

Multi-year stacking searches for solar system bodies

ALEX GERINGER-SAMETH,¹ NATHAN GOLOVICH,¹ AND KEITA IWABUCHI²

¹*Physics Division, Lawrence Livermore National Laboratory, 7000 East Avenue, Livermore CA, 94550, USA*

²*Center for Applied Scientific Computing, Lawrence Livermore National Laboratory, 7000 East Avenue, Livermore CA, 94550, USA*

ABSTRACT

Digital tracking detects faint solar system bodies by stacking many images along hypothesized orbits, revealing objects that are undetectable in every individual exposure. Previous searches have been restricted to small areas and short time baselines. We present a general framework to quantify both sensitivity and computational requirements for digital tracking of nonlinear motion across the full sky over multi-year baselines. We start from matched-filter stacking and derive how signal-to-noise ratio (SNR) degrades with trial orbit mismatch, which leads to a metric tensor on orbital parameter space. The metric defines local Euclidean coordinates in which SNR loss is isotropic, and a covariant density that specifies the exact number of trial orbits needed for a chosen SNR tolerance. We validate the approach with Zwicky Transient Facility (ZTF) data, recovering known objects in blind searches that stack thousands of images over six years along billions of trial orbits. We quantify ZTF’s sensitivity to populations beyond 5 au and show that stacking reaches most of the remaining Planet 9 parameter space. The computational demands of all-sky, multi-year tracking are extreme, but we demonstrate that time segmentation and image blurring greatly reduce orbit density at modest sensitivity cost. Stacking effectively boosts medium-aperture surveys to the Rubin Observatory single-exposure depth across the northern sky. Digital tracking in dense Rubin observations of a 10 sq. deg field is tractable and could detect trans-Neptunian objects to 27th magnitude in a single night, with deep drilling fields reaching fainter still.

1. INTRODUCTION

New solar system bodies are discovered primarily in a two-step procedure: source detection in individual images followed by the linking of position changes from night to night in a self-consistent, physical orbit. However, this approach is limited to the brightest objects, those that lie above the single-exposure detection threshold. At the same time, solar system bodies adhere to an iron fact of astronomy: for every object large enough and near enough to be detected in current imaging, there are vast numbers that lie below threshold (e.g. W. C. Fraser et al. 2014; M. Granvik et al. 2018; R. L. Jones et al. 2018).

These faint bodies still contribute photons to the camera, and their presence is subtly imprinted on the image data. They can be recovered by stacking images over multiple epochs, centering each image on the predicted position of the moving object, and repeating the procedure for many trial orbits. This approach is known as digital tracking, synthetic tracking, or track-before-detect in the solar system community, but is equivalent

to methods used throughout astronomy (image coaddition, matched filtering). Digital tracking on deep stacks of images offers the opportunity to make striking detections of “invisible objects” that are well below the noise floor in every individual image.

The history of digital tracking goes back at least three decades. It has most often been used in searches for trans-Neptunian and Kuiper Belt objects (TNOs and KBOs) in the outer solar system (A. L. Cochran et al. 1995; B. Gladman et al. 1998; E. I. Chiang & M. E. Brown 1999; R. L. Allen et al. 2001; G. M. Bernstein et al. 2004; M. J. Holman et al. 2004; W. C. Fraser et al. 2008; C. I. Fuentes et al. 2009; P. J. Whidden et al. 2019; M. Rice & G. Laughlin 2020; W. C. Fraser et al. 2024; H. Smotherman et al. 2024), but also for the main belt (A. N. Heinze et al. 2015, 2019; C. Zhai et al. 2020; A. Y. Burdanov et al. 2023) as well as for near-Earth objects (NEOs) and other inner solar system populations (M. Shao et al. 2014; C. Zhai et al. 2014, 2018; B. Wang et al. 2018; N. Lifset et al. 2021). A. Geringer-Sameth & S. M. Koushiappas (2012) explore the related idea of detecting a sub-threshold population of moving sources without identifying any particular object.

Stacking increases sensitivity since detector noise averages out over multiple observations while the signal from the source object adds coherently. This generally leads to the minimum detectable source flux scaling as $F_{\min} \propto 1/\sqrt{N}$, where N is the number of images in the stack. From this scaling, it would seem that digital tracking is particularly well-suited to large-area sky surveys that operate over many years. Such surveys will capture a large fraction of the solar system’s minor planets in thousands of individual exposures each.

