedure (as employed by the SPECULOOS survey) maximizes the signal-to-noise ratio. This is indeed a limitation of cycling observations, as slews interrupt time that would otherwise be spent obtaining exposures. Additionally, cycling observations are at greater risk of systematics from flat-fielding errors, as the fields are generally not placed on the exact same pixels following the slews. In the following section, we discuss our efforts to address the limitations of the cycling strategy.

2.4. Guiding Performance

The PTO has a built-in, off-axis auto-guider, which can be used to keep sources on the same pixels over a set of exposures. Stable field positioning is important for accurate IR photometry, as image motions coupled with flat-fielding errors and intrapixel sensitivity variations can inject spurious signals into the final light curves (e.g., Ingalls et al. 2012). However, the existing auto-guider is inadequate for PINES observations because of the time overhead that it introduces. Following each slew, the observer must shift the field to the desired detector position, a task that can take several minutes to complete. Next, they must identify a suitable guide star and initiate the guiding system, which also requires about a minute of active user involvement. Because the PINES cycling strategy necessitates dozens of slews over the course of a night, the use of the existing PTO auto-guider would be prohibitively time-consuming and prone to human error. For this reason, we designed a custom guiding procedure for PINES that is quick, automatic,

[9] https://jgagneastro.com/list-of-ultracool-dwarfs/

Table 1
Targets in the PINES Sample First by Group Number and Then by R.A. (Full Table Available Online)

2MASS Name	PINES Group ID	R.A. (J2000)	Decl. (J2000)	2MASS m_J	2MASS m_H	SpT Optical	SpT NIR
00001354 + 2554180	0	00:00:14	+25:54:17	15.063 ± 0.039	14.731 ± 0.074	⋯	T4.5
00154476 + 3516026	0	00:15:45	+35:16:02	13.878 ± 0.028	12.892 ± 0.035	L2	L1
00193275 + 4018576	0	00:19:33	+40:18:54	15.544 ± 0.057	14.928 ± 0.068	⋯	L2
00250365 + 4759191	0	00:25:04	+47:59:19	14.840 ± 0.036	13.667 ± 0.030	L4	⋯
00144919 − 0838207	1	00:14:50	−08:38:22	14.469 ± 0.026	13.950 ± 0.026	L0	M9
00191165 + 0030176	1	00:19:12	+00:30:18	14.921 ± 0.035	14.180 ± 0.040	L1	L0.5
00242463 − 0158201	1	00:24:25	−01:58:17	11.992 ± 0.033	11.084 ± 0.020	M9.5	L0.5
00261147 − 0943406	1	00:26:11	−09:43:40	15.601 ± 0.067	14.999 ± 0.081	L1	⋯
00320509 + 0219017	1	00:32:05	+02:19:01	14.324 ± 0.023	13.386 ± 0.023	L1.5	M9
00345684 − 0706013	1	00:34:57	−07:06:00	15.531 ± 0.059	14.566 ± 0.041	L3	L4.5
00011217 + 1535355	2	00:01:12	+15:35:36	15.522 ± 0.061	14.505 ± 0.052	⋯	L4
00302476 + 2244492	2	00:30:25	+22:44:47	14.586 ± 0.036	13.975 ± 0.047	⋯	L0.5
00304384 + 3139321	2	00:30:44	+31:39:31	15.480 ± 0.052	14.617 ± 0.057	L2	L3
00361617 + 1821104	2	00:36:16	+18:21:10	12.466 ± 0.025	11.588 ± 0.028	L3.5	L4
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮

(This table is available in its entirety in machine-readable form.)

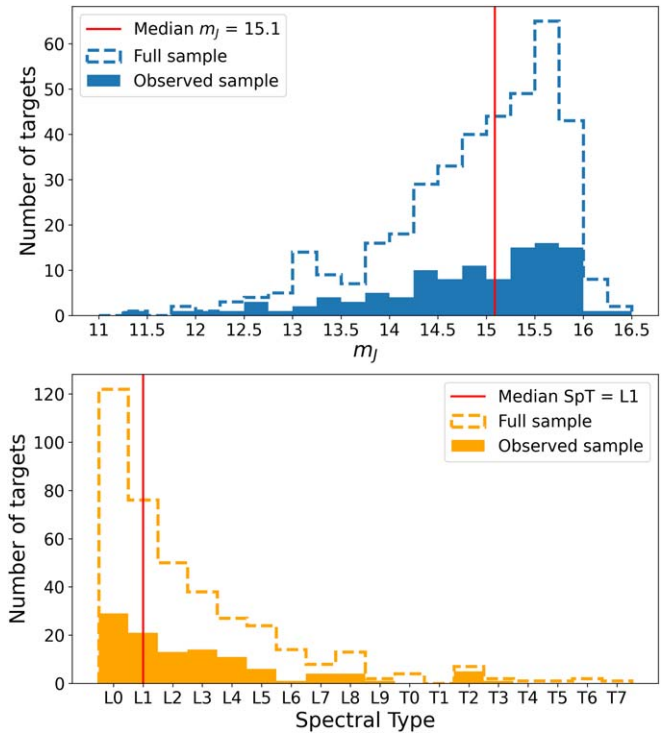

Figure 2. Top: histogram of apparent *J*-band magnitudes of targets in the PINES sample. Bottom: histogram of spectral types.

and places targets on roughly the same pixels throughout observations.

At the beginning of each observing run, we position each field such that the target is placed on a region of the detector that is known to be free of bad pixels (see Figure 1). We acquire a "guide field" image at this location, and detect sources in this image with the `DAOStarFinder` routine from the `photutils` Python package (Bradley et al. 2020). We use the detected source positions to create a synthetic guide field image, which consists of 2D Gaussians at the measured source positions (see Figure 4). We use the synthetic guide field image in subsequent positioning calculations to measure the shifts that the telescope needs to execute to achieve the guide field position. A synthetic image is used to prevent image features like the crack seen in Figure 1 from dominating the fast Fourier transform (FFT) convolution with subsequent images.

When a new science image is saved, the guiding program first determines which field is being imaged by examining the telescope's reported R.A. and decl. coordinates. It then performs a source detection on the newly written image and creates a synthetic science field image using the measured source positions. Subsequently, an FFT convolution is performed between the synthetic science field image and the synthetic guide field image, an example of which is shown in Figure 5. This figure shows many locations where the cross-correlation power is moderate, which is expected and is the result of image shifts that align a single pair of sources in the science image and the guide image. However, there is a clear peak in the cross-correlation power at +1.9 pixels in x and −5.5 pixels in y. This is the shift that, when applied to this particular science image, would align it with the guide field position. The measured shifts are passed to the telescope to adjust its pointing back to the field's guide position. The entire process of measuring image shifts and adjusting the telescope position takes only a few seconds, so the corrections are passed following every science image.

In Figure 6, we show a comparison between the performance of the PTO auto-guider and the PINES guiding system. The top two rows show the x and y shifts (in pixels) between the measured field position and the guide position for about four hours on two different observing nights. On the first night, the target was observed with the PTO auto-guider, while on the second, it was positioned using the PINES guiding system. The PINES guiding observations have worse x pointing precision, with a standard deviation of 1.3 pixels compared to 0.3 using the auto-guider. The y pointing precision is comparable between the two guiders, however. We attribute the slightly worse pixel positioning from PINES guiding to the fact that pointing corrections are only issued on image readout (i.e., the correction frequency is set by the image exposure time, which is typically $\geqslant$30 s), whereas the PTO auto-guider passes corrections to the telescope about every second. Despite the worse pointing performance, the light curves created using the two guiding approaches are essentially the same, as

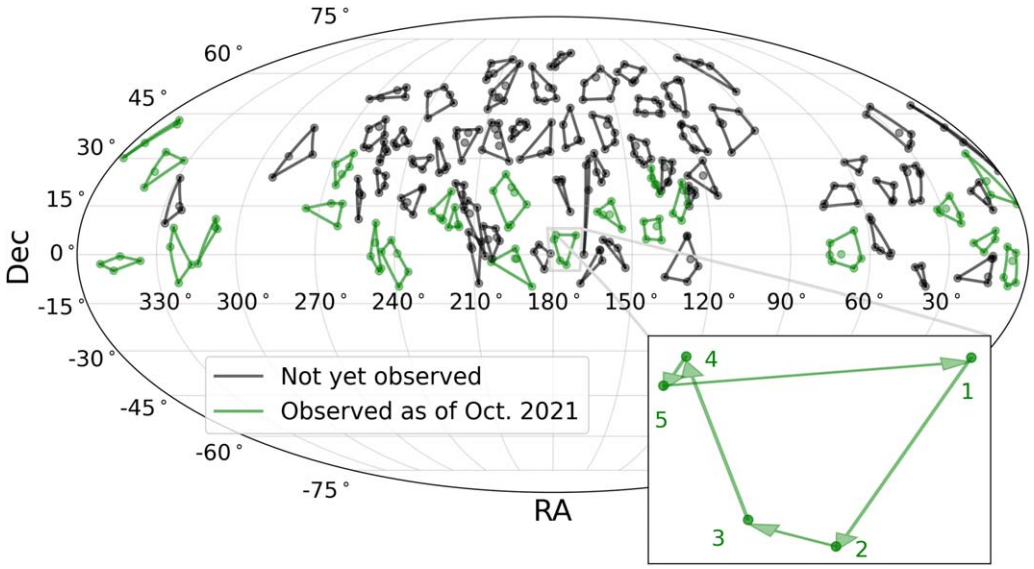

Figure 3. A sky map of the PINES sample. Targets that have been observed as of 2021 October are shown in green, while those that have yet to be observed are shown in black. Individual targets are shown as filled circles, while groups of targets are outlined with solid lines. The inset shows a zoom-in on one group and visualizes the order in which its members were targeted during PINES observations.

demonstrated by the standard deviations of the light curves in the bottom row of Figure 6.

2.5. Log Files

We automatically log the file name, date, target name, filter, exposure time, airmass, x/y pixel shifts, and seeing FWHM of each image. The seeing estimate is given by the average FWHM of 2D, circular Gaussian functions that are fit at the locations of detected sources in the image. These log files provide supporting information for the purposes of centroiding (Section 3.5) and performing time-variable aperture photometry (Section 3.6).

3. The PINES Photometric Pipeline

We process PINES data using a custom photometric pipeline to create light curves that can be searched for potential transit events. We call the pipeline the `PINES Analysis Toolkit` (`PAT`). We designed `PAT` to handle the peculiarities of Mimir data that, if unaccounted for, would result in systematic errors in the final target light curves. It was also designed to function as automatically as possible, to facilitate the analysis of the large quantity of survey data. The software is available on GitHub[10] under an MIT License and version 1.0.0 is archived on Zenodo (Tamburo 2022). We detail the major steps of the pipeline in the following sections.

3.1. Calibration Images

Calibrations are typically taken once per run and consist of flat fields and dark exposures. Flat fields are obtained with a lamp-illuminated dome flat screen, with 100 lamp-on and 100 lamp-off images obtained in the same photometric filters as the science images. Combined lamp-on and lamp-off images are created by looping through each pixel, sigma clipping the 100 brightness measurements for that pixel with a 3σ clipping threshold, and taking the mean of the remaining measurements. We then construct a flat field by subtracting the combined lamp-off image from the combined lamp-on image, the values of which are then median-normalized to 1.

Darks are constructed in a similar fashion, with sets of 10–20 individual dark exposures obtained using exposure times that match those of the science images. The same sigma-clipping procedure used to create the flats is used to average the individual dark exposures together into a dark at each exposure time. A *dark standard-deviation* image is also produced for each exposure time, consisting of the standard deviation of the sigma-clipped values for each pixel. These capture the typical brightness fluctuations of detector pixels and can be used for bad-pixel identification.

3.2. Bad-pixel Identification and Data Reduction

The flats and darks are used to identify variable, hot, and dead pixels on the detector. These pixels are spread more-or-less uniformly across the chip and produce unreliable flux measurements and would result in inaccuracies if included in a photometric aperture or background annulus. We therefore mask out these pixels and correct their values at a later stage in the pipeline.

First, a variable pixel mask is created for each science exposure time using the dark standard-deviation images. Pixels with high values in these images vary significantly between one exposure and the next. Pixels that have standard deviations over five times the average standard deviation of the dark standard-deviation image are flagged as variable.

We then identify hot pixels using the dark images. Each pixel in a dark image is iterated over, and its value is compared to a box of its neighboring pixels. The pixel is flagged as hot if its value is significantly higher than the median of the neighboring pixels. Surrounding pixels can themselves be hot, in which case they will raise the standard deviation of the box and possibly prevent the identification of a hot pixel. The hot pixel identification is therefore performed as an iterative process, with the previously identified hot pixels being excluded from subsequent calculations. Hot pixels are identified with a given box size and significance threshold

[10] `PINES Analysis Toolkit` codebase: https://github.com/patricktamburo/pines_analysis_toolkit.

