MNRAS **000**, 1–?? (2017) Preprint 2 October 2017 Compiled using MNRAS LATEX style file v3.0

A 3π Search for Planet Nine at 3.4μm with WISE and NEOWISE

A. M. Meisner,[1,2⋆] B. C. Bromley[3] ...
[1] *Berkeley Center for Cosmological Physics, Berkeley, CA 94720, USA*
[2] *Lawrence Berkeley National Laboratory, Berkeley, CA, 94720, USA*
[3] *Department of Physics & Astronomy, University of Utah, Salt Lake City, UT, 84112, USA*

2 October 2017

ABSTRACT
The recent 'Planet Nine' hypothesis has led to many observational and archival searches for this giant planet proposed to orbit the Sun at hundreds of astronomical units. While trans-Neptunian object searches are typically conducted in the optical, models suggest Planet Nine would likely be self-luminous and potentially bright enough at ∼3–5μm to be detected by the Wide-field Infrared Survey Explorer (WISE). We have previously demonstrated a Planet Nine search methodology based on time-resolved WISE coadds, allowing us to detect moving objects much fainter than would be possible using single-frame extractions. In the present work, we extend our 3.4μm (W1) search to cover more than three quarters of the sky and incorporate four years of WISE observations spanning a seven year time period. This represents the deepest and widest-area WISE search for Planet Nine to date. We rule out the presence of Planet Nine in the parameter space searched at W1 < 16.7 (90% completeness).

Key words: surveys: trans-Neptunian objects — methods: data analysis — techniques: image processing

1 INTRODUCTION

Over the past fifteen years, observations have emerged which challenge the conventional understanding of our solar system's outer regions. Chief among these are the discoveries of Sedna (Brown et al. 2004) and 2012 VP_{113} (Trujillo & Sheppard 2014), trans-Neptunian objects (TNOs) with orbits that cannot be explained by perturbations from the solar system's known planets. These two sednoids are often grouped together with other high perihelion ($q > 30$ AU), large semi-major axis ($a > 250$ AU) bodies into a class of objects referred to as extreme TNOs (ETNOs). Thirteen such ETNOs are now known (Brown et al. 2004; Trujillo & Sheppard 2014; Sheppard & Trujillo 2016; Bannister et al. 2017). These objects have been found to cluster in argument of perihelion (ω; Trujillo & Sheppard 2014) or longitude of perihelion (Batygin & Brown 2016a; Brown 2017). To explain the clustering in ω, Trujillo & Sheppard (2014) proposed the existence of a super-Earth mass perturber at ∼250 AU. In a related but different analysis, Batygin & Brown (2016a) suggested that an unseen giant planet ('Planet Nine') of mass 5–20$M_\oplus$ can simultaneously explain many mysteries of the outer solar system. These include the longitude of perihelion clustering among ETNOs, the existence of Sedna and 2012 VP_{113}, the population of high-inclination TNOs (Batygin & Brown 2016b), and the solar obliquity (Bailey et al. 2016; Gomes et al. 2017; Lai 2016). Criticisms of this Planet Nine narrative have also been put forth. For instance, the ETNO clustering in longitude of perihelion has been argued to result from observational bias (Shankman et al. 2017), and the presence of such a massive unseen planet has been claimed inconsistent with Cassini ranging measurements (Folkner et al. 2016).

Nevertheless, the Planet Nine hypothesis makes concrete predictions for the proposed planet's orbital parameters ($380 \lesssim a/\mathrm{AU} \lesssim 980$, $i \sim 30^\circ$; Brown & Batygin 2016), facilitating searches which can directly observe or else place limits on its presence. Planet Nine is expected to be faint in the optical, but not prohibitively so ($22 < V < 25$; Brown & Batygin 2016). Studies have also suggested the possibility of detecting Planet Nine at a range of other wavelengths (Cowan et al. 2016; Fortney et al. 2016; Linder & Mordasini 2016), depending on its detailed physical properties. Owing in part to these encouraging forecasts, many archival and/or observational searches for Planet Nine are already underway (Sheppard & Trujillo 2016; Weryk et al. 2016; Fortney et al. 2016; Kuchner et al. 2017; Meisner et al. 2017a).

In this work, we focus on searching for Planet Nine in

⋆ ameisner@lbl.gov

the infrared. Evolutionary models of Planet Nine predict that it would be self-luminous (Linder & Mordasini 2016; Fortney et al. 2016). In particular, Fortney et al. (2016) suggested that planet nine may be much brighter than blackbody expectations at 3–5μm, based on a detailed analysis of model atmospheres. Fortney et al. (2016) found that Planet Nine may be detectable by the Wide-field Infrared Survey Explorer (WISE) in its bluest channel (W1), although the models span an enormous range of luminosities at the relevant wavelengths.