However, digital tracking has so far been applied over relatively short time periods and narrow fields of view. The reason is that the computations are greatly simplified when the possible trajectories through a stack of images are linear and parallel over the duration of the observations. This limits the time baseline to around an hour for NEOs, one night for the main belt, and several days for KBOs (A. N. Heinze et al. 2015). Beyond these limits, curvature due to parallax and the ellipticity of the orbit cannot be ignored, the mapping from orbital parameters to pixel coordinates becomes expensive, and the stacking can no longer be done with an efficient shift-and-add procedure over whole images.

In this work, we consider the general digital tracking problem, in which solar system bodies must be modeled as moving on physical, nonlinear orbits across the full sky. The aim is to quantify (1) the sensitivity of optimal stacking for an arbitrary observing strategy and (2) the scale of the computational challenge of searching the parameter space of possible orbits.

We do this by first deriving how the detection significance varies as the trial orbit used to stack images departs from a moving object’s true orbit. This result is used to define a metric tensor on the space of orbital parameters, which quantifies how “distinguishable” a pair of nearby orbits are in a given survey. This helps define the proper spacing of trial orbits: it is wasteful to search along on multiple orbits that are indistinguishable, while a search that does not sample orbits densely enough will miss detectable objects. The metric can be used to construct a local coordinate system for each region of parameter space. In local coordinates, a unit-spaced lattice corresponds to a dense set of “barely distinguishable” trial orbits. The metric also leads to a natural notion of trial orbit density. This density, integrated over a region of orbital parameter space, gives the number of trial orbits that must be searched to fully cover the region, i.e. to detect all moving objects to within some tolerance of their maximum possible detection significance.

The spacing of trial orbits is a decision that must be made in every digital tracking search, whether in the linear regime or not. In previous work, these spacings,

for the most part, have been derived using heuristic arguments, e.g. that the on-sky tracking error be about a PSF width over the time span of the observations. Our theory gives a result similar to the heuristic methods in the linear regime and also applies to arbitrarily complicated orbital motion.

A few previous studies have considered nonlinear digital tracking. G. M. Bernstein et al. (2004) notably carried out a nonlinear TNO search using the Hubble Space Telescope in stacks of $N = 55$ images taken over approximately 24 hours. The time span and observing geometry were such that the 4-d parameter space of position and velocity used in linear tracking had to be augmented with a distance dimension. A heuristic argument similar to that used for the linear regime was employed to determine the gridding in distance.

A. H. Parker & J. J. Kavelaars (2010) develop a technique to construct sets of trial shift vectors used for image stacking. They first derive a scaling relation on the number of required trial orbits in the linear regime by considering, as we do, a tolerance on reduction of signal-to-noise ratio (SNR). For nonlinear motion, their constructive method works by generating a large library of random orbits and then prunes these to an optimum minimal set that maintains a maximum tracking error in every image. In comparison, our local coordinates can be used to construct a grid of orbits directly, while alternatively, our metric density can be used to “fill up” a region with a suitable number of trial orbits, similar to the Parker and Kavelaars method but without the pruning step.

A third work on nonlinear digital tracking is P. S. Gural et al. (2018). They present a method for constructing shift vectors used to stack space-based observations over short time intervals, where the nonlinear motion is due entirely to the spacecraft orbiting Earth. They construct a “zero-motion template” corresponding to a stationary asteroid at a distance of 1 au and then apply a series of linear transformations to generate other possible shifts to stack on. The procedure is highly efficient for this specific use case. Our method, by contrast, is agnostic to the parameterization and behavior of the motion model and could accommodate an Earth-orbiting observer as a particular case.

The mathematics of our metric resembles the linear covariance approximation in nonlinear least squares fitting of asteroid orbits to observations (e.g. K. Muinonen & E. Bowell 1993). In our case, we start from the expected SNR of a matched filter in image space rather than the χ^2 of sky-coordinate residuals, although the practical difference is minor. Our transformation to local coordinates by diagonalizing the metric has come up

previously in mapping out the nonlinear sky-plane uncertainty for known asteroids, the so-called Line (or Volume) of Variations (K. Muinonen 1996; A. Milani 1999; A. Milani et al. 2005; K. Muinonen et al. 2006). The two applications have many parallels, but we are concerned with efficiently covering a large parameter volume to discover new solar system bodies in a blind search rather than characterizing known ones.