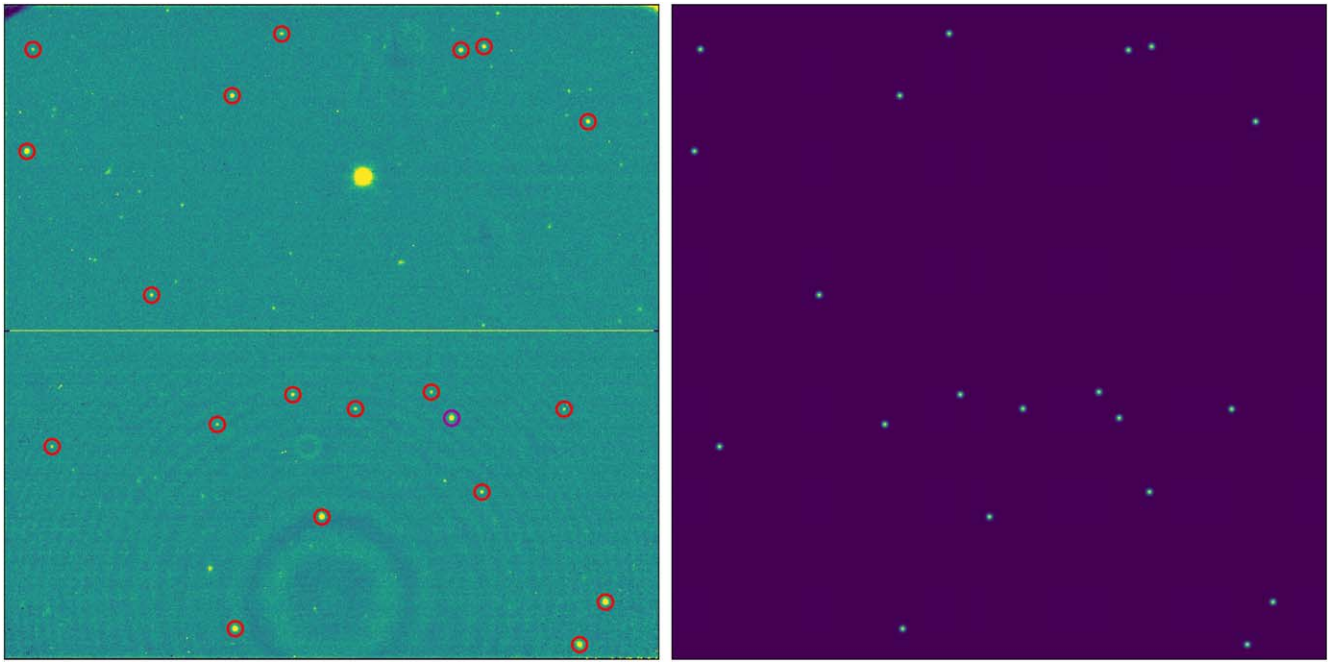

Figure 4. Left: an example of a guide field image, with guide stars indicated with circles (with purple indicating the target and red indicating other sources). Bad pixels in this image have been interpolated with a 2D-Gaussian kernel to improve source-detection efficiency. Right: the guide field synthetic image for this field.

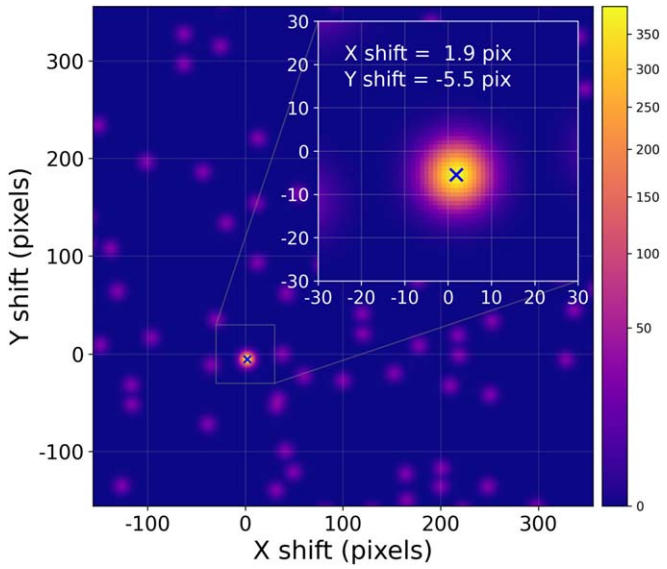

Figure 5. The FFT convolution of a science field synthetic image and a guide field synthetic image, which is used to determine the shifts needed to correct the PTO's pointing during PINES observations.

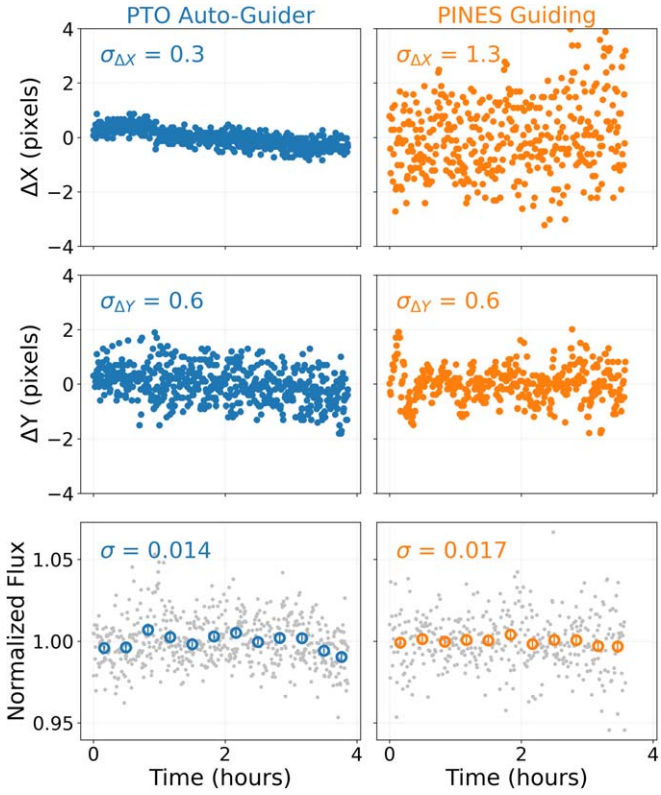

Figure 6. Comparison of the PTO auto-guider performance (left column) vs. the PINES guiding system (right column). Top: measured x pixel shifts. Middle: measured y pixel (decl. axis) shifts. Bottom: light curves resulting from the two guiding approaches, which exhibit similar scatter.

until no more are found, at which point the box is shrunk and the process is repeated. Through testing, we found that the following combinations of box sizes and significance thresholds successfully flag almost all visible hot pixels: 13×13 pixels/10σ, 11×11 pixels/10σ, 9×9 pixels/10σ, 7×7 pixels/9σ, and 5×5 pixels/8σ.

Dead pixel masks are created using the same iterative shrinking box approach, but flagging pixels if they have values significantly less than that of their neighbors in the flat images. Finally, we create a bad-pixel mask (BPM) for each exposure time and filter combination using the variable, hot, and dead pixel masks. We dark subtract and flat field the data to reduce the raw science images, ignoring the flagged pixels in the BPM.

3.3. Target and Reference Identification

The reduced science images are used in all of the following steps of the pipeline, beginning with the identification of sources in the field that are used for photometry. In addition to

the target L or T dwarf, we perform photometry on several reference stars (requiring at least three in every frame), which we use to remove the flux changes that are common to all objects in the field. To choose a set of suitable reference stars, we first generate a list of sources in a user-chosen science image using the `DAOStarFinder` program from `photutils`. We identify the target from this list of sources based on its location on the detector, which is known a priori from the positioning of the field during observations (see, e.g., Figure 1). Suitable reference stars are then chosen, avoiding the following:

1. Those that have counts in the nonlinear regime of the Mimir detector (with pixel values $\gtrsim$4000 ADU, Clemens et al. 2007).
2. Those that are much fainter than the target, and would introduce unnecessary noise to the target light curve. Typically, we set the dimness threshold to 0.3× the target's flux in the source-detection image.
3. Those that are near the edge of the detector, where shifts from guiding could move them off the chip.
4. Those that are located close to other sources ($\lesssim$10 pixels).

On average, there are 24^{+10}_{-6} reference stars that meet these criteria in each field ($16th - 84th$ percentile range). The large number of suitable reference stars generally allows us to perform differential photometry on our targets without inflating the scatter of the target's light curve, except for our brightest targets (see Section 4.1).

3.4. Astrometric Solution and Source Spectral Types

Mimir images do not have a built-in World Coordinate System (WCS) that can map x and y pixel coordinates to R.A. and decl. coordinates. We use a web application programming interface (API)[11] to pass the source-detection image to Astrometry.net (Lang et al. 2010), which solves for the WCS of the field. With the WCS data, we then convert the pixel locations of reference stars to R.A. and decl. coordinates. For every reference star in the field, we perform a $5'' \times 5''$ box search in Gaia Early Data Release 3 (Gaia Collaboration et al. 2021) around its coordinates using the `Astroquery` Python package (Ginsburg et al. 2019), and save its ($G_{BP} - G_{RP}$) color. We also calculate the absolute G magnitude of the reference using its measured parallax, and discard references that are likely white dwarfs or giants. A sufficient number of dwarf stars are present in every field such that we can discard white dwarfs and giants as references without inflating the scatter of our final light curves, and we do so as a precaution against the systematic effects that they could introduce through photometric variability or spectral mismatches with our targets. Finally, we use the ($G_{BP} - G_{RP}$) color to estimate the spectral type (SpT) and T_{eff} of each reference star using a ($G_{BP} - G_{RP}$)-SpT relation, limited to M9V stars and earlier (Pecaut & Mamajek 2013).[12]

3.5. Centroiding

Next, we measure the centroid locations of the target and reference stars in every image. For efficiency, we measure source centroids using small image cutouts (typically 16×16 pixels), which we create using the positions in the source-detection image and the pixel shifts from the observing logs. If there are known bad pixels in the source cutouts, their values are replaced with a 2D-Gaussian convolution to improve the centroiding performance (note that this replacement is not performed in the reduced image itself). We then apply a 2D-Gaussian convolution to the entire cutout, which smooths out any remaining spurious pixel values (from, e.g., cosmic rays or unidentified bad pixels). We measure the centroid on the smoothed cutout with the `centroid_2dg` function from `photutils`. The measured centroid in the cutout is then translated back into the full-image pixel coordinates.

3.6. Aperture Photometry and Bad-pixel Correction

We perform simple aperture photometry with both fixed and time-variable apertures. We do fixed aperture photometry using circular apertures with radii of 3.0, 4.0, 5.0, 6.0, and 7.0 pixels, which encompass the range of typical FWHM seeing values at the PTO (∼2″–3″.5; the Mimir plate scale in wide-field imaging mode is 0″.579 pixel^{-1}). We also do time-variable aperture photometry, with aperture radii equal to 0.5×, 0.75×, 1.0×, 1.25×, and 1.5 × a smoothed trend of the seeing FWHM.

On average, we find that 2.7% of the detector pixels are flagged as hot, dead, or variable in the BPM. Therefore, before we calculate the flux within an aperture, we check if it contains any known bad pixels. If it does, we fit a 2D-Gaussian function to a cutout centered on the source position and replace the bad pixels with the values from the fitted model. We purposefully position the target in a detector region that is mostly devoid of bad pixels, and as such it rarely requires a bad-pixel correction. The reference stars, however, are frequently placed over bad pixels, and the values of these pixels have to be modeled to perform aperture photometry. The other possibility, of discarding the reference stars that fall on bad pixels, is not feasible, as it would eliminate most (if not all) of the potential reference stars in any given field.

We performed the following test to gauge the performance of our 2D-Gaussian replacement procedure. We identified a source in one of our fields that has no bad pixels under its point-spread function. We found the pixel that was nearest to the source centroid position on average, and calculated σ_i, its expected noise in frame i, using the photon noise from the star and sky background, dark current, and read noise. We then created a new set of images, adding random noise to the value of the *target* pixel by drawing from a normal distribution with $\sigma = \sigma_i$. We performed aperture photometry on this data set using a time-variable aperture with radius 1× a smoothed trend of the seeing FWHM. We then created another new set of images, this time flagging the target pixel as bad and replacing its value with our 2D-Gaussian fitting procedure. We performed photometry on these images with the same aperture size.

Taking the ratio of the flux values measured from these two data sets, we found an average value of 1.0005 ± 0.0029. This suggests that our 2D-Gaussian replacement procedure produces values that are consistent with the noise sources in our data. We experimented with a second-order correction to our 2D-Gaussian approach, using the averaged residuals from Gaussians fit to all sources in an image to estimate updated replacement-pixel values. However, we found no evidence of

[11] http://astrometry.net/doc/net/api.html

[12] Updated relations including the Gaia photometric bands are available online at http://www.pas.rochester.edu/~emamajek/EEM_dwarf_UBVIJHK_colors_Teff.txt.