It happens that the W1-brightest model from Fortney et al. (2016) remains fainter than the single-exposure detection limit, but would be readily detectable in coadds (at a fiducial distance of $d_9 \sim 600$ AU). In Meisner et al. (2017a), hereafter M17, we developed a methodology to convert the WISE W1 imaging into a serendipitous search for Planet Nine using a custom set of time-resolved coadd images. We applied this procedure to ∼2,000 square degrees of sky which, at the time, was thought to represent a likely Planet Nine location (Holman & Payne 2016). Here, we extend our W1 search to cover more than three quarters of the sky and include an additional year of WISE data. Although several predictions have been made for the most likely present-day Planet Nine sky location (Holman & Payne 2016; de la Fuente Marcos & de la Fuente Marcos 2016; Millholland & Laughlin 2017), our current strategy is to uniformly search the widest area possible. The very high and homogeneous quality of the space-based imaging delivered by WISE is conducive to deriving detailed constraints on Planet Nine's apparent mid-infrared brightness in the event of a non-detection.

In §2 we describe the archival WISE data used in the course of this search. In §3 we review the Planet Nine search methodology presented in M17. In §4 we provide details of the footprint searched during the course of this work. In §5 we explain our orbit linking procedure. In §6 we discuss the sensitivity of our search. In §7 we describe the results of our search. We conclude in §8.

2 DATA

We begin by very briefly reviewing key aspects of the WISE mission and its timeline of observations. WISE is a 0.4 meter telescope onboard a satellite in low-Earth orbit. The satellite was launched in late 2009. WISE has surveyed the entire sky in four broad bandpasses centered at 3.4μm (W1), 4.6μm (W2), 12μm (W3) and 22μm (W4). WISE successfully completed a seven month full-sky survey in all four of its channels beginning in 2010 January. During 2010 August and September, WISE continued surveying the sky in its bluest three bands. After that, it performed an asteroid hunting mission called NEOWISE until 2011 February (Mainzer et al. 2011), with only the W1 and W2 channels remaining operational. Following a nearly three-year hibernation period, the WISE instrument recommenced survey operations in 2013 December (Mainzer et al. 2014). This post-hibernation mission is referred to as NEOWISE-Reactivation (NEOWISER). WISE continues to survey the sky, and has already completed approximately eight full W1/W2 sky passes over the course of its first seven years in orbit.

NEOWISER publicly releases its data, consisting of single-exposure images and source extractions, in chunks that each span one year of observations. The data releases occur on a roughly annual basis. Thus far three such data releases have been published, in 2015 March, 2016 March and 2017 June.

A major difference between the present work and our initial search presented in M17 is that we have folded in an additional year of NEOWISER exposures obtained from the third-year NEOWISER data release, which became publicly available in 2017 June. In this work we employ all publicly available W1 Level 1b (L1b) single-exposure images that fall within our search footprint described in §4. The acquisition dates range from 7 January 2010 to 13 December 2016 (UTC). In other words, we use all pre-hibernation W1 L1b frames, plus those frames included in the first *three* annual NEOWISER data releases. This amounts to ∼49 months of WISE observations. While the input data spans a roughly seven year time period, the time baseline at a given sky location is typically 6.5 years.

We have dowloaded local copies of all publicly available W1 L1b sky images (`-int-` files) along with their corresponding uncertainty maps (`-unc-` files) and bitmasks (`-msk-` files). We refer to one such group of three W1 files resulting from a single WISE pointing as a 'frameset'. The coadds presented in this work make use of ∼6 million such framesets, for a total of ∼40 TB and ∼18 terapixels of single-exposure inputs.

3 OVERVIEW OF SEARCH METHODOLOGY

Our search strategy is modeled after Brown et al. (2015), although our analysis begins at the pixel level rather than the catalog level. Brown et al. (2015) developed a method for linking transients with arbitrary time sampling into Keplerian orbits for solar system objects beyond 25 AU. We employ a similar orbit linking approach which we describe fully in §5. The Brown et al. (2015) analysis starts with large catalogs of transient candidates provided by the Catalina Real-time Transient Survey (CRTS; Drake et al. 2009). On the other hand, the starting point for our analysis is a set of ∼6 million single-exposure WISE/NEOWISE images. Therefore, much of our analysis is devoted to image processing steps which allow us to obtain a clean list of transients from the available WISE images. In brief, our method stacks W1 exposures into a set of epochal coadds, extracts transients via difference imaging, and finally attempts to link these transients into Keplerian orbits beyond 250 AU. The following subsections briefly review each step of our Planet Nine search procedure. Full details are provided in M17.

