

UC Berkeley

UC Berkeley Previously Published Works

Title

SEARCHING FOR PLANET NINE WITH COADDED WISE AND NEOWISE-REACTIVATION IMAGES

Permalink

<https://escholarship.org/uc/item/4c93q8h9>

Journal

The Astronomical Journal, 153(2)

ISSN

0004-6256

Authors

Meisner, Aaron M
Bromley, Benjamin C
Nugent, Peter E
[et al.](#)

Publication Date

2017-02-01

DOI

10.3847/1538-3881/153/2/65

Peer reviewed

SEARCHING FOR PLANET NINE WITH COADDED WISE AND NEOWISE-REACTIVATION IMAGES

AARON M. MEISNER^{1,2}, BENJAMIN C. BROMLEY³, PETER E. NUGENT^{2,4}, DAVID J. SCHLEGEL², SCOTT J. KENYON⁵,
EDWARD F. SCHLAFLY^{2,6}, AND KYLE S. DAWSON³

Draft version November 2, 2016

ABSTRACT

A distant, as yet unseen ninth planet has been invoked to explain various observations of the outer solar system. While such a ‘Planet Nine’, if it exists, is most likely to be discovered via reflected light in the optical, it may emit much more strongly at 3–5 μ m than simple blackbody predictions would suggest, depending on its atmospheric properties (Fortney et al. 2016). As a result, Planet Nine may be detectable at 3.4 μ m with WISE, but single exposures are too shallow except at relatively small distances ($d_9 \lesssim 430$ AU). We develop a method to search for Planet Nine far beyond the W1 single-exposure sensitivity, to distances as large as 800 AU, using inertial coadds of W1 exposures binned into ~ 1 day intervals. We apply our methodology to ~ 2000 square degrees of sky identified by Holman & Payne (2016) as a potentially likely Planet Nine location, based on the Fienga et al. (2016) Cassini ranging analysis. We do not detect a plausible Planet Nine candidate, but are able to derive a detailed completeness curve, ruling out its presence within the parameter space searched at $W1 < 16.66$ (90% completeness). Our method uses all publicly available W1 imaging, spanning 2010 January to 2015 December, and will become more sensitive with future NEOWISE-Reactivation releases of additional W1 exposures. We anticipate that our method will be applicable to the entire high Galactic latitude sky, and we will extend our search to that full footprint in the near future.

Subject headings: surveys: trans-Neptunian objects – techniques: image processing

1. INTRODUCTION

Beginning with the discovery of Sedna (Brown et al. 2004), distant trans-Neptunian objects (TNOs) with perihelia far beyond Neptune have posed a challenge to existing models of the solar system. Among other possibilities, Brown et al. (2004) suggested that Sedna’s orbit may be indicative of an as yet undiscovered planet in the outer solar system. Upon discovering the second sednoid, 2012 VP₁₁₃, Trujillo & Sheppard (2014) proposed the presence of a super-Earth perturber at several hundred AU (see also Sheppard & Trujillo 2016). Batygin & Brown (2016) elaborated on this scenario by showing that a ‘Planet Nine’ could simultaneously account for sednoids and the clustering of distant TNOs, using extensive simulations to significantly constrain the hypothetical planet’s feasible mass range (5–20 M_{\oplus}) and orbital parameters ($380 < a/\text{AU} < 980$, Brown & Batygin 2016).

As discussed in Brown & Batygin (2016), Planet Nine, if it exists, is most likely to be detected via reflected sunlight in the optical, where its plausible range of brightnesses generally falls within reach of large ground-based telescopes ($22 < V < 25$). Nevertheless, predictions of detectability at other, longer wavelengths have also been made. Cowan et al. (2016) showed that a Neptune-sized

~ 40 K blackbody at many hundred AU would be detectable at millimeter wavelengths with current and future CMB mapping experiments. Linder & Mordasini (2016) suggested that Planet Nine would likely be self-luminous. Fortney et al. (2016) considered detailed atmospheric models, finding that for certain levels of methane depletion, Planet Nine could have enormously enhanced 3–5 μ m emission relative to simple blackbody predictions. In the most optimistic models, Planet Nine would be detectable in the WISE W1 (3.4 μ m, Wright et al. 2010) channel, but would be too faint for single exposure detection unless relatively nearby ($d_9 \lesssim 430$ AU).

Here we develop a method to search for Planet Nine using coadded W1 images, extending ~ 1.4 magnitudes below the single-exposure detection limit. Although Planet Nine’s infrared emission may well be orders of magnitudes too faint for detection by WISE, this full-sky dataset nevertheless deserves to be systematically searched; W1 images spanning a nearly six year baseline have already been acquired and publicly released, and provide time sampling sufficient to enable Keplerian orbit linking. Moreover, the W1 image quality is both extremely high in general, and very uniform with time (Mainzer et al. 2014; Cutri et al. 2015). The WISE imaging data are therefore conducive to establishing fully-characterized constraints on Planet Nine’s apparent infrared brightness over most of the sky in the event of a nondetection.

Given the limited range of likely Planet Nine inclinations, its potential paths across the sky are fairly well-constrained (Brown & Batygin 2016). Several predictions have been made for Planet Nine’s present-day location (Fienga et al. 2016; Holman & Payne 2016; de la Fuente Marcos & de la Fuente Marcos 2016). Holman & Payne (2016) identified ~ 1200 square degrees

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

⁴ Astronomy Department, University of California Berkeley, Berkeley, CA, 94720, USA

⁵ Smithsonian Astrophysical Observatory, 60 Garden Street, Cambridge, MA 02138, USA

⁶ Hubble Fellow

of southern sky toward the constellation Cetus as a high-probability Planet Nine location, by building on the Cassini ranging analysis of Fienga et al. (2016). Although alternative analyses of the Cassini ranging data remain ongoing, we nevertheless target the Holman & Payne (2016) preferred sky region, which at the very least provides a testbed area for developing and demonstrating our methodology.

The only WISE search for Planet Nine has been conducted by Fortney et al. (2016), filtering the AllWISE source catalog in/near the Holman & Payne (2016) favored sky region. However, the AllWISE catalog only utilizes WISE data spanning 2010 January to 2011 February, representing just $\sim 1/3$ of currently available WISE imaging. Moreover, it was not possible to characterize the completeness of this search – doing so would require simulating the AllWISE catalog-making software’s interaction with a faint moving object that appears at certain epochs but disappears at others.

Our basic search strategy is to generate inertial coadds of W1 exposures binned into ~ 1 day intervals, create a reference stack for each epochal coadd, extract a transient catalog via difference imaging, and finally link transients into Keplerian orbits. Searching for transients within time-resolved coadds allows for deeper and cleaner difference images than would be possible with single exposures, though this choice does restrict the parameter space over which we are sensitive to TNOs. Because our search consists of a custom reprocessing that begins with raw WISE exposures, we can use all publicly available imaging, spanning 2010 January through 2015 December, and can perform injection tests to thoroughly characterize our completeness.

In §2 we briefly review the WISE and NEOWISE-Reactivation missions. In §3 we describe the design considerations and forecasted sensitivity of our search method. In §4 we list the data products that form the basis for our search. In §5 we detail our image coaddition procedures. In §6 we describe our difference imaging pipeline. In §7 we explain our orbit linking methodology. In §8 we analyze our completeness. In §9 we discuss the results of our search, and conclude in §10.

2. WISE & NEOWISE-REACTIVATION OVERVIEW

The Wide-field Infrared Survey Explorer (WISE; Wright et al. 2010) has surveyed the full sky at four infrared wavelengths: W1= $3.4\mu\text{m}$, W2= $4.6\mu\text{m}$, W3= $12\mu\text{m}$ and W4= $22\mu\text{m}$. The satellite resides in a ~ 95 minute period Sun-synchronous low-Earth polar orbit, and the field of view is $0.78^\circ \times 0.78^\circ$.

Launched in late 2009, WISE mapped the entire sky in all of W1-W4 between 7 January 2010 and 6 August 2010. As solid hydrogen cryogen was exhausted, W4 became unusable in early August 2010, as did W3 in late September 2010. During the NEOWISE mission phase (Mainzer et al. 2011), WISE continued surveying the sky in W1 and W2 until early February 2011, when the spacecraft was placed in hibernation. WISE was reactivated in October 2013, and returned to surveying the sky in W1 and W2. This W1/W2 survey, dubbed NEOWISE-Reactivation (NEOWISER), is planned to last through the end of 2016. NEOWISER first and second year data products, including all single-exposure images, have been publicly released in March 2015 and March 2016, respec-

TABLE 1
SEASON DEFINITIONS

Season	MJD _{min}	MJD _{max}	Calendar Months	Parity
a	55203	55239	2010 Jan – 2010 Feb	0
b	55370	55430	2010 Jun – 2010 Aug	1
c	55550	55593	2010 Dec – 2011 Feb	0
d	56648	56704	2013 Dec – 2014 Feb	0
e	56835	56895	2014 Jun – 2014 Aug	1
f	57012	57063	2014 Dec – 2015 Feb	0
g	57197	57256	2015 Jun – 2015 Aug	1

tively. NEOWISER images, especially in W1, have essentially the same sensitivity and overall high quality as those of the primary WISE mission (Mainzer et al. 2014).