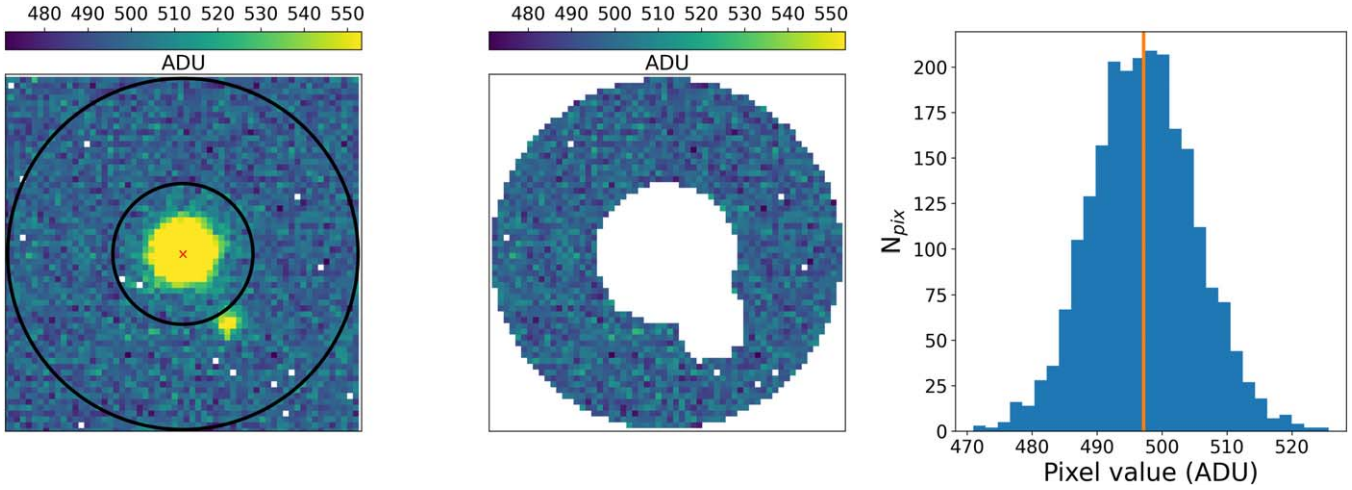

Figure 7. The background estimation process used for aperture photometry. Left: a circular annulus is placed at the measured source position (red x), with bad pixels shown in white. Middle: pixels outside of the annulus (and if necessary, those containing sources) are ignored. Right: a histogram of the resulting pixel values. The background is estimated as the sigma-clipped mean of these values (orange line).

improvement over the first-order 2D-Gaussian replacement, which is not surprising, as the previous test demonstrated that this approach is already consistent with the noise in the data.

With the bad-pixel values modeled with our 2D-Gaussian procedure, the flux within the target aperture is summed up. We then estimate the contribution from sky background through a process that is illustrated in Figure 7. First, a circular annulus centered on the source position is placed, with an inner radius of 12 pixels, and an outer radius of 30 pixels. Pixels outside of the annulus are ignored, as are the pixels inside the annulus that contain any nearby sources (this is done with the `photutils` function `make_source_mask`). The remaining pixels are then sigma clipped with a threshold of 4σ to remove the outliers. We then take the mean of the sigma-clipped values as the representation of the per-pixel sky background. This value, multiplied by the source aperture area, is subtracted from the source aperture sum to derive the background-corrected source flux.

We tested for spatial inhomogeneities over scales comparable to the sky annulus by splitting 60×60 pixel source cutouts (like the one shown in the left-hand panel of Figure 7) into 15×15 pixel subregions, and compared the background value that we measure with the annulus approach described above versus the sigma-clipped means of the individual subregions (masking out the central source). In the several thousand cutouts that we tested, we found no evidence for a significant difference between the background value measured in the annulus versus the background measured in the subregions. We therefore conclude that spatial inhomogeneities in the background across the annulus and aperture should not be a concern for our background correction procedure.

3.7. Light Curves

Finally, we create light curves using the background-subtracted fluxes from the target and reference stars. To remove trends that are common to all sources in the field due to changing observing conditions, an artificial light curve (ALC) is created using the weighted mean of reference star fluxes, in a procedure analogous to the one used by SPECULOOS (Murray et al. 2020). This process is intended to variable reference stars that could otherwise introduce spurious variability to the ALC-corrected target light curve.

First, we restore the raw source fluxes from one of the aperture options described in Section 3.6. We normalize the raw target flux ($F_{\rm T}$) using the weighted mean of the time series:

$$\hat{F}_{\rm T} = \frac{F_{\rm T}}{\left[\frac{\sum_{j=1}^{N_{\rm frames}} F_{{\rm T}_j} W_{{\rm T}_j}}{\sum_{j=1}^{N_{\rm frames}} W_{{\rm T}_j}}\right]}. \quad (1)$$

Here, $W_{{\rm T}_j} = 1/\sigma_{{\rm T}_j}^2$, where $\sigma_{{\rm T}_j}$ is the uncertainty on the target flux in frame j, which is calculated using photon noise from the target, sky background, dark current, and read noise.

Next, the measured flux from the ith reference star ($F_{{\rm R}_i}$) is read in, and normalized in the same way:

$$\hat{F}_{{\rm R}_i} = \frac{F_{{\rm R}_i}}{\left[\frac{\sum_{j=1}^{N_{\rm frames}} F_{{\rm R}_{ij}} W_{{\rm R}_{ij}}}{\sum_{j=1}^{N_{\rm frames}} W_{{\rm R}_{ij}}}\right]}. \quad (2)$$

In this expression, $W_{{\rm R}_{ij}} = 1/\sigma_{{\rm R}_{ij}}^2$, where $\sigma_{{\rm R}_{ij}}$ is the uncertainty on the flux from the ith reference star in frame j. A weighted mean of the $\hat{F}_{{\rm R}_i}$ is used to construct the ALC:

$$ALC = \frac{\sum_{i=1}^{N_{\rm refs}} \hat{F}_{{\rm R}_i} W_i}{\sum_{i=1}^{N_{\rm refs}} W_i}. \quad (3)$$

W_i is the total weight given to reference star i, the calculation of which is detailed in Appendix. Assuming for the moment that the W_i values are known and the ALC has been created, the corrected, differential target flux is then given simply by the following:

$$\hat{F}_{\rm T}^{\star} = \frac{\hat{F_T}}{ALC}. \quad (4)$$

Next, we correct $\hat{F}_{\rm T}^{\star}$ for linear correlations with airmass, seeing, and the target's centroid x and y pixel positions. We also include the target's intrapixel centroid location as a regressor, as NIR-array detectors can exhibit sensitivity variations across single pixels that can impact differential photometry at the $\sim$1% level (e.g., Lauer 1999; Ingalls et al. 2012). Before the regression correction is performed, we calculate the Pearson correlation coefficient between the target flux and each regressor, along with the two-tailed p-value,

which tests for noncorrelation. If the regressor's p-value is less than 0.01, it is retained in the regression; otherwise, it is discarded. We then perform a linear regression fit between the significantly correlated regressors and the corrected differential target flux and divide the resulting fit from the target flux. Optionally, a linear or quadratic function can be fit to the corrected target flux to remove any remaining long-term trends throughout the night, but these options are turned off by default.

We create a target light curve for every photometry file described in Section 3.6. We then choose the best aperture by minimizing the average standard deviation of the corrected target flux evaluated over the duration of individual blocks of data. This is done for both nightly normalized light curves, where each night of data is normalized to 1, and globally normalized light curves, where the entire data set is normalized to 1. The globally normalized light curves allow us to search for variability trends with longer time baselines.

3.8. Precipitable Water Vapor Corrections

A known limitation to the accuracy of differential photometry in the NIR is *second-order extinction*. Because the extinction coefficient depends on wavelength, stars with different SEDs will experience different amounts of wavelength-integrated extinction (Bailer-Jones & Lamm 2003). By design, PINES targets are typically much redder than the reference stars in any given field, and as such are potentially susceptible to second-order extinction effects.

This does not pose an issue if the extinction is stable in time. Unfortunately, at NIR wavelengths, the extinction is primarily driven by precipitable water vapor (PWV), which is notoriously time variable. Changing PWV levels can induce signals on the order of 1% in target light curves at NIR wavelengths, and they can change rapidly enough throughout the night to mimic transit events (Blake et al. 2008; Baker et al. 2017; Murray et al. 2020).

For this reason, we investigated the potential effect of second-order extinction on PINES observations, by examining the change in flux in the Mimir MKO J- and H-band filters as PWV levels are increased for different spectral types. We obtained NIR water-vapor-absorption spectra from ATRAN[13] (Lord 1992) ranging from 0 to 10 mm of PWV, and used low-resolution ($R = 100$) solar-metallicity BT-Settl models (Allard 2014) with $T_{\rm eff}$ ranging from 1500 to 6000 K to model the stars (1500 K is approximately the temperature of an L6 dwarf; 90% of our targets have SpTs of L6 or earlier). For each combination of stellar spectrum, PWV spectrum, and filter bandpass, we interpolated onto a common wavelength grid, multiplied the three together, and integrated over wavelength to measure the total response. We then calculated ΔF for each spectral type, the difference (in percent) between its normalized total response at a given PWV level and its response at a PWV level of 0 mm.

The results of this analysis are shown in panels (a)–(e) in Figure 8, along with calculations for three other NIR filters: 2MASS J, 2MASS H, and a model of the custom $I+z'$ filter used by the SPECULOOS transit survey (Delrez et al. 2018). The colored lines show the response of the different spectral types to changing PWV in the respective band. Panels (a)–(e) are shown with the same y-range, which reveals significant differences in the *absolute* flux change of sources in different

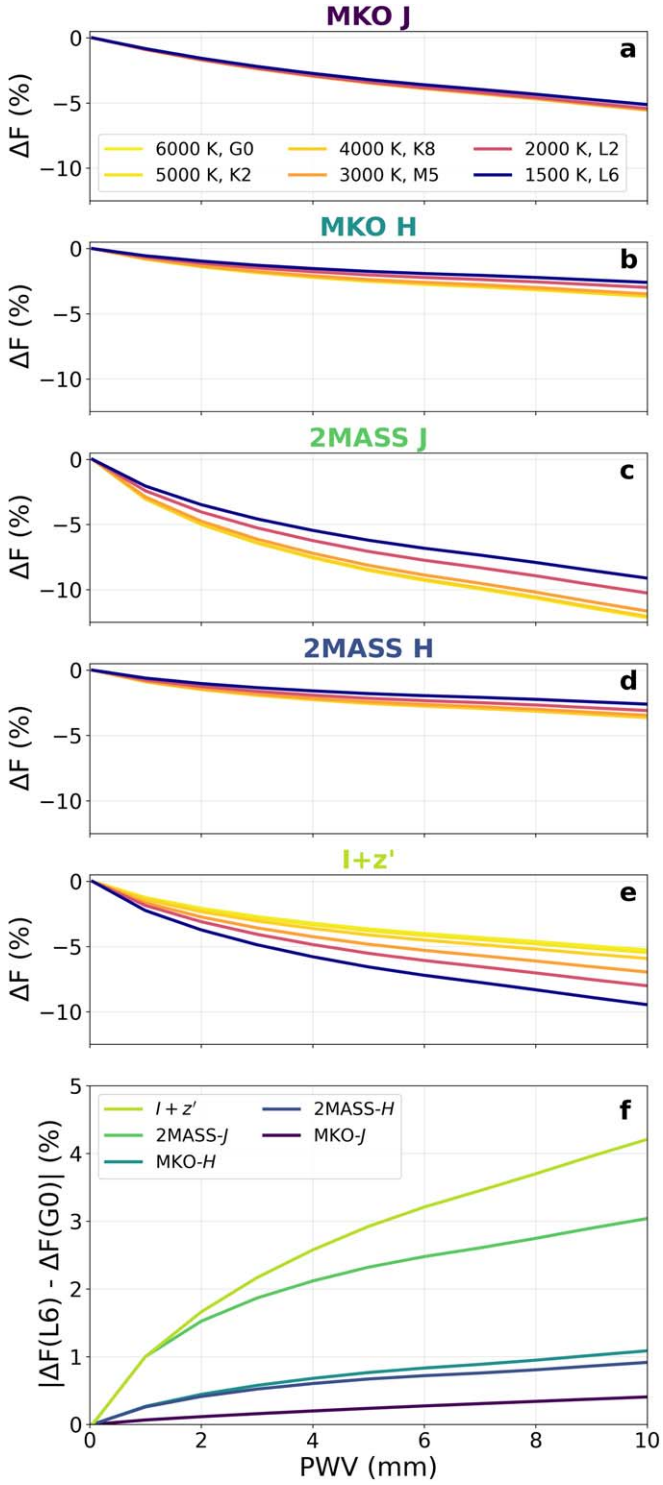

Figure 8. Panels (a)–(e): theoretical response to changing PWV levels for different spectral types in various NIR filters. Panel (f): the magnitude of the PWV response of an L6 dwarf calibrated to the response of a G0 in each filter. The second-order extinction effect is minimized in MKO J band.