3.1 Tiling

Our search region (§4) is composed of a series of 'tiles', with each tile corresponding to a unique ∼2.5 square degree astrometric footprint. These are the footprints onto which we reproject L1b exposures during coaddition. We adopt tile footprints matching those of the WISE team's Atlas stacks (Cutri et al. 2011). The Atlas tiling covers the entire sky with a set of 18,240 astrometric footprints. Atlas tile centers trace out a series of isodeclination rings. Each Atlas tile footprint

is 1.56° on a side, and is identified by a `coadd_id` value. The `coadd_id` is a string encoding the tile's central (RA, Dec) coordinates. For instance, the Atlas tile centered at (RA, Dec) = (324.8°, −19.7°) has `coadd_id` = '3248m197'. For our time-resolved coadds, we adopt a pixel scale matching that of the L1b images, 2.75″/pixel. Thus, each of our tiles has dimensions of 2048 pixels × 2048 pixels. In our terminology, 'tiles' are distinct from 'coadds': there are exactly 18,240 unique tile footprints, whereas we have typically constructed 8–9 epochal coadds for each such tile footprint within our search region. Note that although the 18,240 astrometric footprints of the Atlas tiles are unique, neighboring tiles overlap at their boundaries. The amount of overlap is ∼3′ near the celestial equator.

3.2 Time Slicing

We seek to extract transients from time-resolved coadds of the W1 imaging rather than from single exposures themselves. This choice has two main advantages: the coadds are ∼1.3 magnitudes deeper than single frames, and are also much cleaner because fast transients such as asteroids, cosmic rays, and satellite streaks are eliminated.

The WISE survey strategy is such that a given sky region is 'visited' for roughly a day-long period once every six months during which WISE is operational. We segment W1 exposures into such visits, stacking together those frames grouped within the same visit as each other. The typical timespan of such a 'coadd epoch' at a particular sky location is roughly (1 day)/cos($|\beta|$). We implement special handling of situations in which WISE deviates from its usual scan pattern to avoid pointing too closely toward the Moon. Our careful handling of such Moon avoidance maneuvers splits what would be overly long coadd epochs into pairs of standard coadd epochs occurring before/after each maneuver. WISE scans converge at the ecliptic poles, which therefore receive nearly continuous coverage.

The primary disadvantage of our coaddition is that the parallactic motion within a single coadd epoch's timespan can smear out Planet Nine's appearance and therefore limit the sensitivity of our search. The effect of this limitation on the parameter space we are able to probe is illustrated in Figure 1 of M17.

3.3 Coaddition

We use the unWISE code (Lang 2014) to generate W1 coadds which incorporate the time-slicing described in §3.2 and retain the native WISE angular resolution. Complete details of the unWISE code and the particular options we use in the context of our W1 Planet Nine searches can be found in Lang (2014), Meisner et al. (2017b) and M17. The number of coadd epochs generated per sky location is shown in the top panel of Figure 1. In total, our present search incorporates 129,483 time-resolved coadds. For comparison, the bottom panel of Figure 1 shows the number of coadd epochs per sky location in the M17 search, which included only 6,219 such coadds.

3.4 Reference Templates

We create a reference template image corresponding to each time-resolved coadd. The reference template is constructed by combining all coadd epochs of the relevant `coadd_id` footprint which are temporally separated by more than 5 months relative to the epoch being searched for transients. Typically, 6–8 epochal coadds contribute to each reference template.

3.5 Image Differencing

Our image differencing pipeline is built around a set of standard astronomical image analysis software packages. Source Extractor (Bertin & Arnouts 1996) and SCAMP (Bertin 2006) are used to derive astrometric tweaks that ensure the best possible alignment of the reference and search images. SWARP (Bertin et al. 2002) is used to warp the reference based on the SCAMP outputs. HOTPANTS[1] is used to perform kernel-matched subtractions that account for small (percent level) changes in the PSF and/or multiplicative differences between the reference stack and the search image.

3.6 Extracting Transients

We use Source Extractor to detect transients within our difference images. Centroids are measured using the `XWIN_IMAGE` and `YWIN_IMAGE` parameters.

3.7 Filtering Transients

We go to great lengths to clean identifiable artifacts out our transient catalogs while masking the smallest possible fraction of sky area. The transients we remove include those associated with:

- static sources, as identified by 2MASS counterparts.
- kernel matching artifacts.
- the cores and/or diffraction spikes of bright star profiles.
- latents. These are persistence artifacts which appear at the *detector* position of saturated stars in subsequent exposures. Latents have the appearance of diffuse smudges, and are detected as spurious transients because their positions vary with WISE's scan direction.

At high $|b_{gal}|$ ($|b_{gal}| > 30°$), our filtering masks 1.0% of the area. This fraction varies dramatically over the sky, mainly due to variations in the number of background static sources within the Galaxy along different sightlines (see Figure XX). In §6, we attempt to account for the impact of the spatially varying fraction of area masked on our completeness.