Because W1 single-exposure images represent the primary input to our search, we can make use of WISE data spanning all mission phases, from 2010 January to 2015 December. Importantly, the W1 PSF is very broad relative to typical optical data, with $6.1''$ FWHM. This means that Planet Nine could appear very nearly point-like, even in W1 coadds which bin exposures into $\gtrsim 1$ day intervals.

3. SURVEY GEOMETRY/CADENCE CONSIDERATIONS

3.1. WISE “Visits”

WISE surveys the sky by remaining pointed at very nearly 90° solar elongation, completing one 360° scan per orbit while approximately following a great circle of constant ecliptic longitude. As a result, away from the ecliptic poles, a given sky region will intersect the WISE field of view at six-monthly intervals, each time remaining in the 0.78° FOV for ~ 1 day, during which one exposure is acquired every ~ 95 minutes. We refer to such day long periods of WISE observation as “visits”.

Here we search for Planet Nine by creating per-visit inertial coadds of W1 single exposures. Because we use data from all phases of the WISE+NEOWISER mission, we typically have 7 such epochal coadds available per sky location within our search footprint. This number of coadd epochs results from the fact that WISE has visited our search footprint during seven distinct “seasons”, which we label with letters *a-g* as defined in Table 1. A linkage of transients at four epochs along a Keplerian orbit is generally considered the minimum threshold to identify a TNO candidate. Therefore, our ~ 7 coadded W1 epochs are sufficient to identify or place meaningful constraints on the presence of Planet Nine candidates.

Table 1 also includes a “parity” column, which essentially serves as a binary indicator of Earth’s location relative to the Sun during each season. Because WISE always points at very nearly solar elongation of 90° , observations of a particular sky location with the same parity have almost exactly zero relative parallax, whereas observations with opposite parity have relative parallax of approximately 2 AU.

The foregoing depiction of WISE’s survey strategy is overly simplistic because it ignores Moon avoidance maneuvers, during which WISE can point many degrees away from elongation 90° . Our handling of these events is discussed in detail in §5.2. The remainder of this section’s formulae and forecasts will ignore the complication of Moon avoidance maneuvers, although in practice these actually provide us with the benefit of extra coadd

epochs. As a result, we have on average 7.3 epochal coadds per sky location in the search region studied.

3.2. Parallax Within a Coadd Epoch

Apparent motion due to parallax during the course of a coadd epoch places restrictions on the range of distances at which a TNO will be detectable in our coadds, since moving sources become smeared out over an area larger than that of a point source, effectively decreasing their signal-to-noise. To determine the extent of parallactic smearing, we must first calculate the expected timespans of our epochal coadds.

Consider an idealized version of the WISE survey strategy, assuming WISE always points at exactly 90° solar elongation. At ecliptic latitude (for which we use the symbol β) of zero, a given sky location will be observed over a period of $0.776^\circ \times (365.25 \text{ days}/360^\circ) = 0.787$ days during a WISE visit. In reality, WISE does not point at precisely 90° elongation, but toggles back and forth by a few tenths of a degree on even/odd orbits⁷. This effectively lengthens the $\beta=0^\circ$ visit duration to 1.0 ± 0.05 days, based on examination of per-pixel maps of minimum and maximum MJD for real coadds.

WISE obtains a constant number of frames per unit ecliptic latitude, and therefore the number of exposures per unit area within a visit scales as $1/\cos(|\beta|)$, as does the time spanned by exposures overlapping a given point on the sky:

$$\tau_{\text{coadd}} = \frac{1.0 \text{ days}}{\cos(|\beta|)} \quad (1)$$

Ignoring the curvature of the Earth's orbit, Planet Nine's parallax during a single visit of WISE observations is:

$$\pi_9 = \frac{dl_\perp}{d_9} = \frac{(2\pi \text{ AU}) \times \frac{\tau_{\text{coadd}}}{365.25 \text{ days}} \times \sin(|\beta|)}{d_9} \quad (2)$$

Where d_9 is the distance to Planet Nine in AU, and the $\sin(|\beta|)$ factor projects the Earth's motion onto the direction perpendicular to the line of sight. The combination of the previous two equations can be rewritten as:

$$\pi_9 = 5.91'' \times \frac{600 \text{ AU}}{d_9} \times \tan(|\beta|) \quad (3)$$

The smearing caused by parallactic motion means that there will be a *minimum* distance at which our coadd-based search is sensitive to Planet Nine. For instance, even in the absence of orbital drift, an object in the Kuiper belt at $|\beta|=15^\circ$ and 35 AU from the Earth would appear to move by $\sim 27''$ within a coadd epoch, making it very highly streaked. On the other hand at a fiducial distance of 600 AU and the same $|\beta|$, Planet Nine would experience only $1.6''$ of parallactic motion, rendering its W1 profile very nearly pointlike.

3.3. Effect of Parallactic Smearing

To quantify and account for parallactic smearing, we compute the effective number of noise pixels of a TNO profile as a function of parallax within a coadd epoch.

The effective number of noise pixels for a particular profile is given by:

$$n_{\text{eff}} = \frac{(\sum f_i)^2}{\sum f_i^2} \quad (4)$$

Where f_i is the flux in each pixel i . We simulate Planet Nine profiles that translate by an amount π_9 during the course of a coadd epoch with 12 exposures, for $0'' < \pi_9 < 8.25''$. For these simulations, we use the W1 PSF constructed according to Meisner & Finkbeiner (2014), as shown in Figure 4 of Lang (2014). The resulting n_{eff} values are well fit by:

$$\frac{n_{\text{eff}}}{n_{\text{eff}}(\pi_9 = 0)} = 1 - (1.42 \times 10^{-3})\pi_9 + (9.59 \times 10^{-3})\pi_9^2 - (2.91 \times 10^{-4})\pi_9^3 \quad (5)$$

With π_9 in arcseconds. $8.25''$ corresponds to $3 \times$ the native WISE pixel size, and $1.35 \times$ the W1 FWHM, but even so $n_{\text{eff}}(\pi_9=8.25'')$ remains less than $1.5 \times$ that of a W1 point source. We choose $\pi_9=8.25''$ as a threshold beyond which we will not claim any sensitivity of our search, since we do not aim to push into the regime of detecting/photometering highly elongated sources. We nevertheless believe that we have considerable sensitivity to less pointlike sources (see §6.3.7). We note that a typical galaxy ($r_{\text{half}}=0.45''$) in $1''$ seeing has n_{eff} $1.99 \times$ larger than that of a Gaussian PSF – by restricting to $\pi_9 \leq 8.25''$, we remain strongly within the quasi-pointlike regime for which standard source detection and photometry techniques are applicable.

3.4. Forecasted Sensitivity

Our per-visit coadds very closely resemble the WISE Preliminary Release Atlas stacks in terms of depth of coverage and sensitivity. In unconfused regions of the ecliptic plane with typical integer coverage (12 exposures), the Preliminary release 5σ detection limit for static point sources is $W1_{\text{lim}, \text{PSF}} = 16.5$ (Cutri et al. 2011). We always quote WISE magnitudes in the Vega system⁸. We can calculate $W1_{\text{lim}}(|\beta|, d_9)$ by scaling this static point source limiting magnitude to account for parallactic smearing via its influence on $n_{\text{eff}}(|\beta|, d_9)$, as well as the increasing number of exposures with increasing $|\beta|$:

$$W1_{\text{lim}}(|\beta|, d_9) = W1_{\text{lim}, \text{PSF}} + \Delta_{\text{ref}} - 2.5 \log_{10} \left[\sqrt{\frac{n_{\text{eff}}(|\beta|, d_9)}{n_{\text{eff}}(\pi_9 = 0)}} \sqrt{\frac{n_{\text{exp}}(\beta = 0)}{n_{\text{exp}}(|\beta|)}} \right] \quad (6)$$

Where n_{exp} is the number of W1 exposures contributing to a single sky location within an epochal coadd. Allowing n_{exp} to represent an average capable of taking on fractional values, we adopt:

$$n_{\text{exp}}(|\beta|) = \frac{n_{\text{exp}}(\beta = 0)}{\cos(|\beta|)} \quad (7)$$

We will search for transients in difference images, rather than the epochal coadds themselves (§6). The

⁷ wise2.ipac.caltech.edu/docs/release/allsky/expsup/sec3.4.html

⁸ wise2.ipac.caltech.edu/docs/release/allsky/expsup/sec4.4h.html




FIG. 1.— Parameter space accessible to our Planet Nine search based on time-resolved W1 coadds. The black region is expected to be accessible based on the forecast of §3.4. Both the black and dark gray regions are accessible, based on the final completeness analysis of §8. The light gray area is also accessible, extending up to 2250 AU, but would require a W1 luminosity brighter than the most optimistic Fortney et al. (2016) model.