NIR bands. As PWV levels increase from 0 to 10 mm, sources in MKO J band will experience an average flux decrease of 6%, while those in MKO H band will decrease by just 3%. In 2MASS J, 2MASS H, and $I+z'$, sources will undergo average flux decreases of 11%, 3%, and 7%, respectively. We note that our $I+z'$-band results closely reproduce those from Murray et al. (2020).

[13] https://atran.arc.nasa.gov/cgi-bin/atran/atran.cgi

However, the relevant metric for second-order extinction is the *relative* flux changes of sources in these various bands. This is illustrated in panel (f) of Figure 8, which shows the magnitude of the difference in the response of an L6 target with that of a G0 reference star versus PWV. This panel demonstrates that second-order extinction is minimized in MKO *J* band, with a flux difference of just 0.4% between an L6 and G0 over a change of 10 mm in PWV. In other bands, changing PWV can result in much larger flux mismatches: the effect is 0.9% in 2MASS *H* band, 1.1% in MKO *H* band, 3.0% in 2MASS *J* band, and 4.2% in $I+z'$ band. This result suggests that MKO *J* band is the best choice for minimizing systematic flux variations from second-order extinction effects, mainly because the filter was designed to avoid water-vapor-absorption lines (Simons & Tokunaga 2002; Tokunaga et al. 2002). For this reason, we have elected to perform PINES observations exclusively in MKO *J* band, contrary to the MKO *H*-band recommendation in Tamburo & Muirhead (2019).

While the scatter of our *J*-band light curves (see Section 4.1) is generally higher than the magnitude of potential second-order extinction effects, we implemented an optional procedure for correcting PINES light curves for changing PWV. Because the PTO does not have a direct way to measure PWV on-site, we obtain PWV estimates using the `Fyodor` Python package (Meier Valdés et al. 2021). `Fyodor` utilizes publicly available vertical temperature and moisture profiles from Geostationary Operational Environmental Satellites (GOES) imaging data, which can be used to calculate PWV levels through the line of sight to a target. Typically, we find PWV changes of only a few mm throughout observing nights with `Fyodor`. With these values and the spectral types of our sources (see Section 3.4), we can estimate the expected response of stars in PINES images to changing PWV levels throughout observations and correct for it. However, because downloading GOES data is time-intensive and we expect the magnitude of second-order extinction to be small in MKO *J* band, this procedure is not enabled in `PAT` by default; instead, we use it as a check on light curves that show transit or variability signatures. Additionally, because the spectral types of our reference stars are known (see Section 3.4), the user can also limit the ALC creation to using only the reddest reference stars, instead of doing a full PWV correction.

4. Photometric Performance and First Results

We have applied `PAT` to observations of 83 PINES targets in *J* band, to date. In Section 4.1, we provide an assessment of the noise performance of the survey using the light curves of these targets. We also provide examples that validate the performance of `PAT` for the purpose of transit detection (Section 4.2) and the identification of variable L and T dwarfs (Section 4.3). Finally, in Section 4.4, we identify a new variable L/T transition object: the T2 dwarf WISE J045746.08-020719.2.

4.1. Photometric Performance

In Figure 9, we show the average standard deviation of the *J*-band light curves analyzed with `PAT` over 5 minutes intervals ($\bar{\sigma}_{5-\mathrm{min}}$) . This timescale was chosen because all PINES objects are observed for at least 5 minutes in a single block, so the standard deviation over 5 minutes intervals can be evaluated for every block for every target. The expected noise performance from Tamburo & Muirhead (2019) is indicated with a dashed line, including contributions from photon noise, sky background, dark current, read noise, and scintillation. The measured noise performance is worse than these expectations, and we attribute this mainly to poorer site seeing and lower net throughput than was anticipated. Our simulations assumed an average-seeing-FWHM value of 1″.5, based on an optical site survey of the PTO location (Tsay et al. 1990). Instead, we have measured an average-seeing FWHM of $2\farcs6^{+0\farcs5}_{-0\farcs4}$, necessitating larger photometric apertures, which result in higher contributions from sky background, read noise, and dark current. Additionally, we measure the net throughput in *J* band to be 4.1%, compared to 35% in Clemens et al. (2007). The discrepancy could be due to degraded optical coatings in the instrument and/or degraded detector quantum efficiency. We plan to address the throughput in a future instrument upgrade. If our measured throughput can be increased to the value measured by Clemens et al. (2007), we could expect an improvement in our average-light-curve scatter by a factor of $\sim\sqrt{35/4.1}$. Measured performance characteristics of the survey are given in Table 2 in comparison to the expectations from Tamburo & Muirhead (2019).

The solid line in Figure 9 shows an updated calculation of the PINES noise performance using the measured seeing,

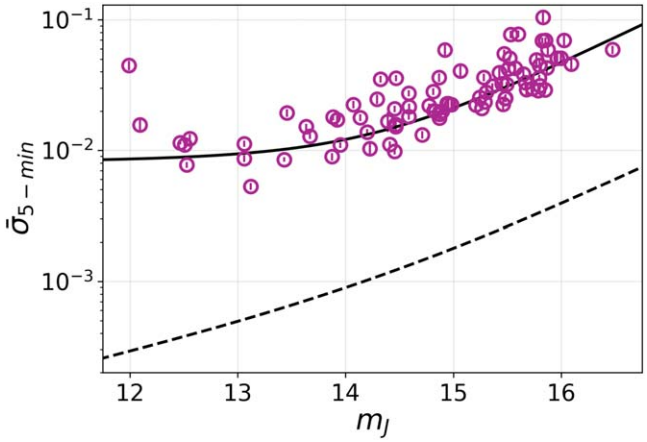

Figure 9. The average 5 minutes standard deviation ($\bar{\sigma}_{5\,\mathrm{min}}$) of PINES targets observed in *J* band (points with error bars). The dashed line indicates the expected noise performance from Tamburo & Muirhead (2019), while the solid line shows an updated calculation using measured seeing, background, and throughput values, along with a logistic noise floor for targets brighter than $m_J \approx 14$.

Table 2
Performance Characteristics of the PINES Survey, as of 2021 October, Compared to Expected Values from Tamburo & Muirhead (2019)

Parameter	Measured	Expected
Seeing	$2\farcs6^{+0\farcs5}_{-0\farcs4}$	1″.5
Nights requested	207	220
Nights assigned	181	220
Total night loss rate[a]	46%	30%
Weather loss rate[b]	36%	30%
J-band bkgd. (ADU/pix/s)	$17.9^{+8.0}_{-5.7}$	22
J-band throughput	4.1%	35.0%

Notes.
[a] This value accounts for whole-night losses from weather, forest closures, instrument failures, etc.
[b] This value accounts for whole-night losses from weather only.

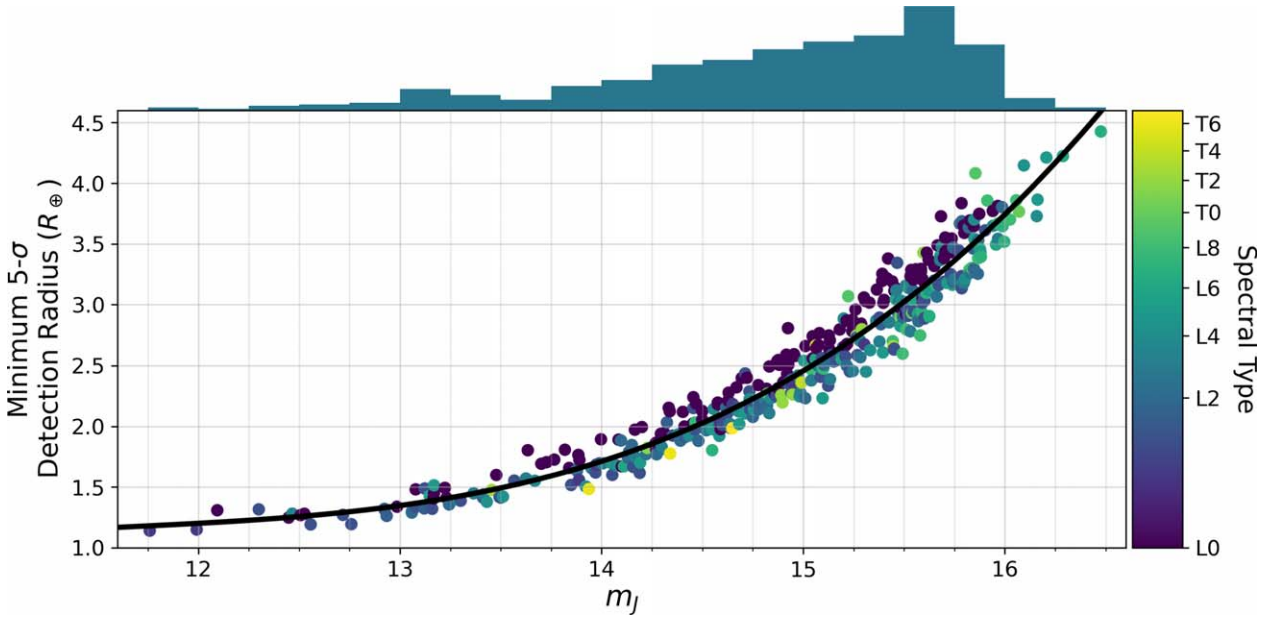

Figure 10. Estimated *J*-band transit-detection sensitivity for all of the targets in the PINES sample (colored points), assuming purely Gaussian noise as determined by the *J*-band noise model shown in Figure 9. Astrophysical or systematic noise sources, which are not captured in this model, will degrade this sensitivity estimation. Targets in this panel are colored by their spectral type, and the black line shows a third-degree polynomial fit to the points. The histogram above the axis shows the *J*-band magnitudes of targets in the PINES sample (the binning matches that in the top panel of Figure 2).

background, read noise, and throughput values given in Table 2. Additionally, the noise model incorporates a logistic function to capture the noise floor that we find for the targets brighter than $m_J \approx 14$. On average, there are fewer suitable reference stars available in the fields of bright targets. This tends to increase the noise of the ALC in these cases, which in turn inflates the scatter of the light curves of our brightest targets compared to expectations.

We used the *J*-band noise model in Figure 9 to estimate our transit-detection sensitivity. We simulated cycling observations for every target in our sample assuming an exposure time of 60 s, a block length of 12 minutes, and a cycle length of 1 hr (typical of the PINES cycling strategy). We converted the known spectral types of our targets to $T_{\rm eff}$ values using the $T_{\rm eff}$–SpT relation from Faherty et al. (2016) for field M6-T9 dwarfs, then selected a random age for each from a uniform distribution from 0.25 to 10 Gyr. These temperatures and ages were used to obtain realistic radius estimates for each target (given the random age assumptions) using the evolutionary models from Baraffe et al. (2015). We then calculated the minimum planet radius that could decrement one block of data to a significance of 5σ, assuming the *J*-band noise model shown in Figure 9. We emphasize that this noise model does not include sources of astrophysical or systematic noise, which would strictly increase the size of the smallest detectable planet around a given PINES target. The results of this estimation, then, should be interpreted as lower limits to our planet detection sensitivity, and a full injection/recovery simulation will be required to establish our true sensitivity on a target-by-target basis.

We show the results of this analysis in Figure 10. The targets are colored by their spectral type, and a third-degree polynomial fit to the results is shown as a black line. A histogram of target *J*-band magnitudes is shown on the top axis. A clear dependence on the spectral type can be seen in this plot. Early L dwarfs are found in larger quantity above the polynomial fit than below, while the opposite is true for later-type sources. This is because substellar objects shrink with time; if two objects have the same age, the earlier spectral type will have a larger radius than the later spectral type, requiring larger-radius planets to meet the 5σ detection criterion.

Under the assumption of purely Gaussian noise, this plot shows that PINES is sensitive to the 5σ detection of planets smaller than Neptune ($R = 3.86\ R_\oplus$) for all but the faintest targets in the sample. In particular, for the median *J*-band magnitude of the PINES sample ($m_J = 15.1$), we find a lower radius sensitivity limit of $2.5 \pm 0.2\ R_\oplus$. For our brightest targets, we calculate a lower sensitivity limit of $1.2\ R_\oplus$. We note that this lies within the 1σ uncertainty range of the "Terran-Neptunian" boundary at $R = 1.23^{+0.44}_{-0.22}\ R_\oplus$ identified by Chen & Kipping (2017), which marks the division between rocky and volatile-dominated planets.