3.8 Orbit Linking

4 SEARCH REGION

Although WISE observes the entire sky, some regions have been omitted from our search. As illustrated in Figure 1 of M17, the minimum d_9 for which our search method is sensitive increases much faster with $|\beta|$ at high ecliptic latitude

[1] www.astro.washington.edu/users/becker/v2.0/hotpants.html

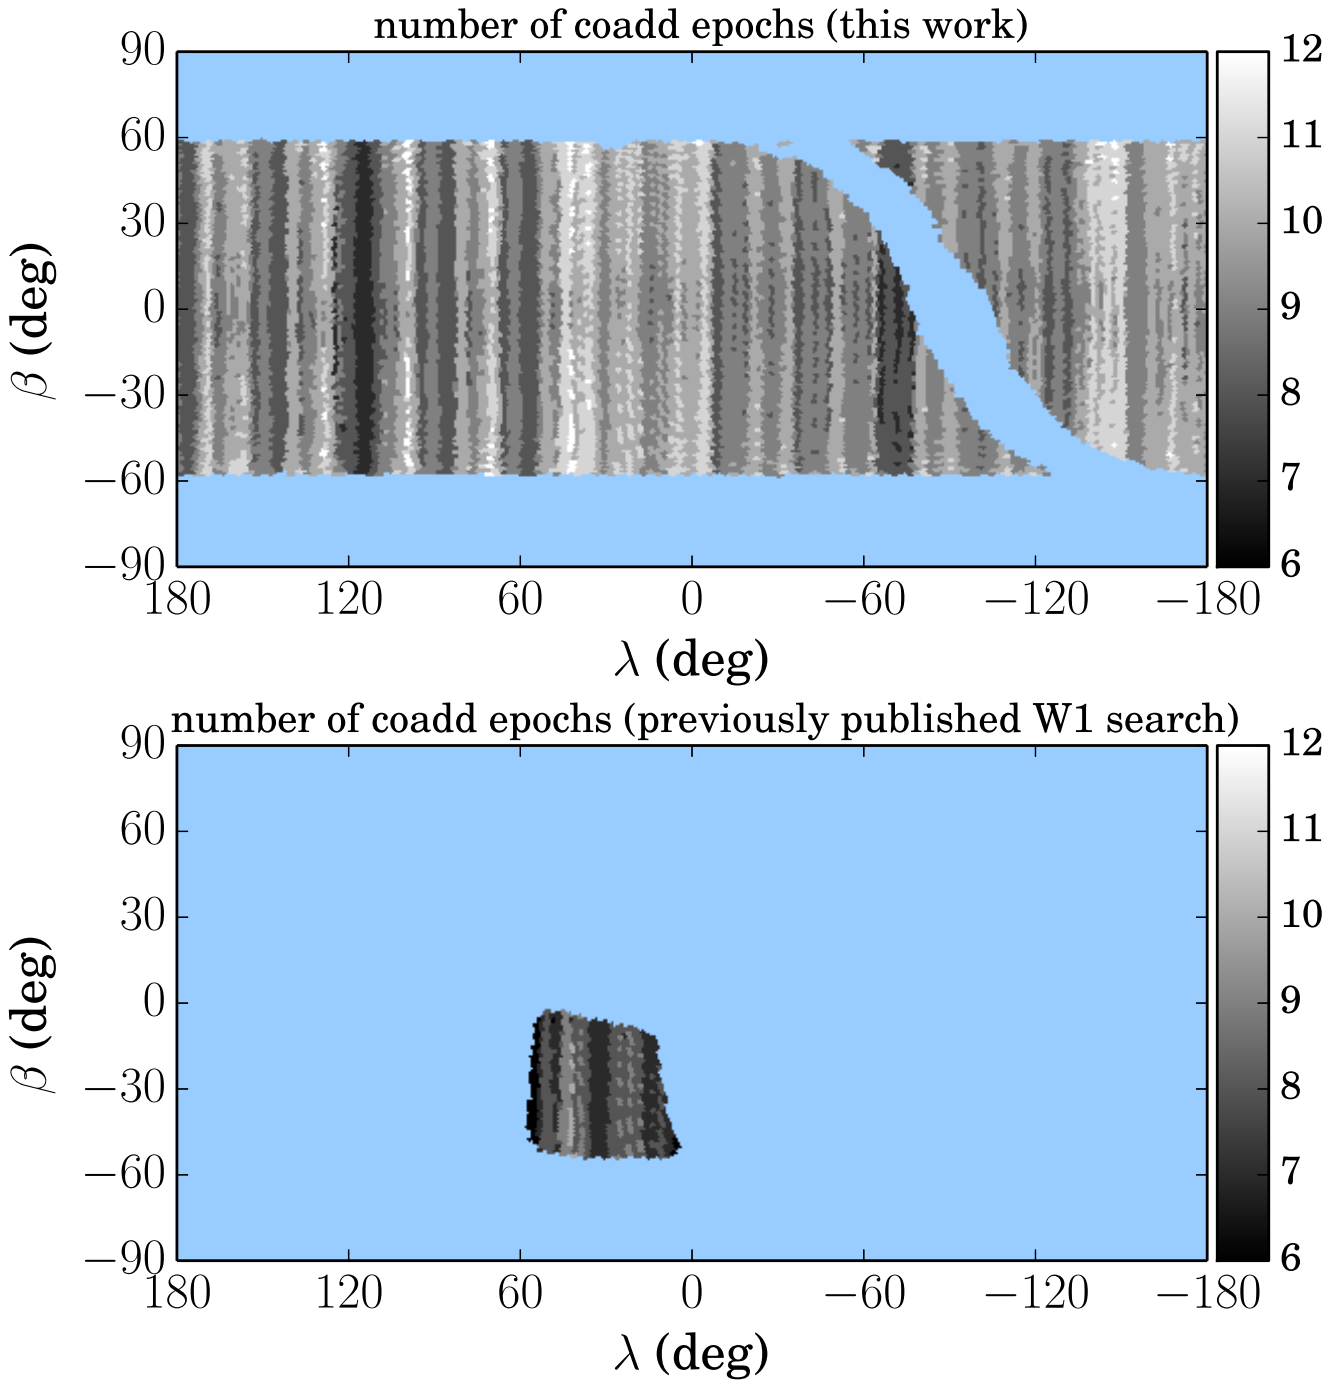

Figure 1. Maps of the number of W1 coadd epochs available. Top: present search, covering 76% of the sky and incorporating ∼4 years of W1 observations. Bottom: corresponding map for previously published M17 search, for the sake of comparison. M17 searched a much smaller footprint (∼1,840 square degrees), and even within that area employed fewer epochs per sky location because the third-year NEOWISER data release was not yet available. The present search incorporates 129,483 time-resolved coadds, whereas that of M17 incorporated only 6,219 epochal coadds. The entire sky is shown in both panels. As can be seen in the top panel, our present search omits tiles at $|\beta| > 55°$. The diagonal missing stripe centered near $(\lambda, \beta) = (-90°, 0°)$ is due to our omission of low $|b_{gal}|$ regions toward the inner Galaxy.

than does the maximum d_9 for which our search is sensitive. There exists such a minimum d_9 value because the timespan of a given W1 coadd epoch increases as $1/\cos(|\beta|)$, resulting in excessive parallactic smearing for relatively nearby objects. Furthermore, our orbit linking (§5) is designed to work only for $d_9 > 250$ AU, as parallactic smearing would render us relatively insensitive at smaller distances, and resolving orbits at smaller distances requires linking transients over