Δ_{ref} term represents a small loss in sensitivity to due to the noise added by subtracting the reference coadd, which typically has $6\times$ more contributing exposures than the epochal coadd. In that case:

$$\Delta_{ref} = -2.5 \log_{10} \left(\sqrt{1 + 1/6} \right) = -0.084 \text{ mag} \quad (8)$$

We can now evaluate $W1_{lim}(|\beta|, d_9)$ and compare against Planet Nine’s minimum W1 apparent magnitude as a function of d_9 , yielding a mask for accessible versus inaccessible regions of $(|\beta|, d_9)$ parameter space. We calculate the minimum W1 magnitude as a function of d_9 by assuming a maximum W1 luminosity of $H_{W1}=2.13$ ($W1=16.1$ at $d_9=622$ AU, Fortney et al. 2016, Table 2). The results are shown in Figure 1.

Per the discussion in §3.3, regions with $\pi_9 > 8.25''$ are marked as inaccessible to our search in Figure 1. The $d_9 < 250$ AU region of parameter space is also marked as inaccessible for multiple reasons. Planet Nine could only be so nearby if very close to perihelion, which is unlikely (Brown & Batygin 2016). Additionally, the total apparent motion over 5.5 years can be larger for smaller d_9 , which requires fitting orbits to very large numbers of transient tuples. As a result, purely by chance, we expect to become disproportionately flooded with spurious linkages as we push toward lower d_9 . Because the marginal benefit of pushing to $d_9 < 250$ AU is outweighed by the increase in false positives, we perform our search only for $d_9 \geq 250$ AU.

In summary, we expect substantial sensitivity to the maximally W1-luminous Fortney et al. (2016) model, especially at relatively low $|\beta|$. For instance, at $|\beta| \leq 30^\circ$ the $H_{W1}=2.13$ model is detectable at all distances between 250 and ~ 700 AU. We also forecast sensitivity out

to $d_9 \approx 800$ AU at high $|\beta|$. We note that during orbit linking (§7), no upper limit is imposed on d_9 . We are therefore able to detect Planet Nine out to thousands of AU, provided it could be sufficiently more luminous than any of the Fortney et al. (2016) models. Our actual sensitivity turns out to be ~ 0.2 mag better than the present forecasts, essentially because of the distinction between the 5σ detection threshold considered in this section and our 90% completeness threshold for acquiring a quadruplet of detections (§8).

Alternative search strategies based on linking single-exposure detections (for instance a W1 variant of Luhman 2014) are not expected to be sensitive beyond $d_9=430$ AU, given the single-exposure limiting magnitude of $W1=15.3$ and adopting the maximum W1 luminosity from Fortney et al. (2016). Still, a search based on single-exposure transients would be complementary to our coadd-based analysis, filling in some of the low d_9 parameter space which is problematic for our approach.

3.5. Drift within a Coadd Epoch

We have thus far ignored the complication of orbital drift within the timespan of a coadd epoch. If large enough, orbital drift could compromise our sensitivity by further smearing the appearance of Planet Nine in our coadds. To determine the maximum intra-coadd orbital drift, we calculate the maximum intra-coadd drift within the mask shown in Figure 1, for each (a_9, e_9) pair in a set of orbits tracing the two loci of a_9 versus e_9 provided in §2 of Brown & Batygin (2016). We find that the maximum intra-coadd drift across all orbits is $1.29''$, which occurs near $|\beta|=30^\circ$, for orbits currently very close to perihelion and with perihelia very similar to the inner edge of our search space ($q \approx d_9 \approx 250$ AU).

Thus, the intra-coadd orbital drift is always expected to be less than $1/4$ of the W1 PSF FWHM. Because of the small amplitude of intra-coadd drift, and because it does not in general add constructively with parallactic motion, we do not attempt to account in detail for its minimal effect on our sensitivity.

3.6. Drift Between Coadd Epochs

We note that the upper limit of d_9 for which our §8 completeness curve calculation applies is dictated by *lack* of sufficient orbital drift for very distant, eccentric objects near aphelion. Because a given coadd epoch’s reference image (§5.3) will in general be constructed using other coadd epochs with zero relative parallax, an orbiting object drifting sufficiently slowly will behave somewhat like a static source, and at least partially cancel itself out in our difference images. Assuming $e \leq 0.75$, the annual drift⁹ does not become smaller than the W1 FWHM until $d_9 \gtrsim 2250$ AU. Our pixel-level simulations suggest a negligible loss in sensitivity even in the case of $6.1''/\text{yr}$ drift, and indeed the linear drift model of §7.3 recovers faint high proper motion stars among our list of transients down to $\mu \approx 0.2''/\text{yr}$. To be conservative, we claim our completeness curve of §8 to be applicable only for $d_9 \leq 2250$ AU, though we have considerable sensitivity at larger distances in the very unlikely event

⁹ Annual drift is the relevant quantity because W1 coadd epochs with zero relative parallax are spaced almost exactly an integer number of years apart.




FIG. 2.— Sky region searched in equatorial, ecliptic and Galactic coordinate systems. Black lines trace edges of the 805 unWISE tile footprints which are analyzed throughout this work. Green and red contours are those of Holman & Payne (2016), Figure 9. These 805 unWISE tiles were chosen to fill the outer red contour, plus some margin. The area searched totals 1840 square degrees.

that Planet Nine could be so far away and detectably luminous.

4. INPUT DATA DETAILS

The WISE “Level 1b” (L1b) single-exposure images constitute the primary input data for our search. For every L1b frameset in the search region of interest, we employ the `-w1-int-1b`, `-w1-msk-1b` and `-w1-unc-1b` images, which respectively provide the sky intensity and corresponding per-pixel bitmask and uncertainty values. We have downloaded a copy of these files for every publicly available frameset in our search region (shown in Figure 2), including those from the primary WISE mission and the first two NEOWISER releases. In total, $\sim 315,000$ W1 L1b framesets contribute to our analysis. The dates of observation of the WISE exposures analyzed range from 2010 January 7 to 2015 August 22 (UTC).

We also make use of the AllWISE Source Catalog and 2MASS All-sky Release source catalog (Skrutskie et al. 2006), to aid in masking static compact sources (§6.3).

5. IMAGE COADDITION METHODOLOGY

To stack the W1 single exposures, we make use of the Lang (2014) unWISE coaddition framework. We employ the Meisner et al. (2016) adaptation of the original unWISE codebase. The Meisner et al. (2016) unWISE coaddition pipeline includes updates to handle NEOWISER imaging and better correct time-dependent artifacts. We briefly review a few notable aspects of the unWISE coaddition procedure here; refer to Lang (2014) and Meisner et al. (2016) for a full discussion.

unWISE coaddition divides the sky into the same set of $18,240$ $1.56^\circ \times 1.56^\circ$ astrometric footprints defined by the WISE team’s Atlas stacks. The tile centers trace out a series of iso-declination rings, and their axes are aligned along the equatorial cardinal directions. While the Atlas coadds are intentionally convolved by the WISE PSF, our unWISE code preserves the native WISE angular resolution using Lanczos interpolation, resulting in outputs with $2.75''$ pixels and 2048 pixels on a side. We note that these footprints are not mutually exclusive. Neighboring tiles typically overlap each other by $\sim 3'$ along each boundary.

In the Lang (2014) and Meisner et al. (2016) coadds, there is no notion of a coadded epoch like the one we introduced in §3.1. In these previous works, every W1 coadd is simply identified by its `coadd_id` value, which

is a string encoding the tile center RA and Dec (e.g. ‘0000p000’). For our epochal coadds, we define a zero-indexed integer called `epoch` such that a W1 epochal coadd is uniquely defined by its (`coadd_id`, `epoch`) pair. For each `coadd_id`, `epoch` increases with time. However, a particular `epoch` value will not in general correspond to a single season for all `coadd_id` values; `epoch=1` of a given `coadd_id` may occur in a different season than `epoch=1` of another `coadd_id`.