4.2. Recovery of a Transit of WASP-2 b

In Figure 11, we show an example of the light curve creation process for WASP-2, a binary star whose K1V primary is known to host a transiting $1.1\ R_{\rm Jup}$ planet on a 2.15 days orbit (Collier Cameron et al. 2007; Daemgen et al. 2009). We selected this target for PINES observations to test our ability to recover known transits with the analysis pipeline. WASP-2 is not an L or T dwarf, and its *J*-band magnitude of 10.2 is about five magnitudes brighter than the median *J*-band magnitude of the PINES sample. However, WASP-2 b's transit depth of 1.7% is comparable to that of a super-Earth around a typical L or T dwarf (a $1.4\ R_\oplus$ planet would exhibit this transit depth around a $1.0\ R_{\rm Jup}$ L or T dwarf), and for that reason, it serves as an appropriate test case for transit recovery.

We observed WASP-2 on the night of UT 2021 July 28 in *J* band with a 5 s exposure time. We switched between WASP-2 and one other target throughout the night, observing each in 15 minutes blocks before switching to the other. Observing conditions were poor with intermittent clouds throughout the night, in some cases for extended periods, which interrupted

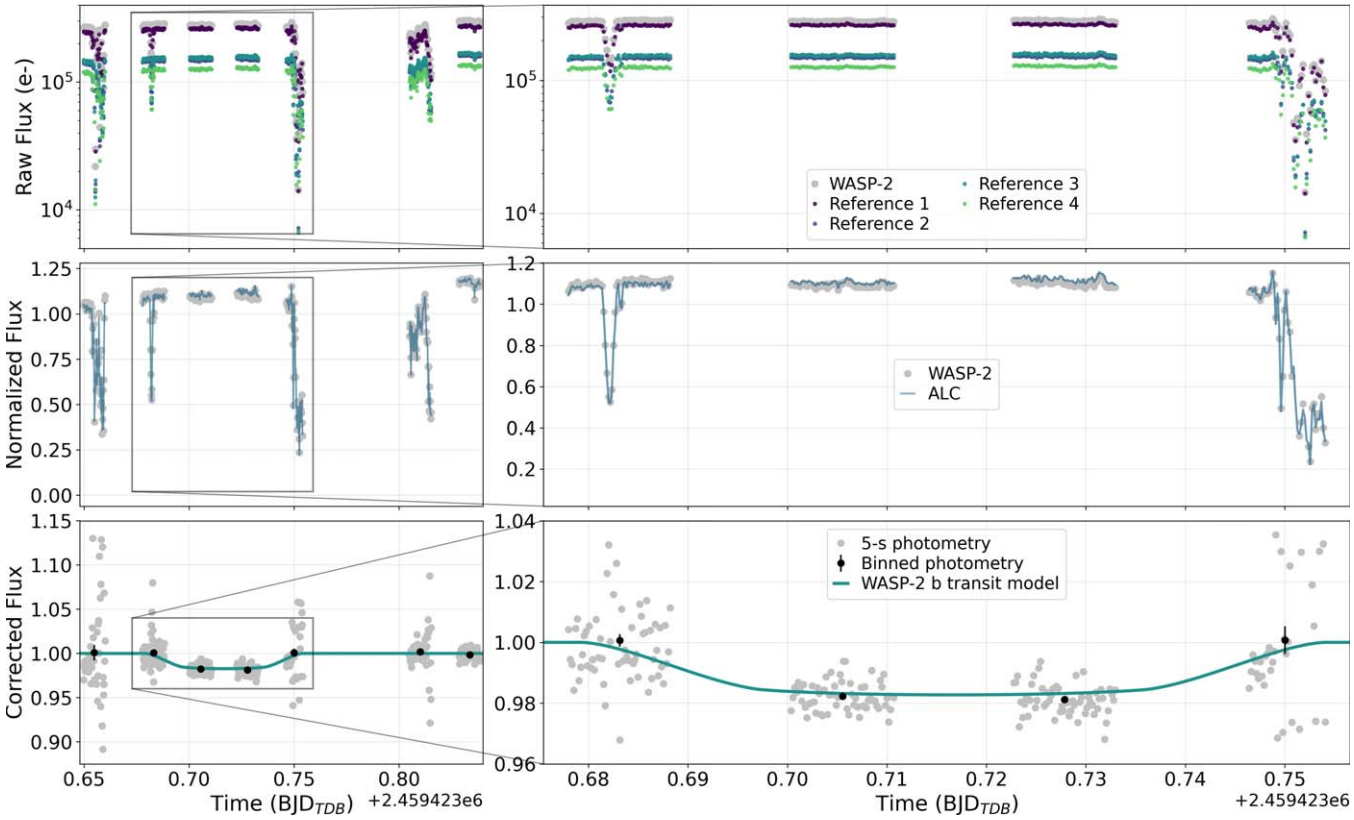

Figure 11. Transit recovery of WASP-2 b on UT 2021 July 28. The rows show different stages of PAT, going from raw flux to a final light curve. The left column shows the full night of data, while the right column shows a zoom-in on the four blocks that cover the transit event. Top: raw flux measured for WASP-2 and four reference stars. Middle: the normalized flux of WASP-2 and the ALC constructed from the weighted mean of the four reference stars. Bottom: the normalized flux for WASP-2 corrected by the ALC and a linear baseline, along with a transit model of WASP-2 b.

the desired cycle time of 30 minutes. Despite these interruptions, we managed to obtain seven blocks of data for WASP-2.

The top row of Figure 11 shows the raw flux measured for WASP-2 and four reference stars, using variable apertures with radii set to 1.5 × a smoothed trend of the measured seeing. The short-timescale variability in the raw flux is due to time-variable cloud cover. The middle row shows the normalized target flux ($\hat{F}_T$) and the ALC, created using the weighted-mean procedure described in Section 3.7. The bottom row shows the target flux corrected by the ALC and a linear baseline, with a transit model of WASP-2 b overplotted. The corrected flux shows excellent agreement with the transit model, which validates the observing strategy and pipeline for the purpose of detecting ∼1% transit depths.

4.3. Recovered Variables

In this section, we describe the recovery of variability signatures from PINES targets that are known to be rapid rotators. Because of the discontinuous time coverage of PINES data, our observations are generally inadequate for the accurate characterization of the rotational variability of L and T dwarfs; rather, we show that we can readily *detect* variability, and can use our observations to identify potential new variables, which can be confirmed with follow-up observations (see Section 4.4).

4.3.1. SIMP J16291840+0335371

SIMP J16291840+0335371 (SIMP 1629+0335) is a T2-type BD (Deacon et al. 2011) that had a measured rotation period of 6.9 ± 2.4 hr and a peak-to-peak variability amplitude of $4.3\% \pm 2.4\%$ in Radigan et al. (2014). SIMP 1629+0335 was observed by PINES for three nights in 2020 June in J band, and its light curve is shown in the top row of Figure 12. Images were taken with a 60 s exposure time, and the target was observed for 8 minutes approximately every hour.

We created a Lomb-Scargle (LS) periodogram of the SIMP 1629+0335 light curve using the implementation in the astropy Python package with frequencies from 0.01 to 1 hr^{-1} (equivalent to periods ranging from 1 to 100 hr). This periodogram, which is shown in the second row of Figure 12, shows a clear peak at 0.209 hr^{-1} (4.8 hr). We calculated the false alarm probability (FAP) of this peak using the approximation described in Baluev (2008) and as implemented in astropy, and found it to be 1.5e-7.

The period estimate from the LS periodogram is consistent with the rotation period measured in Radigan et al. (2014) to within 1σ. To quantify the uncertainty on our measured period, we performed a Markov Chain Monte Carlo (MCMC) simulation using emcee to fit our data with a sine curve of the form

$$1 + \frac{A_i}{2} \sin\left(2\pi\left(\frac{t}{P} + \phi\right)\right), \quad (5)$$

where A_i is the peak-to-peak amplitude on night i, P is the period, and ϕ is the phase. We allow the fitted amplitude to vary from night-to-night because the variability of L/T transition objects can evolve significantly on short timescales (e.g., Artigau et al. 2009; Gillon et al. 2013b). The allowed

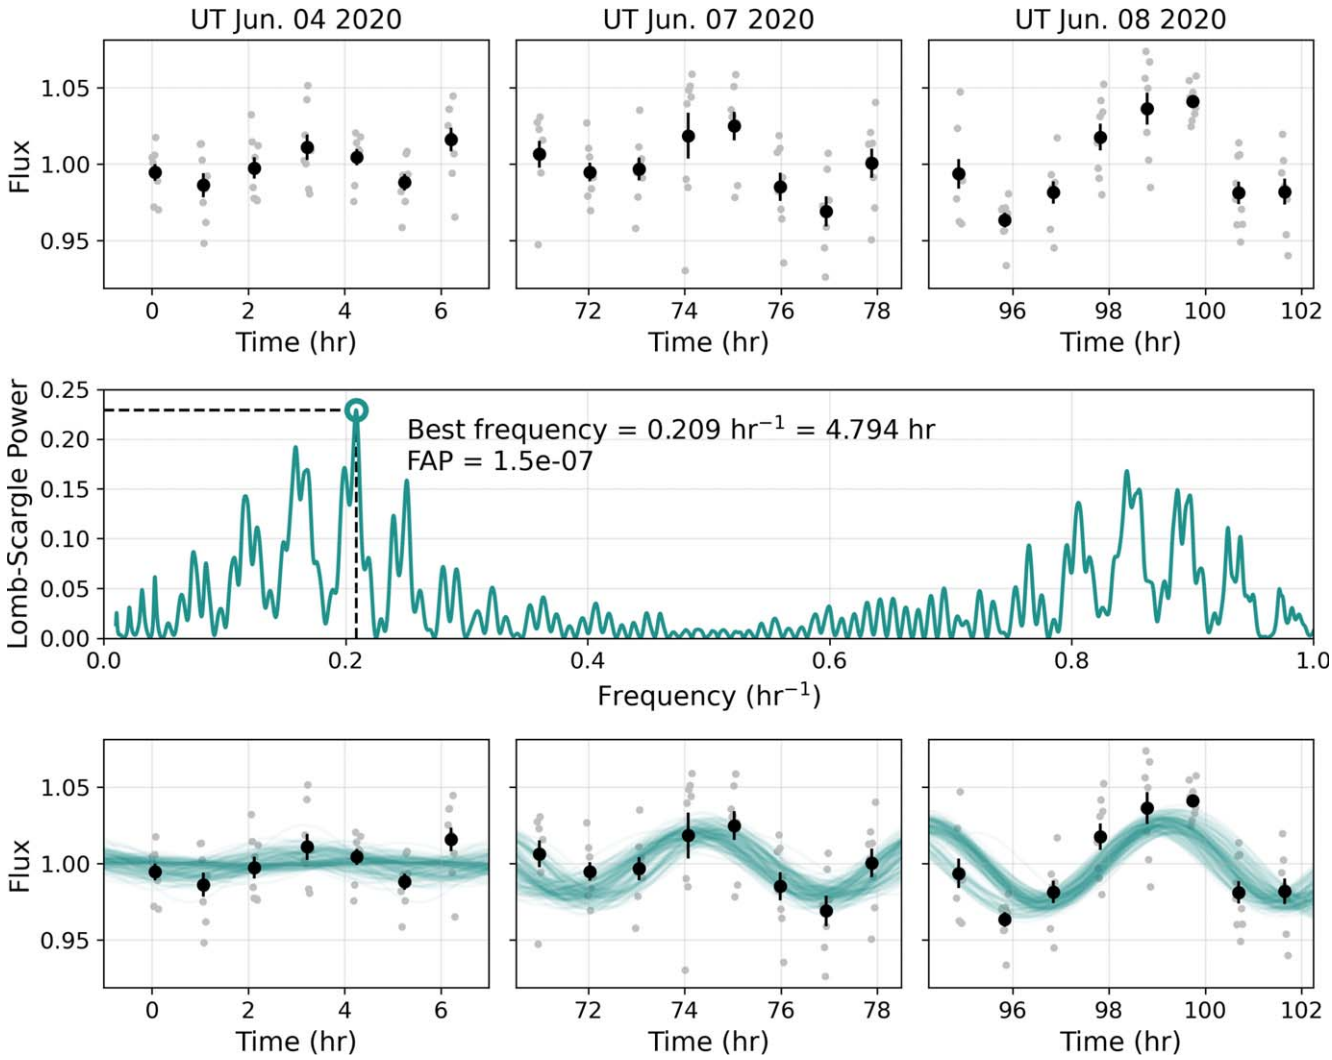

Figure 12. Variability recovery for SIMP 1629+0335. Top: PINES observations taken in *J* band with a 60 s exposure time. The 60 s exposures are shown in gray, while the data binned over 8 minutes blocks are shown in black with error bars. Middle: a Lomb-Scargle periodogram of the unbinned data. Bottom: photometry with random MCMC models overplotted.

periods were restricted to within ±3 hr of the peak of the LS periodogram. We used 128 walkers for a total of 10,000 steps, discarding the first 1000 steps of each chain for burn-in. Random samples from this MCMC simulation are shown in the third row of Figure 12. We find a period posterior that is bimodal with peaks at $4.97^{+0.08}_{-0.14}$ and $6.22^{+0.13}_{-0.19}$ hr, both of which are consistent with Radigan et al. (2014). The peak-to-peak amplitude posteriors reveal a slight increase in the variability from night to night: $1.24^{+1.08}_{-0.84}$% on night 1, $4.18^{+1.02}_{-1.01}$% on night 2, and $5.46^{+0.94}_{-0.89}$% on night 3.