progressively wider sky areas, which becomes computationally prohibitive. At $|\beta| \approx 30°$, the minimum d_9 imposed by parallactic smearing becomes larger than 250 AU and therefore sets our search's lower distance limit at $|\beta| > 30°$. The maximum d_9 probed is determined by the maximum plausible W1 luminosity of Planet Nine, which corresponds to W1 = 16.1 at a fiducial distance of 622 ($H_{W1} = 2.13$) according to Fortney et al. (2016). This translates to a maximum detectable distance that ramps upward from ∼800 AU at $|\beta| = 0°$ to ∼900 AU at $|\beta| \approx 65°$.

For $|\beta| < 30°$, we are sensitive throughout the distance range $250 \lesssim d_9/\mathrm{AU} \lesssim 800$. As $|\beta|$ increases beyond 30°, the distance range within which we are sensitive to Planet Nine narrows rapidly. By $|\beta| \sim 65°$ we effectively lose all sensitivity (the maximum and minimum detectable d_9 converge at $|\beta| \sim 65°$, assuming $H_{W1} \geq 2.13$). At $|\beta| \sim 55°$, only a relatively narrow range of d_9 is probed ($650 \lesssim d_9/\mathrm{AU} \lesssim 850$). We therefore restrict our search to tiles with central coordinates interior to the region defined by $|\beta| < 55°$, or within 1.1° of the $|\beta| = 55°$ boundary[2]. This cut removes 3,182 of 18,240 tiles, corresponding to 17.4% percent of the sky. We note that by inventing specially tailored time-slicing rules to define our coadd epoch boundaries at very high ecliptic latitude, we could make use of all WISE data up to $|\beta| = 90°$. However, the Planet Nine narrative disfavors very high inclinations, and $i > 56°$ would be required to observe Planet Nine in the high $|\beta|$ sky area we have excluded.