To create the time-resolved coadds analyzed in this work, as opposed to the full-depth stacks of Meisner et al. (2016), we have made the following additional updates to the unWISE pipeline:

1. We segment the exposures which overlap each `coadd_id` footprint by MJD, into a series of ~ 1 day coadd epochs. The boundaries of each coadd epoch coincide by default with the beginning and end of a WISE visit. However, a visit affected by a Moon avoidance maneuver is split into two when enforcing the rules stipulated in §5.2.
2. Each quadrant of *every* L1b exposure receives a fourth-order polynomial background correction as described in §6 of Meisner et al. (2016) before being accumulated into its epochal coadd. The Meisner et al. (2016) full-depth W1 stacks are used as templates with respect to which the polynomials are computed. These corrections help to avoid sharp edge-like features in our coadds near the boundaries of contributing L1b exposures, which could potentially lead to large numbers of spurious difference image detections.
3. After creating an initial version of each epochal coadd, we generate a final version by ignoring any exposures which were flagged as having an excessive fraction of outlier pixels ($>1\%$) during creation of the initial coadd. Such exposures are rejected midway through the first iteration of coaddition, but their outliers compromise our flagging of outlier pixels in *other* exposures. By performing a second coadd iteration that completely ignores problematic exposures from the start, we are able to greatly reduce the number of spurious transients due to cosmic rays, particularly in sky regions affected by the South Atlantic Anomaly.




FIG. 3.— Map of the number of coadd epochs per `coadd_id` footprint throughout our search region.

4. We generate two additional metadata images per epochal coadd. First, we create a map of the mean MJD of the contributing exposures at each pixel. This data product is crucially important, because its values at the positions of detected transients are used during Keplerian orbit linking (§7). Second, we make a map of the time spanned by the exposures contributing to each coadd pixel. We use these maps to ensure that the coadd pixel timespans never significantly exceed the expectation from Equation 1.

The search region shown in Figure 2 consists of 805 distinct `coadd_id` astrometric footprints. In total, we generated 6,219 epochal coadds. Each `coadd_id` footprint has between 6 and 10 epochs available. The number of coadd epochs per `coadd_id` footprint throughout our search region is shown in Figure 3.

Figure 4 shows an example of the benefits of our coadds relative to single exposures. The coadds are ~ 1.2 magnitudes deeper than L1b images, and the numerous main-belt asteroids and satellites found in single exposures are completely removed by the unWISE outlier rejection.

5.1. Outlier Rejection Tests

One might reasonably ask whether the same outlier rejection that removes satellite streaks and main-belt asteroids could also erase Planet Nine from our epochal coadds. We have performed injection tests to address this concern. We consider the worst-case (fastest-moving) scenario by injecting point sources with parallax of $8.25''/\text{day}$ into individual exposures, and then running the modified exposures through our coaddition process. We inject 100 such fakes per $W1$ magnitude, for dozens of $W1$ magnitudes spanning $14.0 \leq W1 \leq 16.8$. The outlier rejection should be most problematic for the brightest Planet Nine injections. This motivated our choice to extend these fakes as bright as $W1=14.0$, since the maximally $W1$ -luminous Fortney et al. (2016) model would




FIG. 4.— Small $9.0' \times 6.7'$ field extracted from epochal coadd with $(\text{coadd_id}, \text{epoch}) = (0211\text{p}000, 1)$ and one of its contributing exposures, illustrating the benefits of our coaddition process. Top: contributing $W1$ exposure 06337b163. Bottom: time-resolved coadd resampled onto the 06337b163 astrometry. Both panels have identical grayscale stretches. The coadd is visibly much deeper and less noisy, as expected. In addition, a satellite streak which would likely lead to many spurious transient detections in the L1b image has been completely eliminated. The coaddition is also highly effective at rejecting fast-moving main-belt asteroids – the red circle marks the location of (6136) Gryphon during exposure 06337b163.

have $W1 \approx 14.2$ at the inner edge of our search space ($d_9 = 250$ AU).

We find that the outlier rejection causes a very small, but nonetheless measurable, reduction in the average fluxes we recover relative to the true fluxes of the injections. Near our 90% completeness threshold of $W1=16.66$ (§8), on average 7 mmag of flux is lost. The fraction of flux lost ramps up slowly for brighter injections. At the bright end, $W1=14.0$, 19 mmag of flux is lost on average.

Because the impact of unWISE outlier rejection on a maximally fast-moving Planet Nine near our 90% completeness threshold is less than a hundredth of a magnitude, we neglect this effect in our completeness analysis. In short, Planet Nine is expected to be too slow-moving and faint to trigger unWISE outlier rejection at a level which significantly impacts our sensitivity.

5.2. Moon Avoidance Maneuvers

The largest departures of WISE’s pointing relative to solar elongation of 90° occur during so-called “Moon-

TABLE 2
MOON AVOIDANCE MANEUVERS

Season	MJD _{start}	MJD _{end}	λ_{approx}
a	55219.5	55220.075	38.0°
b	55381.28	55381.88	12.0°
b	55410.91	55411.53	41.0°
c	55573.63	55574.18	26.5°
d	56665.1	56665.62	20.0°
d	56694.815	56695.27	50.0°
e	56856.71	56857.16	23.0°
e	56886.10	56886.55	51.0°
f	57019.74	57020.21	9.5°
f	57049.49	57049.97	43.0°
g	57211.67	57212.08	16.0°
g	57240.87	57241.335	43.5°

avoidance maneuvers”. These maneuvers coincide with the roughly monthly occasions on which the Moon crosses the WISE 90° elongation sight line. At the start of such a maneuver, WISE will gradually shift several degrees ahead in ecliptic longitude (λ) relative to 90° elongation, then avoid the Moon by jumping ~ 5 -10° backward in λ . This backward jump in ecliptic longitude has the effect of lengthening WISE visits near in time to a Moon avoidance maneuver to last up to ~ 5 -7 days. Coadds spanning such long periods of time would dramatically dilute the signal from an object orbiting at hundreds of AU, which could undergo many dozens of arcseconds of parallax during such a timespan.

In order to avoid long coadd timespans that are strongly inconsistent with Equation 1, we enforce the following rules:

1. No epochal coadd may include exposures during a Moon avoidance maneuver. A list of start/end MJD values for the maneuvers relevant to our search region is provided in Table 2. By discarding the frames acquired during Moon avoidance maneuvers, we lose 1.6% of exposures.
2. No epochal coadd may include exposures acquired both before and after the same Moon avoidance maneuver.

This time-slicing splits the 5-7 day visits created by Moon avoidance maneuvers into two short coadd epochs which span lengths of time consistent with or shorter than Equation 1. As a result, even though there are only 7 seasons of W1 observations available, we end up with an average of 7.7 coadd epochs per `coadd_id`, thanks to the “extra” epochs created by Moon avoidance maneuvers. Once the above rules are enforced, the MJD ranges spanned by exposures contributing to actual pixels in our epochal coadds closely follow Equation 1. The largest measured value of τ_{coadd} for any pixel in our analysis is 1.65 days, which occurs at $\beta = -52.3^\circ$ (the expectation from Equation 1 at this $|\beta|$ is 1.64 days). The mean measured value of τ_{coadd} across all pixels in our analysis is 1.01 days.

Note that WISE Moon avoidance maneuvers are spaced such that a single `coadd_id` can have at most two coadd epochs during a single season. For cases in which one `coadd_id` has two coadd epochs in one season, these epochs are separated by a small handful of days.

5.3. Creating Reference Stacks

For each time-resolved coadd identified by its unique (`coadd_id`, `epoch`) pair, we generate a corresponding reference stack. To do so, we combine all epochal coadds for the relevant `coadd_id` astrometric footprint, excluding the epoch under consideration. We also exclude any other epoch during the same season as the epoch for which the reference is being created. This latter constraint is important because, in cases where a single `coadd_id` has two epochs during the same season, these epochs are typically separated by only a couple of days, and Planet Nine’s apparent motion could be slow enough that its detections at these two epochs have some overlap. To combine the epochal coadds into a reference image, we apply a weighted median filter on a per-pixel basis. The mean integer coverage of our reference stacks is 71.5 exposures, $\sim 6\times$ the typical integer coverage of our epochal coadds.

6. DIFFERENCE IMAGING PIPELINE

We use a suite of standard image differencing and source extraction software to obtain a catalog of transient detections in each epochal coadd.

6.1. Generating Difference Images

We begin by astrometrically calibrating the reference coadd to 2MASS using SCAMP (Bertin 2006). We next use SCAMP to astrometrically calibrate the epochal coadd to 2MASS by comparing its source positions against the 2MASS-calibrated reference stack source positions. We then use SWARP (Bertin et al. 2002) to warp the reference and its uncertainty image to match the astrometry of the epochal coadd.