The Spitzer archive possesses time-series observations of SIMP 1629+0335, which we downloaded and analyzed to constrain the true period of this target. We analyzed the data from two programs available in the archive—PID: 80213 (PI: J Radigan) obtained in 2012 and PID: 11132 (PI: J Radigan) obtained in 2015. The 2012 data set consists of ∼24 hr of continuous monitoring at 3.6 μm, and the 2015 data set consists of ∼7 hr of 4.5 μm monitoring followed by ∼7 hr of 3.6 μm monitoring.

We reduced and analyzed the three light curves following the methods presented in Vos et al. (2020) and briefly explain the steps here. We measured aperture photometry from the Basic Calibrated Data images produced by the Spitzer Science Center. We used a variety of apertures ranging from 2.0 to 5.0 pixels in steps of 0.25, and ultimately choose the aperture resulting in the highest signal-to-noise ratio final light curve. We corrected for the pixel-phase effect using a cubic function of the *x* and *y* coordinates, as described by previous light curve studies (e.g., Knutson et al. 2008; Metchev et al. 2015; Vos et al. 2020), and find that this correction significantly reduces correlations between the flux and pixel position for our target. We use two methods to check our light curves for significant variability—a periodogram analysis method and a Bayesian method making use of the Bayesian Information Criterion. Both methods find that SIMP 1629+0335 appears variable at all three epochs. To measure the rotation period of SIMP 1629 +0335, we used `emcee` to fit a sinusoidal curve to each Spitzer data set. We use 1000 walkers for a total of 1000 steps and discard 100 burn-in samples from each chain. In Figure 13, we

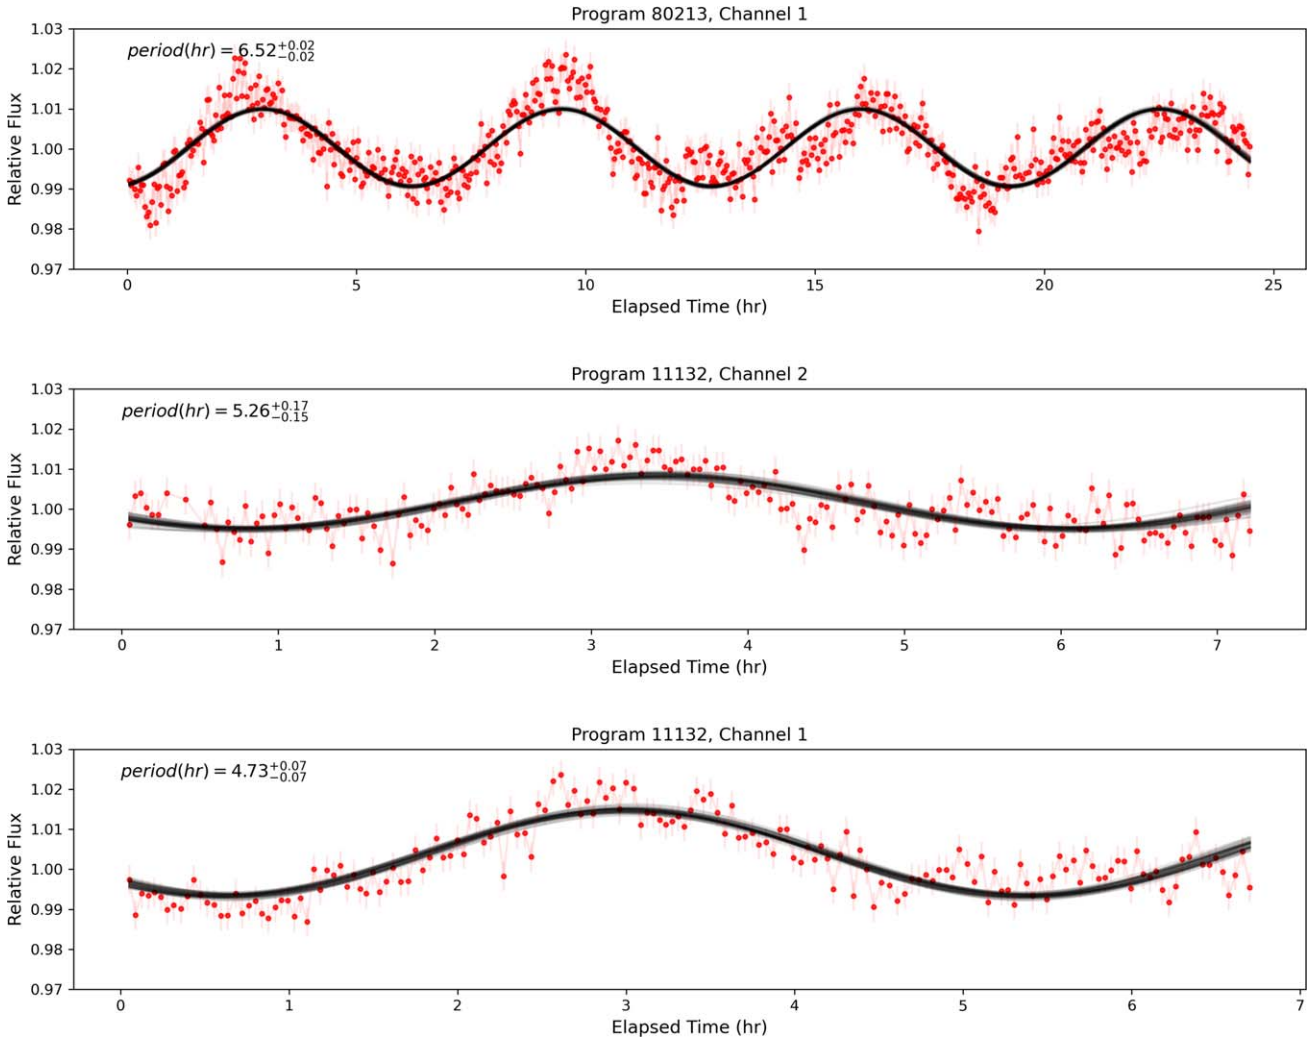

Figure 13. Spitzer light curves for the variable brown dwarf SIMP 1629+0335. The top panel displays the data taken in Channel 1 (3.6 μm) in 2012. The middle and bottom panels display the data from Channel 2 (4.5 μm) and Channel 1 (3.6 μm), respectively, taken in 2015. The corrected light curves are shown in red, with 50 random draws from our MCMC fit shown in black. The measured rotation period is significantly different in all three epochs.

show the three reduced light curves and their sinusoidal fits. We find that the measured rotation period changes significantly between the three different light curves: 6.52 ± 0.02 hr in Program 80213, $5.26^{+0.17}_{-0.15}$ hr in the first set of Program 11132 observations, and 4.73 ± 0.07 hr in the second set of Program 11132 observations. Programs 80213 and 11132 occurred three years apart, and a change in period over that time is not necessarily surprising. On the other hand, the two observations from Program 11132 are separated by the time span of hours, representing two consecutive rotations of the BD. The change in the measured period is likely due to rapidly moving and/or evolving structures in the atmosphere of the dwarf, as have previously been noted in the light curves of variable, L/T transition BDs (Apai et al. 2017). The long-duration, continuous observations of SIMP 1629+0335 using Spitzer suggest that the bi-modal period distribution measured from the PINES light curve is astrophysical in nature, and not a consequence of our observing strategy.

Finally, it should be noted that inconsistent variability renders planet searching prohibitively difficult, as it is not possible to remove a signal that wholly accounts for the measured variability. For this reason, targets like SIMP 1629 +0335 (and 2MASS 2139+0220, in the following section) are not included in searches for transiting planet signals in our pipeline.

4.3.2. 2MASS J21392676+0220226

2MASS J21392676+0220226 (2MASS 2139+0220) was found to be significantly variable in Radigan et al. (2012), who measured a period of 7.721 ± 0.005 hr and peak-to-peak amplitudes up to 26% in J band. We observed 2MASS 2139 +0220 for three nights in J band in 2020 August and performed the same variability analysis as in the previous section. The results of this analysis are shown in Figure 14.

In the LS periodogram for the 2MASS 2139+0220, the highest peak is at 6 hr, with an FAP of 6.7e-42. An MCMC fit with the function given in Equation (5) reveals a multimodal period distribution. The closest peak to the true period, at $7.91^{+0.05}_{-0.04}$ hr, is still significantly discrepant with the period measured in Radigan et al. (2012), potentially because the variability of 2MASS 2139+0220 is known to not be perfectly sinusoidal (as assumed in our fitted model). Despite the period ambiguity, however, variability is clearly visible in the PINES light curve of 2MASS 2139+0220. Were this an unknown variable, our pipeline would have flagged it for follow-up

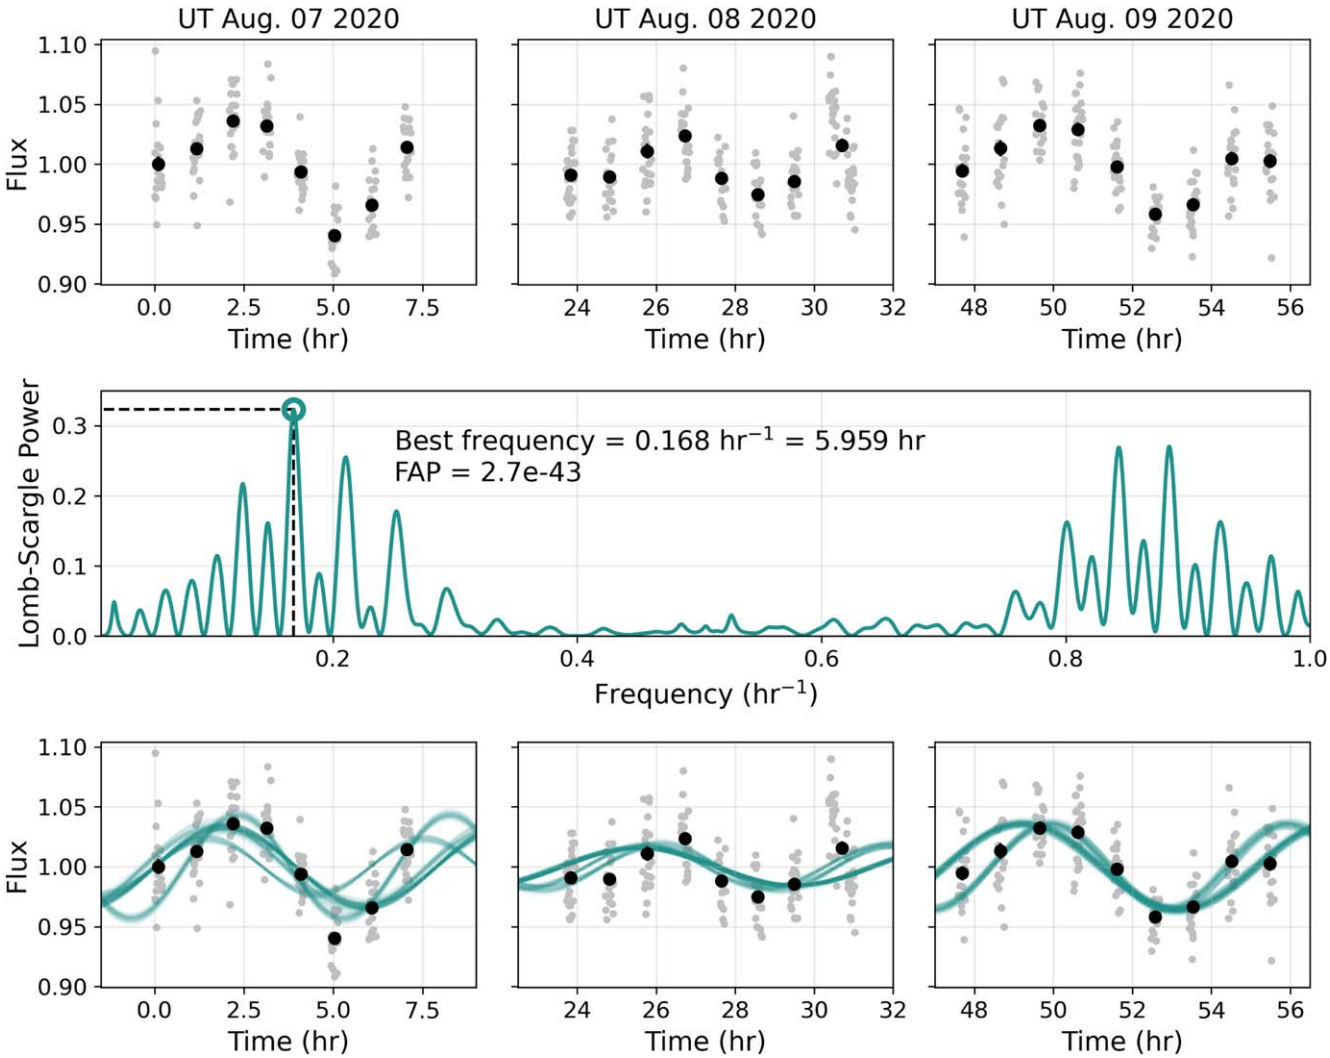

Figure 14. PINES light curve of the known variable 2MASS 2139+0220. The panels match those in Figure 12.

observations, which could be performed at a more suitable cadence for characterizing the variability.