We ran all stages of our processing up to and inclusive of transient filtering for all 15,058 tile footprints satisfying our ecliptic latitude cut. However, we found that an excessive number density of transients remained in our filtered catalogs in low Galactic latitude regions toward the inner Galaxy. The increased transient number density near the Galactic center is due to an increase in the overall number density of static sources in these regions, and also the increased number of bright stars in these regions. Extreme, localized increases in transient density are particularly problematic because of the combinatorics associated with orbit linking – the number of transient n-tuples increases with the n^{th} power of transient number density. We therefore sought to excise regions near the Galactic center containing huge numbers of static background sources. We did so by making a cut on number of 2MASS (Skrutskie et al. 2006) sources per tile footprint. We removed tiles that contained >90,000 2MASS sources within their ∼2.5 square degree area. This corresponds to 1 source per 47 pixels. For comparison, the effective number of pixels corresponding to the W1 PSF is 15.8, meaning that our chosen source density threshold translates to roughly 0.34 2MASS sources per WISE beam. We used 2MASS source density as a proxy for crowding rather than WISE source density because 2MASS has far superior angular resolution; in general sources detected by 2MASS are also detected by WISE (WISE is deeper), though deblending issues may modulate/reduce the number of WISE sources actually catalogued in the densest regions. 1,131 tiles are rejected due to our cut on number of 2MASS sources. These tiles trace out a roughly diamond-shape area centered around $(l_{gal}, b_{gal}) = (0°, 0°)$ when viewed in Galactic coordinates. This region is shown as the empty (light blue) stripe passing through the Galactic center at $(\lambda, \beta) = (267°, -6°)$ in Figures 1 and 3. The region excised due to high source density reaches a maximum $|b_{gal}|$ of ∼14° near $l_{gal} = 0°$, then tapers toward the Galactic plane, disappearing by $l_{gal} = 90°$ toward the east and $l_{gal} = 282°$ toward the west. We are able to perform our search all the way into the Galactic plane (through $b_{gal} = 0°$) in the outer Galaxy – our background source density cut is lenient enough to retain most of the outer Galactic plane.

We also found that a region of the Galactic plane near $(l_{gal}, b_{gal}) = (109°, -2°)$ exhibited an unusually high number density of transients associated with diffuse emission. We presume this to result from time-variable illumination of dust clouds. As a result, we removed eight nearby tiles, with `coadd_id` values 3461p575, 3484p590, 3489p575, 3510p605, 3513p590, 3516p575, 3540p605, 3542p590.

External galaxies with exceptionally large apparent sizes also lead to spikes in the transient number density. In particular, we excised elliptical regions surrounding M31, Triangulum, M110 and IC10. The combined area of these elliptical external galaxy masks is 6.6 square degrees.

In M17, we created custom masks for four 'ultra-bright' stars whose profiles extend far beyond the PSF model used in our standard bright star masking approach. In this work, we opted not to continue the practice of making such custom masks for bright stars throughout our search footprint.

After deducting the areas masked at high ecliptic latitude, near the Galactic center, near time-variable ISM reflection features and surrounding external galaxies, we are left with a 31,480 square degree search region. This corresponds to 76.3% of the sky.

5 ORBIT LINKING PROCEDURE

Our orbit linking strategy follows the procedure we outlined in M17, with some modification to handle the increased number of epochs and detections. First, for each detection, we identify a pool of all nearby subsequent detections that might be associated with bound Keplerian orbits and parallax distances of 250 AU or more. We then construct triplets that consist of the first detection and pairs of nearby ones from the pool around it. We then test if each triplet is consistent with a simple model of reflex motion plus linear drift in the sky plane, and we accept it if the model fit's χ^2 is less than 30. We repeat this process for quadruplets built from accepted triplets and a fourth detection from the pool. Quadruplets themselves are tested with the parallax-linear drift model, but with a χ^2 threshold of 150. These steps, which convert the full list of detections into potential linkages of quadruplets, constitute our prescreening process.

We test the prescreening process with mock data. By

[2] 1.1° is the distance between the center of a tile and its corners. Therefore, tiles that have centers up to 1.1° beyond either $|\beta| = 55°$ boundary could potentially contain sub-regions at $|\beta| < 55°$.