We use the warped reference to create two difference images for each epochal coadd. First, we perform a subtraction using the HOTPANTS package¹⁰, which applies the Alard (2000) kernel-matching technique. Second, we perform a direct subtraction without kernel matching. We find that for WISE, which delivers highly stable space-based imaging, kernel matching does not play the same crucial role it does for ground based images, where the PSF shape/size can vary dramatically. When creating the direct subtraction, we do take into account small zero point differences between the epochal and remapped reference images, using the `KSUM00` value obtained by HOTPANTS.

6.2. Extracting Transient Sources

To extract transient sources from the subtracted images, we use Source Extractor (Bertin & Arnouts 1996) in “dual image mode”, which makes use of static source properties in the remapped reference image to inform source detection in the subtracted image.

6.3. Filtering Difference Image Extractions

In addition to legitimate transients, the source extraction performed on our difference images also yields large numbers of false positives. We need to cull the list of transients for several reasons. First, the number of quadruplets that must be fit with Keplerian orbits during linking scales as the fourth power of transient number

¹⁰ www.astro.washington.edu/users/becker/v2.0/hotpants.html

density, meaning that our full list of difference detections would be computationally expensive to link. Aside from computational limitations, we also expect the number of chance occasions on which spurious transients happen to link into Keplerian quadruplets to scale as the fourth power of the transient number density. Thus, eliminating as many spurious transients as possible from the outset allows us to obtain a manageable number of candidate quadruplets well-fit by Keplerian orbits.

To remove false positives, we would ideally employ an approach like that of Goldstein et al. (2015), analyzing a variety of pixel-level features for each transient in order to classify it as legitimate versus spurious. Unfortunately, that methodology has thus far only been developed for the case of pointlike transients, whereas we are explicitly searching for a moving object which may appear slightly streaked. We therefore resort to a number of heuristics. These heuristics are designed to have minimal impact on our completeness, and have effects which can be easily incorporated into our simulations of §8. The following subsections detail each of the transient filtering heuristics.

6.3.1. 2MASS Sources

We discard any transient within a 1 pixel radius ($2.75''$) of a 2MASS source location. This cut eliminates 0.16% of the search area.

6.3.2. Bright Static W1 Sources

We create a mask for each bright W1 compact source in the AllWISE catalog with $w1mpro < 9.5$, and discard all transients that land within these masks. The mask for each bright source is built by determining which pixels in the coadd will be brighter than a threshold value of 100 Vega nanomaggies¹¹, given the source’s total flux, centroid and the Meisner & Finkbeiner (2014) W1 PSF model. These masks include considerable detail, as the $14.9' \times 14.9'$ Meisner & Finkbeiner (2014) PSF model extends quite far into the wings. For instance, our bright star masks capture significant portions of the diffraction spikes. These masks eliminate 0.46% of the search area.

6.3.3. Kernel Matching Artifacts

We find that the HOTPANTS kernel matching produces “ringing” features around moderately bright sources, and that these artifacts are frequently identified as difference image detections. These are difficult to mask without discarding an excessive amount of area because they often extend ~ 10 - $15''$ from their parent source’s centroid. We find that when running source detection on difference images created via direct subtraction, virtually all legitimate transients are still recovered, while the spurious detections associated with moderately bright stars are strongly confined near their centroids. Therefore, we require that transients be detected in both the HOTPANTS and direct subtractions (using a $5.5''$ match radius). This requirement does reduce the overall depth of our search, as the noise properties of the difference images with and without kernel matching can be different. Using the simulations described in §8, we

¹¹ A nanomaggie is a linear unit of flux such that a 1 nanomaggie source corresponds to a magnitude of 22.5.




FIG. 5.— Histogram of W1 magnitudes for detections in our filtered transient catalog. The source counts peak near $W1=16.6$.

find the size of this effect to be 0.048 mags. On the other hand, this requirement leads to a $\sim 2.5\times$ reduction in the number of transients, leading us to deem the depth trade-off acceptable. Note that, although we create two separate difference detection catalogs per epochal coadd, our downstream analyses always employ transient properties (in particular RA and Dec) measured from HOTPANTS subtractions with kernel matching.

6.3.4. Reference Image Brightness

We reject any transient with *total* flux less than the reference coadd brightness at its centroid. Given the W1 PSF and our coadd pixel size, this is roughly equivalent to rejecting any source with peak brightness fainter than the reference image at the same location by a factor of more than 7. For the transient flux, we adopt the Source Extractor `flux_best` measurement. It is not entirely straightforward to translate our reference brightness cut into a fraction of area masked, but we note that the transient magnitude histogram (Figure 5) is sharply peaked near $W1=16.6$, and for this magnitude 0.48% of the total search area is masked.

6.3.5. Negative Fluxes

We discard transient detections with negative `flux_best`, as these are generally spurious extractions near bright source residuals.

6.3.6. Ultra-Bright Stars

A small handful of stars are so exceptionally bright that the approach of §6.3.2 is insufficient, and such stars result in localized regions with extremely high transient number density. In these cases, we construct special masks consisting of a circular core centered on the star, plus straight lines emerging from the centroid at angles of approximately 45° , 135° , 225° and 315° from ecliptic north, representing diffraction spikes. Transients within these masks are discarded. The four sources masked in this way are Mira, τ^4 Eri, α Ceti and V^* EG Cet. 0.07% of the search area is lost.

6.3.7. Latents

The W1 detector suffers from persistence artifacts dubbed ‘latents’, which appear as faint, diffuse positive excursions above background. A latent will occur at the *detector* position of a sufficiently bright ‘parent’ source, in the exposure immediately following imaging of the parent. Extremely bright stars can also yield ‘second latents’, artifacts similar to standard latents, but occurring *two* exposures after the parent was imaged. Because the WISE scan direction varies between coadd epochs, latents do not always appear at the same sky location, and thus are expected to contaminate our list of difference image extractions.

Using the AllWISE catalog and each epochal coadd’s unWISE `-frames` metadata table¹², we can predict the exact positions of all latents in each coadd. We adopt a latent parent brightness threshold of $w1pro \leq 8.3$, and a second latent threshold of $w1pro \leq 6.2$. We discard any transient within a 1 pixel radius of a predicted first or second latent position. $\sim 75,000$ persistence artifacts are thereby removed from our transient list, while masking only 0.03% of the search area.

W1 latents are much more diffuse than the most streaked Planet Nine profile for which we claim to be sensitive, with $n_{eff} > 10\times$ that of a point source. Our extraction of tens of thousands of latents from the difference images demonstrates our ability to detect even highly non-pointlike transients.

6.3.8. Merging Per-coadd Transient Lists

All of the previous filtering steps are applied on a per-coadd basis. However, there will be some transients near, but still within, the edges of multiple `coadd_id` footprints, since these are not mutually exclusive. Because the minimum Moon avoidance maneuver timespan is 0.4 days and WISE visits are spaced by ~ 6 months, transients at the same location but separated temporally by < 0.4 days, even if extracted from different coadd images, will generally be found in the same set of L1b exposures and are therefore not distinct. These multiple detections should be merged prior to orbit linking, to avoid the possibility of a transient being linked with another copy of itself. We create a single merged transient catalog from the per-coadd catalogs by implementing a ‘resolve’ step. The resolve step looks for groups of transients separated spatially by $\leq 1''$ and temporally by < 0.4 days, retaining only the single transient furthest from the boundary of its coadd tile.

6.3.9. Filtering Summary

After filtering, the final transient catalog contains 206,891 detections, corresponding to a mean density of 113 per square degree. Figure 5 shows a histogram of the transient W1 magnitudes, which peaks near $W1=16.6$. The number of transients per $N_{side}=128$ HEALPix pixel (0.21 square degrees, Górski et al. 2005) is shown in Figure 6.

The regions masked due to our various filters are not mutually exclusive, and in total we have masked 1.0% of the search area. This small loss of area is accounted for in the Monte Carlo analysis of §8.




FIG. 6.— Number of transients in the filtered catalog per 0.21 square degree $N_{side}=128$ HEALPix pixel. Red dots mark the locations of transients linked into low- χ^2 Keplerian quadruplets (§7.3).

7. ORBIT LINKING

Our basic orbit linking approach is to enumerate all viable quadruplets of transients, and determine which can be well fit by a Keplerian orbit. The following subsections provide additional details of our orbit linking methodology, which is largely inspired by that of Brown et al. (2015).

7.1. Astrometry

The goal of our astrometric calibration is to ensure that the most appropriate and accurate spatial coordinates are input to our orbit linker (§7). The orbit linker computes model orbits in ICRS coordinates, and therefore, as described in §6 we have calibrated the astrometry of each epochal coadd to ICRS via 2MASS.

It is crucial to have a full accounting of the uncertainties on our astrometric measurements of transients, since these uncertainties will strongly influence the χ^2 values of our best-fit orbits. We include three terms in our astrometric uncertainty model.