4.4. New Variable: WISE J045746.08-020719.2

The previous two sections demonstrated our capacity for detecting variable objects in PINES data. Having performed these validation checks, we searched for new variables in our collection of PINES light curves and identified one previously unknown candidate variable: WISE J045746.08−020719.2 (WISE J0457−0207). Bihain et al. (2013) identified WISE J0457−0207 as a T2 dwarf using low-resolution NIR spectroscopy as part of a search for BDs in the solar neighborhood. Best et al. (2020) measured a parallax of 82.0 ± 2.9 mas for this object, with proper-motion components $\mu_\alpha \cos\delta = 93.0 \pm 3.0$ mas yr^{-1}, $\mu_\delta = -105.2 \pm 2.2$ mas yr^{-1}. We used these values to determine the probability that WISE J0457−0207 is a member of nearby young associations with BANYAN Σ (Gagné et al. 2018), finding a 61.5% probability for membership in the β Pictoris moving group. We have no reason to suspect that this object is young a priori, and an RV is required to determine WISE J0457−0207's potential β Pictoris membership. A literature search for the target revealed that no light curve observations have been published for WISE J0457−0207.

We present the PINES light curve of WISE J0457−0207 in Figure 15, which shows variability signatures consistent with those of the known variables recovered in Section 4.3. The target was observed for four nights in J band in 2020 December with a 60 s exposure time. We flagged the target as a potential variable after visual inspection, and then performed the variability retrieval analysis described in Section 4.3. The LS periodogram peak at 5.7 hr has a FAP of 1.3e-10. In contrast to the results for SIMP 1629+0335 and 2MASS 2139+0220, there is a single dominant peak in the period posterior for WISE J0457−0207, with a value of 5.748 ± 0.011 hr. The measured peak-to-peak variability amplitudes were $6.1\% \pm 0.8\%$ on night 1, $4.0\% \pm 0.6\%$ on night 2, $2.0\% \pm 0.6\%$ on night 3, and $6.5\% \pm 0.5\%$ on night 4, showing evolving variability that is typical of an L/T transition object (Radigan et al. 2014; Metchev et al. 2015). Some blocks of data are significantly discrepant with the pure sine model used in this analysis, however, and follow-up observations will be required to more fully characterize this object's variability.

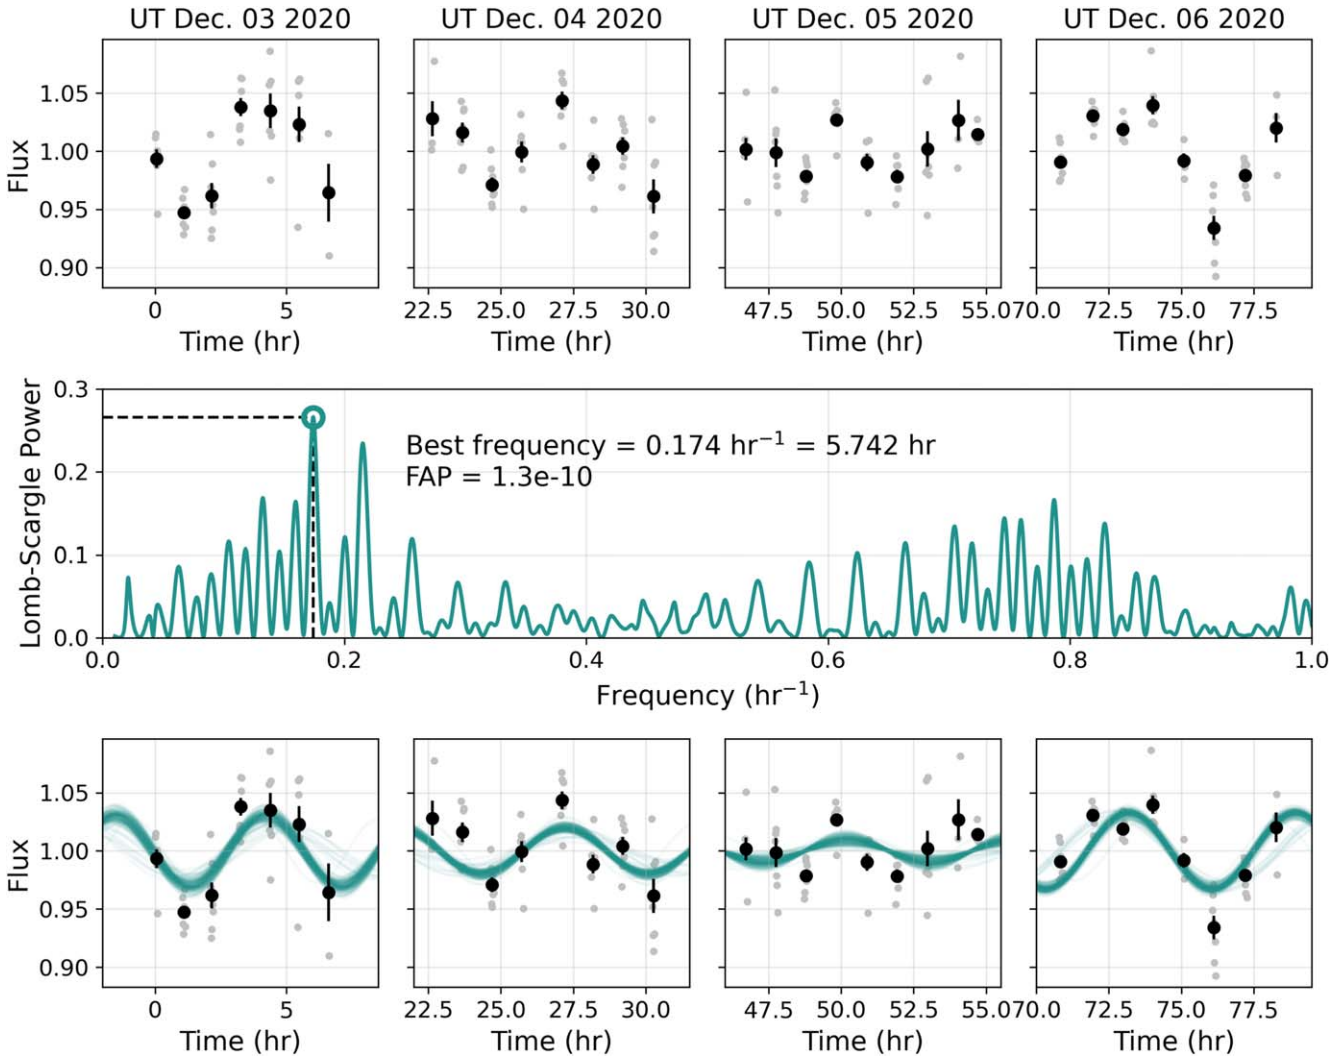

Figure 15. PINES light curve of the new variable WISE J0457−0207. We estimate its variability period to be 5.748 hr. The panels match those in Figure 12.

5. Conclusions

In this work, we provided an overview of PINES, a search for transiting satellites around a sample of almost 400 spectroscopically confirmed L and T dwarfs. Our main results can be summarized as follows:

1. We described the observing facility, construction of the target sample, and observational strategy of PINES, which was designed to maximize the likelihood of detecting short-duration transits around L and T hosts.
2. We detailed the custom guiding procedure that was created to execute the cycling strategy employed by PINES, and compared the performance of this system to the built-in PTO auto-guider. We found that, despite worse pointing performance, light curves created with the PINES guiding system have only marginally higher noise than those created with the auto-guider.
3. We described the major steps of `PAT`, the custom photometric pipeline that we created to analyze PINES data. We showed how we identify bad pixels, reduce data, determine centroids, perform photometry, and create final target light curves using a weighted mean of reference star fluxes.
4. We performed an analysis of the second-order extinction effect, in which changing PWV levels lead to different responses in the flux of different spectral types. We demonstrated that this effect is minimized in MKO *J* band, the filter used by PINES.
5. We applied `PAT` to observations of 83 PINES targets. We used the resulting light curves to measure the noise performance of the survey and found that it was significantly higher than the expectations from Tamburo & Muirhead (2019), due to a combination of worse average site seeing and net throughput. However, the derived *J*-band noise model suggests that PINES will still be sensitive to the detection of sub-Neptune-sized planets around essentially the entirety of the sample.
6. Finally, we detailed test cases that validate the performance of `PAT` for transit recovery and identifying variables. We identified one new variable in our observations thus far, the T2 dwarf WISE J045746.08−020719.2.

The PINES survey is ongoing, with roughly 70% of the sample remaining to be observed. PINES observations will permit stronger constraints on the occurrence rates of short-period planets around L and T dwarfs, which generally lie

beyond the detection limits of previous optical exoplanet surveys. In turn, this effort will help to resolve the tension between competing lines of evidence for the existence of such planets around very-low-mass stars, BDs, and planetary-mass objects.

The authors thank their anonymous referee for a thorough and insightful review that improved the quality of this work.

This material is based upon work supported by the National Aeronautics and Space Administration under grant No. 80NSSC20K0256 issued through the Science Mission Directorate.

This work is based on observations made with the Spitzer Space Telescope, which is operated by the Jet Propulsion Laboratory, California Institute of Technology under a contract with NASA. The Spitzer data is available at the Spitzer Heritage Archive hosted by IRSA (Spitzer Science Center 2020).

Support for this work was provided by NASA through the NASA Hubble Fellowship grant HST-HF2-51447.001-A awarded by the Space Telescope Science Institute, which is operated by the Association of Universities for Research in Astronomy, Inc., for NASA, under contract NAS5-26555.

Facilities: Boston University Perkins Telescope Observatory (Mimir), Spitzer Space Telescope (IRAC).

Software: `Astrometry.net` (Lang et al. 2010), `Astropy` (Astropy Collaboration et al. 2013, 2018), `Astroquery` (Ginsburg et al. 2019), `BANYAN Σ` (Gagné et al. 2018), `BATMAN` (Kreidberg 2015), `emcee` (Foreman-Mackey et al. 2013), `Fyodor` (Meier Valdés et al. 2021), `Matplotlib` (Hunter 2007), `NumPy` (Harris et al. 2020), `photutils` (Bradley et al. 2020), `PINES Analysis Toolkit` (Tamburo 2022).

Appendix
Calculating Reference Star Weights

The reference star weights used to create the ALC in Equation (3), W_i, are calculated to give more weight to bright, stable reference stars, while giving less weight to dim or variable reference stars. The weights are determined through an iterative process, based on a procedure described in Murray et al. (2020). We construct a *special* ALC (SALC) for every reference star in the field using the flux of all of the *other* reference stars, weighted by their calculated uncertainties:

$$\mathrm{SALC}_i = \frac{\sum_{j=1,j\neq i}^{N_{\mathrm{refs}}} \hat{F}_{\mathrm{R}_j} W_{\mathrm{R}_j}}{\sum_{j=1,j\neq i}^{N_{\mathrm{refs}}} W_{\mathrm{R}_j}}. \tag{A1}$$

The SALC is then used to correct the flux of the reference star in question:

$$\hat{F}^{\star}_{\mathrm{R}_i} = \frac{\hat{F}_{\mathrm{R}_i}}{\mathrm{SALC}_i}. \tag{A2}$$

The standard deviation of $\hat{F}^{\star}_{\mathrm{R}_i}$, $\sigma^{\star}_{\mathrm{R}_i}$ is measured; then, W_i is given by

$$W_i = 1/\sigma^{\star 2}_{\mathrm{R}_i}. \tag{A3}$$

This operation is performed for every reference star, and then the process is repeated, with the W_{R} values in Equation (A1) replaced by the calculated W_i values in subsequent loops. This continues until the weights have stabilized to values within 1e–9 of their values in the previous loop. Once the weight values have converged, they are normalized to sum to one.