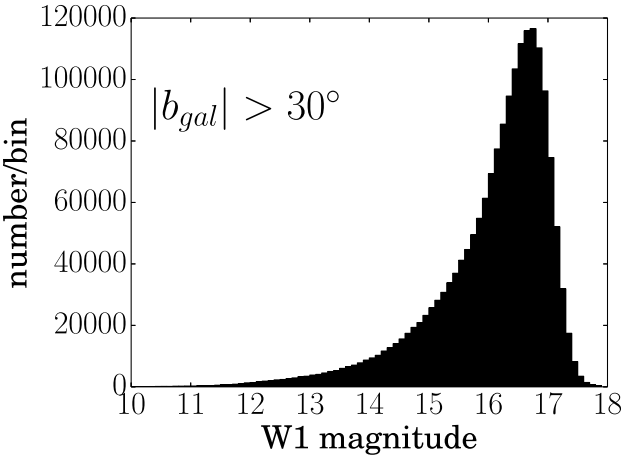

Figure 2. Histogram of W1 magnitudes for detections in our filtered transient catalog. The source counts peak near W1=16.6, far fainter than the single-exposure detection limit at W1 = 15.3.

sampling bound Keplerian orbital parameters for hypothetical bound solar system objects beyond 250 AU, we generate mock detections with similar cadence and astrometric uncertainties as the real data. Applying the same parallax and linear drift tests on triplets and quadruplets as we do for real data, we find that our prescreen step falsely rejects no more than 1 in 10^4 linkages.

In the next step of the linkage procedure, we evaluate each prescreened quadruplet with `orbfit`, a code designed by Bernstein & Khushalani (2000) to fit observations of outer solar system objects with Keplerian orbital elements. We choose a χ^2 threshold of 40 in the decision to accept or reject a quadruplet, irrespective of order or degrees of fitting (`orbfit` uses several different models, depending on the details of the detections). This choice yeilds failure rates that are well below the value of 0.001 adopted in M17.

Finally, we check if the quadruplets that are admitted on the basis of `orbfit` have any detections in common. In cases where they do, we merge them together to form quintuplets, or even higher-order tuples; if the merged `orbfit` χ^2 is less than 40, as before. For all of the linkages that we find, the number of degrees of freedom is low enough that this χ^2 threshold preserves a false detection rate below the 0.001 threshold.

This procedure culls over 5 million detections into 2,560 quintuplets, including the members of 175 of higher-order n-tuples (n = 6–8). Along the way, the prescreen process enumerated over 92 billion triplets, 275 million quadruplets, and ran the `orbfit` code on over 8 million potential linkages.

6 COMPLETENESS ANALYSIS

6.1 Cautionary Notes on Completeness

very edges of footprint – not sure what the limits mean here

7 RESULTS

8 CONCLUSION

We have performed the deepest and widest area WISE-based search for Planet Nine. We rule out the presence of Planet Nine in the parameter space searched at $W1 < 16.7$ (90% completeness).

mention W2

mention that WISE is still going so there will be at least another year of data to fold in

ACKNOWLEDGEMENTS

This work has been supported by grant NNH17AE75I from the NASA Astrophysics Data Analysis Program.

This research makes use of data products from the Wide-field Infrared Survey Explorer, which is a joint project of the University of California, Los Angeles, and the Jet Propulsion Laboratory/California Institute of Technology, funded by the National Aeronautics and Space Administration. This research also makes use of data products from NEOWISE, which is a project of the Jet Propulsion Laboratory/California Institute of Technology, funded by the Planetary Science Division of the National Aeronautics and Space Administration. This research has made use of the NASA/ IPAC Infrared Science Archive, which is operated by the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration.

The National Energy Research Scientific Computing Center, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231, provided staff, computational resources, and data storage for this project.

REFERENCES

Bailey E., Batygin K., Brown M. E., 2016, AJ, 152, 126
Bannister M. T., et al., 2017, AJ, 153, 262
Batygin K., Brown M. E., 2016a, AJ, 151, 22
Batygin K., Brown M. E., 2016b, ApJ, 833, L3
Bernstein G., Khushalani B., 2000, AJ, 120, 3323
Bertin E., 2006, in Gabriel C., Arviset C., Ponz D., Enrique S., eds, Astronomical Society of the Pacific Conference Series Vol. 351, Astronomical Data Analysis Software and Systems XV. p. 112
Bertin E., Arnouts S., 1996, A&AS, 117, 393
Bertin E., Mellier Y., Radovich M., Missonnier G., Didelon P., Morin B., 2002, in Bohlender D. A., Durand D., Handley T. H., eds, Astronomical Society of the Pacific Conference Series Vol. 281, Astronomical Data Analysis Software and Systems XI. p. 228
Brown M. E., 2017, AJ, 154, 65
Brown M. E., Batygin K., 2016, ApJ, 824, L23
Brown M. E., Trujillo C., Rabinowitz D., 2004, ApJ, 617, 645
Brown M. E., et al., 2015, AJ, 149, 69
Cowan N. B., Holder G., Kaib N. A., 2016, ApJ, 822, L2
Cutri R. M., et al., 2011, Technical report, Explanatory Supplement to the WISE Preliminary Data Release Products
Drake A. J., et al., 2009, ApJ, 696, 870
Folkner W., Jacobson R. A., Park R., Williams J. G., 2016, in AAS/Division for Planetary Sciences Meeting Abstracts. p. 120.07