7.1.1. 2MASS Calibration Residuals

We employ the SCAMP outputs `ASTRRMS1`, `ASTRRMS2` as per-coadd measures of the residual astrometric scatter between our calibrated epochal coadds and 2MASS, for relatively bright sources. The typical value of these parameters is $0.28''$. We aim to derive a *per-coordinate* astrometric uncertainty, but only one value per transient, rather than two distinct values for the RA and Dec directions. Therefore, we define $\sigma_{calib} = (\text{ASTRRMS1} + \text{ASTRRMS2})/2$. We calculate one such σ_{calib} value per epochal coadd, which is assigned to all transients within that coadd.

7.1.2. Systematic Proper Motion Shifts

We have not corrected for proper motion when calibrating WISE to 2MASS. As discussed in Cutri et al. (2013), there is a systematic error associated with this

¹² <http://unwise.me/data/-README.txt>

procedure, because calibrator sources in a given sky region may move coherently, with nonzero average proper motion¹³. From Cutri et al. (2013), we adopt a maximum systematic error (per coordinate) of 200 mas between 2000.0 and 2010.0, and scale this linearly based on the MJD of each transient. We call this quantity σ_{pm} , and its mean value is 0.26".

7.1.3. Statistical Scatter

Statistical scatter due to background noise is by far the dominant source of astrometric uncertainty for our typical transient. We use the 300,000 injections described in §8 to fit an analytic relation to the (point source) per-coordinate statistical scatter as a function of W1 magnitude and integer coverage n_{exp} . We find that the results are well-fit by:

$$\sigma_{stat,PSF} = 0.41'' \times \frac{\sqrt{12/n_{exp}}}{(10^{(16.0-W1)/2.5})^{0.69}} \quad (9)$$

To be conservative, we wish to scale this uncertainty up so that it is appropriate even in the case of maximum intra-coadd parallax (8.25"). Typically, for fixed amount of signal, the astrometric uncertainty scales linearly with both FWHM and the effective amount of noise. The FWHM (n_{eff}) for 8.25" of parallax is $1.30 \times (1.48 \times)$ larger than that of a W1 point source, and therefore we adopt $\sigma_{stat} = 1.30 \times \sqrt{1.48} \times \sigma_{stat,PSF}$. Finally, if the positional error ellipse semi-major axis quoted by Source Extractor is larger than our model for σ_{stat} , we assign σ_{stat} to the Source Extractor value.

7.1.4. Total Astrometric Uncertainty

For our per-coordinate uncertainty provided to the orbit linker (σ_{ast}), we adopt the quadrature sum of σ_{calib} , σ_{pm} and σ_{stat} . The median σ_{ast} value is 0.90". The σ_{calib} and σ_{pm} terms effectively provide a floor at $\sigma_{ast} \approx 0.4''$, even for very bright transients.

7.2. Transient Linking Rules

When enumerating quadruplets, we do not simply allow all transients in the merged catalog to be paired with all other transients. We enforce two rules restricting which transients may be linked:

1. No two transients may be linked if they are extracted from the same (coadd_id, epoch) coadd image.
2. No two transients may be linked if they are separated temporally by <0.4 days. Because of the way that our coadd epochs are determined from WISE visit and Moon avoidance maneuver boundaries, coadd pixels with mean MJD separated by less than 0.4 days will in general have been constructed from (at least partially) the same set of single-exposure images.

Otherwise all transients may be linked. Importantly, we allow linkages of transients which occur during the same season but are extracted from different coadd images, provided the second rule above is not violated. Note also that we impose no constraints on the relative fluxes of transients that may be linked.

¹³ wise2.ipac.caltech.edu/docs/release/allwise/expsup/sec2_5.html

7.3. Fitting All Quadruplets

Our basic orbit linking approach will be to enumerate all viable quadruplets of transients, and determine which can be well fit by a Keplerian orbit. We break this process up into two phases, a ‘‘prescreen’’ of a large number of possible candidates, followed by a detailed fit to Keplerian orbital parameters. The prescreen step starts with the complete list of $n_{transient}$ transients, for each one identifying a region of the sky where that object might be at subsequent epochs. The actual shape of this window is set by our choice of a minimum orbital distance, which we take to be 250 AU, and limits on the inclination. Our results of §9 place no restrictions on inclination. In this way we create $n_{transient}$ lists of neighboring transients, with typically 40 neighbors each.

The next phase of our prescreen is to run through the neighbor lists, combining them to form unique quadruplets. We fit the observation times and sky positions of each quadruplet with a simple linear model that includes the magnitude of the parallax and the 2-D plane-of-sky orbital drift; this model is appropriate to distant Solar System bodies observed on time scales that are short compared with their orbital periods. If this model yields a $\chi^2 < 80$, we accept the quadruplet; otherwise we reject it.

The second phase of our linking algorithm is to fit each prescreened quadruplet with Keplerian orbital parameters. For this step we use the Bernstein & Khushalani (2000) code, `orbfit`, designed for objects at Kuiper belt distances and beyond. If `orbfit` successfully resolves a Keplerian orbit (see §7.4), we save the quadruplet. Finally we take the list of saved quadruplets and attempt to combine them into quintuplets, sextuplets, etc., if they have transients in common. In this way we start with $\sim 2 \times 10^5$ W1 transients, assess over 14 billion quadruplets, and perform Keplerian orbit fits to the $\sim 12,000$ that pass prescreening. The final list contains just under 500 successful quadruplet linkages. We always retain every tuple with low `orbfit` χ^2 , never making cuts on the best-fit orbital elements or relative fluxes of constituent transients.

7.4. Keplerian Orbit χ^2 Threshold

We must choose a goodness-of-fit threshold defining the boundary between quadruplets that are well fit versus poorly fit by a Keplerian orbit. We do so by choosing a χ^2 threshold such that the expected false negative rate is 0.001. Our `orbfit` quadruplet (quintuplet) fits always have 2 or 3 (4 or 5) degrees of freedom (DOF). To obtain the desired level of false negatives, we therefore choose our χ^2 threshold to be 13.8, 16.3, 18.5, 20.5 for 2, 3, 4, 5 DOF. The impact of our 0.001 failure rate in fitting orbits can be accounted for in combination with the simulations of §8, but it leads to a negligible shift in our 90% completeness threshold.

8. COMPLETENESS ANALYSIS

8.1. Notation

There are two types of completeness relevant to our search. The first completeness refers to the probability that a transient source of a given brightness is successfully detected in a reference-subtracted epochal coadd. This will be denoted f_{det} . Note that f_{det} will be a func-

tion of coadd integer coverage n_{exp} . The second completeness is the probability with which an object of a particular brightness is detected as a transient in ≥ 4 epochal coadd difference images, given the available number of coadd epochs and the number of contributing exposures within each coadd epoch. Calculating $f_{\geq 4}$ is our ultimate goal, because we require that a Planet Nine candidate consist of ≥ 4 transients which are linked along a low- χ^2 orbit. Calculating f_{det} as a function of integer coverage is one step towards determining $f_{\geq 4}$.

8.2. Parameterizing Completeness Curves

f_{det} is expected to approach unity for very bright transients, and should decrease to zero for very faint transients, which will be indistinguishable from background noise. The situation is similar for $f_{\geq 4}$. The rejection of transients due to e.g. our bright star masking means that $f_{\geq 4}$ will asymptote to a value slightly less than unity for bright objects.

For these reasons, we parameterize completeness curves, in both the case of f_{det} and $f_{\geq 4}$, with the model:

$$f(W1) = \frac{A}{2} \times [1 + \text{erf}(-(W1 - \mu)/\sigma)] \quad (10)$$

Where $W1$ is the $W1$ magnitude, and there are 3 free parameters: A , μ and σ . A is the asymptotic completeness value toward low $W1$, and is typically very close to unity. When fitting (A, μ, σ) we impose the constraint $A \leq 1$. μ is the $W1$ value at which the completeness is $A/2 \approx 0.5$, the center of the transition from $f \approx 1$ to $f = 0$. σ dictates the width of this transition.

8.3. Computing f_{det} Curves

In order to most accurately calculate $f_{\geq 4}$, it is necessary to compile a set of completeness curves $f_{det}(W1; n_{exp})$, for each value of n_{exp} . We have done so by injecting 300,000 pointlike fakes of varying $W1$ into our difference images. 100 real difference images spread throughout our search footprint are used as the basis for 3,000 distinct mock difference images, each with 100 fakes injected. We record whether each fake is recovered using a 1 pixel match radius. We then bin the fakes by n_{exp} , and fit Equation 10 to the fraction of fakes recovered as a function of $W1$ for each n_{exp} value.