We tested this weighting procedure using simulated photometry, and found it to perform as expected. When we injected the variability signatures into the reference stars, we found that those stars were down-weighted appropriately in the calculation of the target ALC, even if the variability was added to the brightest reference star.

ORCID iDs

Patrick Tamburo https://orcid.org/0000-0003-2171-5083
Philip S. Muirhead https://orcid.org/0000-0002-0638-8822
Allison M. McCarthy https://orcid.org/0000-0003-2015-5029
Murdock Hart https://orcid.org/0000-0002-0330-1648
David Gracia https://orcid.org/0000-0002-0447-7231
Johanna M. Vos https://orcid.org/0000-0003-0489-1528
Daniella C. Bardalez Gagliuffi https://orcid.org/0000-0001-8170-7072
Jacqueline Faherty https://orcid.org/0000-0001-6251-0573
Christopher Theissen https://orcid.org/0000-0002-9807-5435
Eric Agol https://orcid.org/0000-0002-0802-9145
Julie N. Skinner https://orcid.org/0000-0002-4794-1591
Sheila Sagear https://orcid.org/0000-0002-3022-6858

References

Allard, F. 2014, in Exploring the Formation and Evolution of Planetary, ed. M. Booth, B. C. Matthews, & J. R. Graham, Vol. 299 (Cambridge: Cambridge Univ. Press), 271
Apai, D., Karalidi, T., Marley, M. S., et al. 2017, Sci, 357, 683
Artigau, É., Bouchard, S., Doyon, R., & Lafrenière, D. 2009, ApJ, 701, 1534
Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, A&A, 558, A33
Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, AJ, 156, 123
Bailer-Jones, C. A. L., & Lamm, M. 2003, MNRAS, 339, 477
Baker, A. D., Blake, C. H., & Sliski, D. H. 2017, PASP, 129, 085002
Ballard, S., & Johnson, J. A. 2016, ApJ, 816, 66
Baluev, R. V. 2008, MNRAS, 385, 1279
Baraffe, I., Homeier, D., Allard, F., & Chabrier, G. 2015, A&A, 577, A42
Béjar, V. J. S., Zapatero Osorio, M. R., Pérez-Garrido, A., et al. 2008, ApJL, 673, L185
Best, W. M. J., Liu, M. C., Magnier, E. A., & Dupuy, T. J. 2020, AJ, 159, 257
Bihain, G., Scholz, R. D., Storm, J., & Schnurr, O. 2013, A&A, 557, A43
Blake, C. H., Bloom, J. S., Latham, D. W., et al. 2008, PASP, 120, 860
Blake, C. H., Charbonneau, D., & White, R. J. 2010, ApJ, 723, 684
Borucki, W. J., Koch, D., Basri, G., et al. 2010, Sci, 327, 977
Bradley, L., Sipőcz, B., Robitaille, T., et al. 2020, astropy/photutils: v1.0.0, Zenodo, doi: 10.5281/zenodo.4044744
Broekhoven-Fiene, H., Matthews, B., Duchêne, G., et al. 2014, ApJ, 789, 155
Burgasser, A. J., Geballe, T. R., Leggett, S. K., Kirkpatrick, J. D., & Golimowski, D. A. 2006, ApJ, 637, 1067
Burrows, A., Marley, M., Hubbard, W. B., et al. 1997, ApJ, 491, 856
Cannon, A. J., & Pickering, E. C. 1901, AnHar, 28, 129
Chabrier, G. 2005, in The Initial Mass Function: From Salpeter 1955 to 2005, ed. E. Corbelli, F. Palla, & H. Zinnecker, Vol. 327 (Berlin: Springer), 41
Chauvin, G., Lagrange, A. M., Dumas, C., et al. 2004, A&A, 425, L29
Chauvin, G., Lagrange, A. M., Dumas, C., et al. 2005, A&A, 438, L25
Chen, J., & Kipping, D. 2017, ApJ, 834, 17
Clemens, D. P., Sarcia, D., Grabau, A., et al. 2007, PASP, 119, 1385
Close, L. M., Zuckerman, B., Song, I., et al. 2007, ApJ, 660, 1492
Collier Cameron, A., Bouchy, F., Hébrard, G., et al. 2007, MNRAS, 375, 951
Crump, C. C. 1929, PA, 37, 553
Daemgen, S., Hormuth, F., Brandner, W., et al. 2009, A&A, 498, 567
Dawson, P., Scholz, A., Ray, T. P., et al. 2013, MNRAS, 429, 903
Deacon, N. R., Liu, M. C., Magnier, E. A., et al. 2011, AJ, 142, 77
Delrez, L., Gillon, M., Queloz, D., et al. 2018, Proc. SPIE, 10700, 107001I
Dressing, C. D., & Charbonneau, D. 2013, ApJ, 767, 95
Dressing, C. D., & Charbonneau, D. 2015, ApJ, 807, 45
Dupuy, T. J., & Liu, M. C. 2012, ApJS, 201, 19

Epchtein, N., de Batz, B., Capoani, L., et al. 1997, Msngr, 87, 27
Faherty, J. K., Riedel, A. R., Cruz, K. L., et al. 2016, ApJS, 225, 10
Fontanive, C., Allers, K. N., Pantoja, B., et al. 2020, ApJL, 905, L14
Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, PASP, 125, 306
Gagné, J., Mamajek, E. E., Malo, L., et al. 2018, ApJ, 856, 23
Gaia Collaboration, Brown, A. G. A., Vallenari, A., et al. 2021, A&A, 649, A1
Gillon, M., Jehin, E., Fumel, A., Magain, P., & Queloz, D. 2013a, in EPJ Web of Conf. 47, Hot Planets and Cool Stars (Les Ulis: EDP Sciences), 03001
Gillon, M., Triaud, A. H. M. J., Jehin, E., et al. 2013b, A&A, 555, L5
Gillon, M., Jehin, E., Lederer, S. M., et al. 2016, Natur, 533, 221
Gillon, M., Triaud, A. H. M. J., Demory, B.-O., et al. 2017, Natur, 542, 456
Ginsburg, A., Sipőcz, B. M., Brasseur, C. E., et al. 2019, AJ, 157, 98
Han, C., Jung, Y. K., Udalski, A., et al. 2013, ApJ, 778, 38
Hardegree-Ullman, K. K., Cushing, M. C., Muirhead, P. S., & Christiansen, J. L. 2019, AJ, 158, 75
Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, Natur, 585, 357
Harvey, P. M., Henning, T., Ménard, F., et al. 2012, ApJL, 744, L1
Hayashi, C., & Nakano, T. 1963, PThPh, 30, 460
He, M. Y., Triaud, A. H. M. J., & Gillon, M. 2017, MNRAS, 464, 2687
Howard, A. W., Marcy, G. W., Bryson, S. T., et al. 2012, ApJS, 201, 15
Hunter, J. D. 2007, CSE, 9, 90
Ingalls, J. G., Krick, J. E., Carey, S. J., et al. 2012, Proc. SPIE, 8442, 84421Y
Johnson, J. A., Aller, K. M., Howard, A. W., & Crepp, J. R. 2010, PASP, 122, 905
Jung, Y. K., Udalski, A., Gould, A., et al. 2018a, AJ, 155, 219
Jung, Y. K., Hwang, K.-H., Ryu, Y.-H., et al. 2018b, AJ, 156, 208
Kirkpatrick, J. D., Reid, I. N., Liebert, J., et al. 1999, ApJ, 519, 802
Kirkpatrick, J. D., Gelino, C. R., Cushing, M. C., et al. 2012, ApJ, 753, 156
Kirkpatrick, J. D., Martin, E. C., Smart, R. L., et al. 2019, ApJS, 240, 19
Klein, R., Apai, D., Pascucci, I., Henning, T., & Waters, L. B. F. M. 2003, ApJL, 593, L57
Knutson, H. A., Charbonneau, D., Allen, L. E., Burrows, A., & Megeath, S. T. 2008, ApJ, 673, 526
Kreidberg, L. 2015, PASP, 127, 1161
Kroupa, P. 2001, MNRAS, 322, 231
Kumar, S. S. 1963, ApJ, 137, 1121
Lang, D., Hogg, D. W., Mierle, K., Blanton, M., & Roweis, S. 2010, AJ, 139, 1782
Lauer, T. R. 1999, PASP, 111, 1434
Lawrence, A., Warren, S. J., Almaini, O., et al. 2007, MNRAS, 379, 1599
Liu, Y., Joergens, V., Bayo, A., Nielbock, M., & Wang, H. 2015, A&A, 582, A22
Lord, S. D. 1992, NASA Technical Memorandum 103957
Luhman, K. L. 2007, ApJS, 173, 104
Luhman, K. L., Mamajek, E. E., Allen, P. R., Muench, A. A., & Finkbeiner, D. P. 2009, ApJ, 691, 1265
Luhman, K. L., Lada, C. J., Hartmann, L., et al. 2005, ApJL, 631, L69
Mace, G. N. 2014, PhD thesis, University of California, Los Angeles
Meier Valdés, E. A., Morris, B. M., & Demory, B. O. 2021, A&A, 649, A132
Metchev, S. A., Heinze, A., Apai, D., et al. 2015, ApJ, 799, 154
Miguel, Y., Cridland, A., Ormel, C. W., Fortney, J. J., & Ida, S. 2020, MNRAS, 491, 1998
Monin, J. L., Guieu, S., Pinte, C., et al. 2010, A&A, 515, A91
Muirhead, P. S., Mann, A. W., Vanderburg, A., et al. 2015, ApJ, 801, 18
Mulders, G. D., Drążkowska, J., van der Marel, N., Ciesla, F. J., & Pascucci, I. 2021, ApJL, 920, L1
Mulders, G. D., Pascucci, I., & Apai, D. 2015, ApJ, 798, 112
Murray, C. A., Delrez, L., Pedersen, P. P., et al. 2020, MNRAS, 495, 2446
Nakajima, T., Tsuji, T., & Yanagisawa, K. 2004, ApJ, 607, 499
Nutzman, P., & Charbonneau, D. 2008, PASP, 120, 317
Payne, M. J., & Lodato, G. 2007, MNRAS, 381, 1597
Pecaut, M. J., & Mamajek, E. E. 2013, ApJS, 208, 9
Radigan, J., Jayawardhana, R., Lafrenière, D., et al. 2012, ApJ, 750, 105
Radigan, J., Lafrenière, D., Jayawardhana, R., & Artigau, E. 2014, ApJ, 793, 75
Rilinger, A. M., & Espaillat, C. C. 2021, ApJ, 921, 182
Rilinger, A. M., Espaillat, C. C., & Macías, E. 2019, ApJ, 878, 103
Sagear, S. A., Skinner, J. N., & Muirhead, P. S. 2020, AJ, 160, 19
Sahlmann, J., Lazorenko, P. F., Ségransan, D., et al. 2013, A&A, 556, A133
Scholz, A., Jayawardhana, R., & Wood, K. 2006, ApJ, 645, 1498
Sebastian, D., Gillon, M., Ducrot, E., et al. 2021, A&A, 645, A100
Sestovic, M., & Demory, B.-O. 2020, A&A, 641, A170
Simons, D. A., & Tokunaga, A. 2002, PASP, 114, 169
Skrutskie, M. F., Cutri, R. M., Stiening, R., et al. 2006, AJ, 131, 1163
Spitzer Science Center 2020, Spitzer Level 1/Basic Calibrated Data Dataset, IPAC, doi: 10.26131/IRSA430
Tamburo, P. 2022, patricktamburo/pines_analysis_toolkit: PINES Analysis Toolkit Release v1.0.0, Zenodo, doi: 10.5281/zenodo.6200054
Tamburo, P., & Muirhead, P. S. 2019, PASP, 131, 114401
Tannock, M. E., Metchev, S., Heinze, A., et al. 2021, AJ, 161, 224
Tokunaga, A. T., Simons, D. A., & Vacca, W. D. 2002, PASP, 114, 180
Tsay, W.-S., Bagnuolo, W. G. J., McAlister, H. A., White, N. M., & Forbes, F. F. 1990, PASP, 102, 1339
van der Plas, G., Ménard, F., Ward-Duong, K., et al. 2016, ApJ, 819, 102
Van Grootel, V., Fernandes, C. S., Gillon, M., et al. 2018, ApJ, 853, 30
Vos, J. M., Biller, B. A., Bonavita, M., et al. 2019, MNRAS, 483, 480
Vos, J. M., Biller, B. A., Allers, K. N., et al. 2020, AJ, 160, 38
Vrba, F. J., Henden, A. A., Luginbuhl, C. B., et al. 2004, AJ, 127, 2948
Wilson, P. A., Rajan, A., & Patience, J. 2014, A&A, 566, A111
Wright, E. L., Eisenhardt, P. R. M., Mainzer, A. K., et al. 2010, AJ, 140, 1868