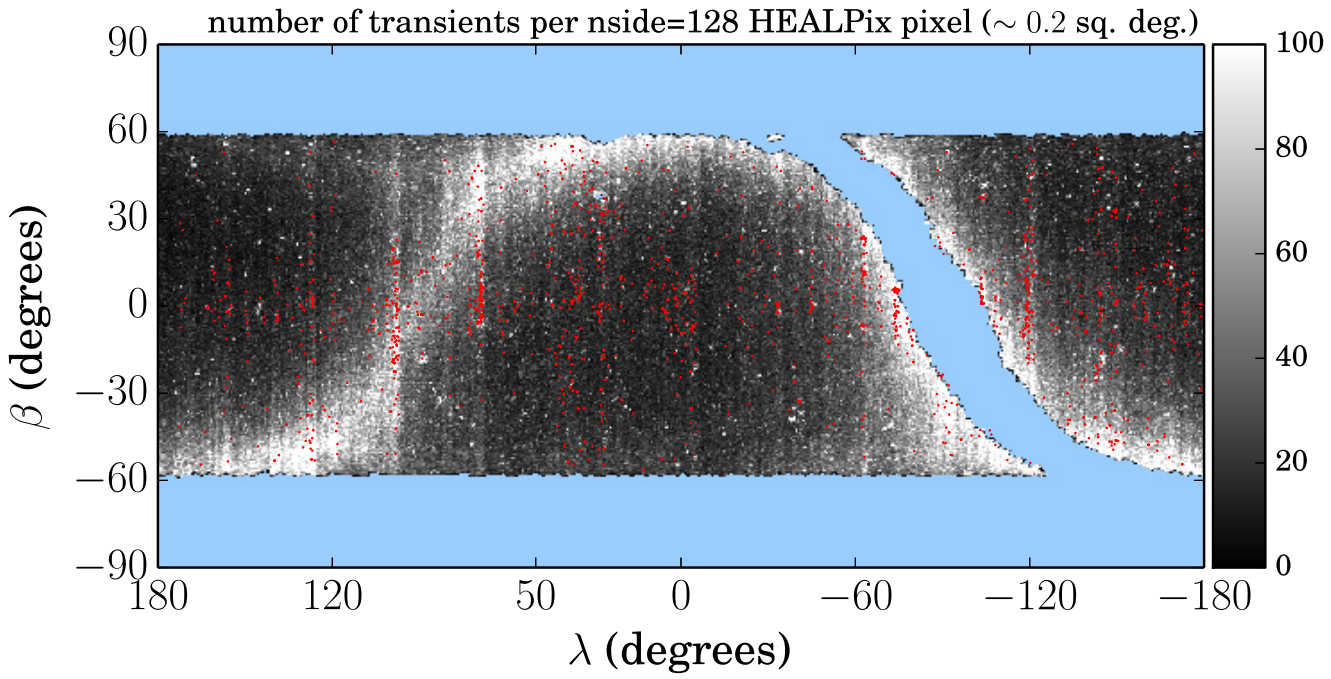

Figure 3. Number density of transients throughout our search footprint. The scale bar is in units of transients per N_{side} = XX HEALPix pixel. Each such pixel is ∼0.2 square degrees in size. The locations of our 2,560 quintuplet linkages with $\chi^2 < 40$ are marked with red dots.

Fortney J. J., et al., 2016, ApJ, 824, L25
Gomes R., Deienno R., Morbidelli A., 2017, AJ, 153, 27
Holman M. J., Payne M. J., 2016, AJ, 152, 94
Kuchner M. J., et al., 2017, ApJ, 841, L19
Lai D., 2016, AJ, 152, 215
Lang D., 2014, AJ, 147, 108
Linder E. F., Mordasini C., 2016, A&A, 589, A134
Mainzer A., et al., 2011, ApJ, 731, 53
Mainzer A., et al., 2014, ApJ, 792, 30
Meisner A. M., Bromley B. C., Nugent P. E., Schlegel D. J., Kenyon S. J., Schlafly E. F., Dawson K. S., 2017a, AJ, 153, 65
Meisner A. M., Bromley B. C., Nugent P. E., Schlegel D. J., Kenyon S. J., Schlafly E. F., Dawson K. S., 2017b, AJ, 153, 65
Millholland S., Laughlin G., 2017, AJ, 153, 91
Shankman C., et al., 2017, AJ, 154, 50
Sheppard S. S., Trujillo C., 2016, AJ, 152, 221
Skrutskie M. F., et al., 2006, AJ, 131, 1163
Trujillo C. A., Sheppard S. S., 2014, Nature, 507, 471
Weryk R. J., et al., 2016, preprint, (arXiv:1607.04895)
de la Fuente Marcos C., de la Fuente Marcos R., 2016, MNRAS, 462, 1972