Not all n_{exp} values are sufficiently sampled by these fakes, because many n_{exp} values represent just a tiny fraction of our search area. Most of the fakes have $6 \leq n_{exp} \leq 15$, and these are the n_{exp} values for which high-quality f_{det} curves can be created. Table 3 lists the best-fit completeness curve parameters (A , σ , μ) for these n_{exp} values.

One might hope that $\mu(n_{exp})$ scales as expected from statistical noise proportional to $1/\sqrt{n_{exp}}$:

$$\mu_{pred}(n_{exp}) = \mu(n_{exp,0}) + 2.5 \log_{10}(\sqrt{n_{exp}/n_{exp,0}}) \quad (11)$$

We take $n_{exp,0}=12$, given that 12 is both the median and mode of n_{exp} in our full set of epochal coadds. $\mu_{pred}(n_{exp})$ calculated according to Equation 11 is also listed in Table 3, and it clearly agrees very well with the measured $\mu(n_{exp})$ values for $6 \leq n_{exp} \leq 15$.

For our downstream Monte Carlo analysis, we employ $f_{det}(W1; n_{exp})$ completeness curves parameterized

TABLE 3
PARAMETERS TABULATED AS A FUNCTION OF n_{exp} , THE NUMBER OF EXPOSURES CONTRIBUTING TO A COADD PIXEL

n_{exp}	$p(n_{exp})$	μ	σ	A	μ_{pred}
3	0.0104	-	-	-	16.156
4	0.0108	-	-	-	16.312
5	0.0117	-	-	-	16.433
6	0.0139	16.512	0.302	0.9970	16.532
7	0.0197	16.624	0.356	0.9914	16.616
8	0.0338	16.689	0.368	0.9940	16.688
9	0.0657	16.766	0.347	0.9955	16.752
10	0.118	16.832	0.342	0.9940	16.809
11	0.175	16.867	0.352	0.9933	16.861
12	0.196	16.908	0.354	0.9945	16.908
13	0.162	16.948	0.355	0.9943	16.952
14	0.0986	16.997	0.360	0.9943	16.992
15	0.0450	17.044	0.368	0.9958	17.030
16	0.0193	-	-	-	17.065
17	9.81×10^{-3}	-	-	-	17.098
18	4.70×10^{-3}	-	-	-	17.129
19	2.29×10^{-3}	-	-	-	17.158
20	1.27×10^{-3}	-	-	-	17.186
21	7.12×10^{-4}	-	-	-	17.212
22	5.42×10^{-4}	-	-	-	17.238

by Equation 10. For these parameterizations, we adopt $\mu = \mu_{pred}(n_{exp})$. Since σ and A appear roughly constant over the range of n_{exp} for which these have been measured, we adopt their median values ($A=0.9943$, $\sigma=0.355$) for all $f_{det}(W1; n_{exp})$ completeness curves during Monte Carlo simulations.

8.4. Coverage Statistics

Our Monte Carlo simulation of $f_{\geq 4}$ requires as input statistics regarding the fraction of sky area in the full search footprint as a function of number of coadd epochs (n_{epochs}) and number of exposures per epochal coadd pixel (n_{exp}). Tables 4 and 3 list these values, referred to as $p(n_{epochs})$ and $p(n_{exp})$ respectively. Note that our analysis discards the small number of coadd pixels with integer coverage ≤ 2 , as it is not possible to perform outlier rejection with fewer than three exposures.

TABLE 4
FRACTION OF PIXELS IN SEARCH REGION AS A FUNCTION OF NUMBER OF COADD EPOCHS

n_{epochs}	$p(n_{epochs})$	n_{epochs}	$p(n_{epochs})$
1	1.32×10^{-5}	6	0.126
2	1.20×10^{-5}	7	0.483
3	1.41×10^{-5}	8	0.327
4	1.50×10^{-5}	9	0.0635
5	2.50×10^{-4}		

8.5. $f_{\geq 4}$ Monte Carlo Calculation

We generate a large number of fake objects, each assigned its own magnitude $W1$. Each object is randomly assigned a number of coadded epochs using a multinomial draw from $p(n_{epochs})$. Each epoch of each object is then randomly assigned a corresponding number of exposures n_{exp} using a multinomial draw from $p(n_{exp})$. Given an object's assigned $W1$ and n_{exp} within each epoch, each appearance of that object will be detected with a probability given by $f_{det}(W1; n_{exp})$, which is specified by the parameters (μ , σ , A) described in the final paragraph of




FIG. 7.— Completeness of obtaining ≥ 4 detections as a function of W1 magnitude. Black plus marks are results of the Monte Carlo simulation described in §8.5. Red is the best-fit erf model, with parameters $(\mu, \sigma, A) = (16.83, 0.19, 0.998)$. The completeness curve shown incorporates corrections for area lost to spatial masking and non-pointlike morphology.

§8.3. Each epoch of each object is then marked as either ‘detected’ or ‘undetected’ by performing a Bernoulli trial with success probability given by $f_{det}(W1; n_{exp})$. $f_{\geq 4}(W1)$ is then calculated by computing the fraction of objects for which ≥ 4 total detections are accumulated, using a sample of 10,000 objects per value of W1, with $W1=16.20, 16.21, \dots, 17.59, 17.60$.

Our measured $f_{\geq 4}$ completeness curve is plotted in Figure 7. Fitting the tabulated $f_{\geq 4}(W1)$ values with Equation 10 gives $(\mu, \sigma, A) = (16.89, 0.18, 0.998)$. This best-fit model reaches $f_{\geq 4}=0.90$ (‘90% completeness’) at $W1=16.72$. Requiring five rather than four detections results in a 0.11 magnitude reduction in sensitivity, with 90% completeness at $W1=16.61$. We can account for the masking of 1.0% of our search area (§6.3) by randomly zeroing out 1% of $f_{det}(W1; n_{exp})$ probabilities during the Monte Carlo. We find that this pushes the 90% completeness threshold brighter by 6.4 (9.5) mmag in the case of demanding four (five) linked detections.

We can scale these measured point source completeness thresholds in a manner consistent with the final term of Equation 6, to account for variations in sensitivity as a function of $|\beta|$ and d_9 . The resulting 90% completeness threshold is shown in Figure 8. Over the parameter space in which we are sensitive to the maximally W1-luminous Fortney et al. (2016) model, the mean 90% completeness threshold is 16.66 (16.54) when demanding at least four (five) linked detections. The ~ 0.05 mag degradation in sensitivity relative to our point source measurement is a result of parallactic smearing, as the average π_9 value within our search space is $3.6''$.

We caution that we might fail to link detections of Planet Nine if it saturates in W1 ($W1 \lesssim 8.25$), though this would require a luminosity hundreds of times greater than any of the Fortney et al. (2016) models. We might also fail to detect Planet Nine if it is brighter than $W1=9.5$, due to the AllWISE catalog masking described in §6.3.2. Finally, we have not presented a detailed accounting of small-scale variations in our 90% completeness threshold, such as those introduced by bright star masking and the spatially varying number of coadd epochs shown in Figure 3.




FIG. 8.— Variation of our 90% completeness threshold for acquiring a quadruplet of detections, due to changes in amount of parallactic smearing and depth of coverage with d_9 and $|\beta|$. The red boundary encloses the parameter space accessible to our search, assuming the maximally W1-luminous Fortney et al. (2016) model. Within the accessible parameter space, the average 90% completeness threshold is 16.66 (16.54) for quadruplets (quintuplets).

9. RESULTS

Orbit linking yields 478 Keplerian quadruplets that pass our χ^2 cut described in §7.4. The locations of all successfully linked quadruplets are indicated by red dots in Figure 6. We visually inspect finder charts for all 478 candidate quadruplets, plus an additional 786 higher- χ^2 quadruplets up to `orbfit` χ^2 of 35.

In situations where all four transients have the same parity, there is very nearly zero relative parallax between our coadd epochs. Given this type of cadence, with observations spaced at integer numbers of years, Planet Nine’s apparent motion could be similar to that of a high proper motion star, if close to aphelion with large enough d_9 and e_9 . As a result, 113 of our 478 low- χ^2 quadruplets are single-parity linkages associated with known high proper motion stars, according to SIMBAD (Wenger et al. 2000). These known objects have $|\mu|$ ranging from ~ 200 mas/yr to $\sim 2,500$ mas/yr. As expected, we do not obtain low- χ^2 cross-parity linkages for these fast-moving stars, as their small parallaxes ($\lesssim 0.1''$) correspond to completely unbound trajectories of the type which `orbfit` is not intended to recover.

All quadruplets not associated with known moving objects appear to be affected, at least in part, by one or more problems with the transient catalog. Among these, the most common contaminants are:

- Diffraction spikes not fully captured by the masks of §6.3.2.
- Kernel matching artifacts that evaded the filtering of §6.3.3.
- Residuals of moderately bright static sources just far enough from the centroid to escape the 2MASS matching of §6.3.1.
- Extractions affected by the wings of bright stars, falling just outside of the §6.3.2 masks.

- Glints near extremely bright stars¹⁴.
- Variable static sources too faint to be captured by our 2MASS rejection step.

Examples of several such features are shown in Figure 9. As we do not find any plausible candidate among linked quadruplets, we quote a quadruplet-based limit of $W1 \geq 16.66$ (90% completeness) over the parameter space searched. Admittedly, it is not possible to precisely characterize the effect of visual inspection on our completeness.

We obtain 12 low- χ^2 quintuplets. All are single-parity linkages in which all five transients are associated with a unique known high proper motion star. Therefore we can rule out the presence of Planet Nine candidates in the parameter space searched with $W1 < 16.54$ at 90% completeness, independent of any visual inspection.

Our recourse to visual inspection is a deeply regrettable and unsatisfying aspect of the present search, and in future extensions of this work we will seek to eliminate this step. We estimate that by ignoring finders associated with known fast-moving stars and better masking diffraction spikes, we can reduce the number of finder chart inspections per unit search area by a factor of ~ 2 . Therefore, extending to the $\sim 80\%$ of sky amenable to our search will likely entail inspection of $\sim 4,400$ quadruplet finders.

10. CONCLUSION

We have developed a method to search for Planet Nine using custom, inertial coadds of WISE and NEOWISE-Reactivation W1 exposures. These epochal coadds extend much deeper than individual W1 exposures, and remove most of the inner solar system objects and instrumental artifacts that would otherwise overwhelm such a search. We have demonstrated our method by searching ~ 2000 square degrees of sky identified as a potentially likely present-day Planet Nine location. We do not find any plausible candidates, but are able to characterize our completeness, thereby ruling out the presence of Planet Nine in the parameter space searched at $W1 < 16.66$ (90% completeness).

Translating this value to an optical equivalent is both model and distance dependent, because the observed flux from reflected sunlight scales like d_9^{-4} , whereas that from intrinsic emission at $3\text{-}4\mu\text{m}$ scales like d_9^{-2} . Adopting the maximally W1-luminous Fortney et al. (2016) model and $d_9 = 650$ AU, $W1 = 16.66$ corresponds to an optical VR magnitude of 22.06. However, alternative models can have far bluer (VR–W1) colors at the same distance, and in such cases the corresponding optical constraint will be much weaker. For instance, assuming a (VR–W1) color characteristic of the Fortney et al. (2016) Table 2 model

with $T_{eff} = 50$ K, CH₄ abundance of $10^{-3} \times$ solar and $d_9 = 650$ AU, the corresponding limit is $VR \geq 19.16$.

In the near future, we will apply our methodology to all of the W1 sky area over which we expect substantial sensitivity. We anticipate significantly reduced sensitivity near the Galactic plane ($|b_{gal}| \lesssim 10^\circ$), due to the high number density of bright static sources and elevated background levels. Also, as illustrated in Figure 1, we expect very little sensitivity at $|\beta| \gtrsim 60^\circ$ given our current prescription for segmenting exposures into coadd epochs. However, by modifying our time-slicing rules near the ecliptic poles, we should in principle be able to derive deeper than usual constraints in these regions thanks to the disproportionately large number of WISE exposures at high $|\beta|$. Our methodology should be readily applicable to WISE W2, although the models of Fortney et al. (2016) suggest that a W2-detectable Planet Nine would likely require a relatively large mass (at least $\sim 35M_\oplus$), in tension with dynamical constraints (Brown & Batygin 2016).

The entire high-probability Planet Nine sky region singled out by de la Fuente Marcos & de la Fuente Marcos (2016) will be searchable with our method, as that footprint has both low $|\beta|$ and high $|b_{gal}|$. We will prioritize a search of that region, only 5% of which overlaps the footprint considered in this work.

Finally, because our analysis proceeds from WISE single-exposure images, we can naturally improve upon the present results by generating additional coadd epochs from future NEOWISER data releases.

We thank Roc Cutri for responding to queries at the WISE Help Desk. We thank Dustin Lang, Peter Eisenhardt and Mike Brown for their feedback.

This research made use of the NASA Astrophysics Data System (ADS) and the IDL Astronomy User’s Library at Goddard.¹⁵

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 publication makes use of data products from the Two Micron All Sky Survey, which is a joint project of the University of Massachusetts and the Infrared Processing and Analysis Center/California Institute of Technology, funded by the National Aeronautics and Space Administration and the National Science Foundation.

REFERENCES

- Alard, C. 2000, A&AS, 144, 363
 Batygin, K., & Brown, M. E. 2016, AJ, 151, 22
 Bernstein, G., & Khushalani, B. 2000, AJ, 120, 3323
 Bertin, E. 2006, in Astronomical Society of the Pacific Conference Series, Vol. 351, Astronomical Data Analysis Software and Systems XV, ed. C. Gabriel, C. Arviset, D. Ponz, & S. Enrique, 112
 Bertin, E., & Arnouts, S. 1996, A&AS, 117, 393

¹⁴ wise2.ipac.caltech.edu/docs/release/allsky/expsup/sec4_4g.html

¹⁵ Available at <http://idlastro.gsfc.nasa.gov>



FIG. 9.— Examples of common contaminants drawn from the transients contributing to our 478 low- χ^2 Keplerian quadruplets. Top row: epochal coadds. Middle row: corresponding reference images. Bottom row: HOTPANTS subtractions. Cutouts are $6.9' \times 6.9'$. Red circles are centered on the locations of linked difference image extractions. Left column: nebulosity associated with a glint. Middle left column: real variability of a faint, static W1 source not cataloged by 2MASS. Middle right column: a diffraction spike not fully masked by our prescription of §6.3.2. Right column: bright static source residuals.

- Bertin, E., Mellier, Y., Radovich, M., et al. 2002, in *Astronomical Society of the Pacific Conference Series*, Vol. 281, *Astronomical Data Analysis Software and Systems XI*, ed. D. A. Bohlender, D. Durand, & T. H. Handley, 228
- Brown, M. E., Bannister, M. T., Schmidt, B. P., et al. 2015, *AJ*, 149, 69
- Brown, M. E., & Batygin, K. 2016, *ApJ*, 824, L23
- Brown, M. E., Trujillo, C., & Rabinowitz, D. 2004, *ApJ*, 617, 645
- Cowan, N. B., Holder, G., & Kaib, N. A. 2016, *ApJ*, 822, L2
- Cutri, R. M., Mainzer, A., Conrow, T., et al. 2015, *Explanatory Supplement to the NEOWISE Data Release Products*, Tech. rep.
- Cutri, R. M., Wright, E. L., Conrow, T., et al. 2011, *Explanatory Supplement to the WISE Preliminary Data Release Products*, Tech. rep.
- . 2013, *Explanatory Supplement to the AllWISE Data Release Products*, Tech. rep.
- de la Fuente Marcos, C., & de la Fuente Marcos, R. 2016, arXiv:1607.05633
- Fienga, A., Laskar, J., Manche, H., & Gastineau, M. 2016, *A&A*, 587, L8
- Fortney, J. J., Marley, M. S., Laughlin, G., et al. 2016, *ApJ*, 824, L25
- Goldstein, D. A., D’Andrea, C. B., Fischer, J. A., et al. 2015, *AJ*, 150, 82
- Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, *ApJ*, 622, 759
- Holman, M. J., & Payne, M. J. 2016, arXiv:1604.03180
- Lang, D. 2014, *AJ*, 147, 108
- Linder, E. F., & Mordasini, C. 2016, *A&A*, 589, A134
- Luhman, K. L. 2014, *ApJ*, 781, 4
- Mainzer, A., Bauer, J., Cutri, R. M., et al. 2014, *ApJ*, 792, 30
- Mainzer, A., Bauer, J., Grav, T., et al. 2011, *ApJ*, 731, 53
- Meisner, A. M., & Finkbeiner, D. P. 2014, *ApJ*, 781, 5
- Meisner, A. M., Lang, D., & Schlegel, D. J. 2016, arXiv:1603.05664
- Sheppard, S. S., & Trujillo, C. 2016, arXiv:1608.08772
- Skrutskie, M. F., Cutri, R. M., Stiening, R., et al. 2006, *AJ*, 131, 1163
- Trujillo, C. A., & Sheppard, S. S. 2014, *Nature*, 507, 471
- Wenger, M., Ochsenbein, F., Egret, D., et al. 2000, *A&AS*, 143, 9
- Wright, E. L., Eisenhardt, P. R. M., Mainzer, A. K., et al. 2010, *AJ*, 140, 1868