The remainder of the paper is structured as follows. Section 2 motivates the metric by using a heuristic notion of distance between two orbits in a stack of images. In Sec. 3 we rederive the metric as measuring the relative reduction in detection significance caused by using the incorrect orbit in the matched filter search. The analytic formulas are verified by comparing them with simulations and Zwicky Transient Facility (ZTF) images of the minor planet Sedna (Sec. 3.4). In Sec. 4 we construct a local coordinate system for the space of possible orbits, and in Sec. 5 introduce the notion of density of trial orbits, again validating the theory in experiments with ZTF. Next, in Sec. 6 we apply the methods to the special case of linear motion, suitable for relatively short time spans and limited sky areas. In Sec. 7, we quantify the reach and difficulty of fully coherent digital tracking with six years of ZTF data, covering about $30,000 \text{ deg}^2$ of the northern sky (E. C. Bellm et al. 2019a). Finally, in Sec. 8 we discuss multi-year digital tracking in relation to the Rubin Observatory’s Legacy Survey of Space and Time (LSST) as well as tradeoffs between computational cost and sensitivity. In the appendix, we give a careful derivation of the expected detection significance for general optimal image stacking, which is used in Sec. 3.

2. A METRIC IN ORBITAL PARAMETER SPACE

2.1. Motivation

Digital tracking can be thought of as a trial-and-error approach to the discovery of moving objects. In each trial, an orbit is suggested and then the imaging data is analyzed to see whether there is evidence of a solar system body on that orbit. This is an example of forward modeling in that we start from the physical parameters describing the orbit and then predict what the signal looks like in data space, i.e. determining which survey images and pixels within those images contain the object. The detection significance is calculated by centering the images on the predicted location of the moving object, stacking them (with weights), and looking for excess flux at the center of the stacked image.

The true orbital parameters of the body are not known beforehand, and the search proceeds by enumerating a long list of trial orbits, testing each one for the presence

of a signal. The question is how to generate a suitable set of trial orbits.

A useful concept here is the density of trial orbits in the parameter space. If the density is too high, two trial orbits may be so similar that they intersect the same survey images and the same pixels within those images. In this case, the detection significance is highly correlated between these orbits. There will be either a detection on both orbits or on neither one. The opposite problem are voids, regions of orbital parameter space far from any trial orbit. If there is an actual solar system body in such a region, it will go undetected.

One approach is to construct a lattice of points in parameter space that have the correct spacing. This can be done by developing a distance measure in the space of possible orbits that quantifies how “distinguishable” two orbits are in the imaging data. The optimal spacing of the lattice is such that two adjacent orbits are barely distinguishable.

2.2. Heuristic derivation of the metric in image space

By what measure should we determine whether two orbits are distinguishable? The answer is given by the angular resolution or point spread function (PSF) of the imaging. The PSF can provide a metric in image space which we can then be pulled back to the parameter space of orbits.

Consider a trial orbit that intersects N survey images at pixel coordinates (x_i, y_i) for $i = 1, \dots, N$. In other words, a real object on that orbit would be present at location (x_i, y_i) in image i . Now imagine a slightly different orbit which shifts the object’s coordinates in each image by a small amount $(\Delta x_i, \Delta y_i)$. Denoting the width of the PSF in image i by b_i , a reasonable metric for the distinguishability of the two orbits is

$$\Delta s^2 = \frac{1}{N} \sum_{i=1}^N \frac{\Delta x_i^2 + \Delta y_i^2}{b_i^2}, \quad (1)$$

with the criterion for distinguishability being

$$\Delta s^2 \gtrsim 1. \quad (2)$$

If $\Delta s^2 < 1/N$, it is guaranteed that the difference between the orbits is less than a PSF width in every image. If there were an actual object on one of the orbits, stacking images along either one will give nearly identical detection significances. As Δs^2 approaches 1, the position shift is on average about 1 PSF width in each image. In this case, an object on one of the orbits will not contribute much flux at the location of the other orbit, and the detection significance for the two orbits will begin to decorrelate.

Setting the distinguishability threshold at $\Delta s^2 \approx 1$ instead of $\Delta s^2 \approx 1/N$ assumes that the shifts in different images are of similar magnitude. This is reasonable as there are generally many more image intersections than the number of degrees of freedom of possible orbits. It would be unusual to adjust one of the few orbital degrees of freedom and have the orbit's location remain unchanged in nearly every image except for a few in which it shifts by a large amount.

2.3. Pulling back to the space of orbits

The metric defined in Eq. 1 induces a metric on the space of possible orbits³. To see this, it will help to rewrite the metric in the language of differential geometry.

For a given trial orbit, let x^i be the image coordinates of an imaginary object on that orbit. The index i is a multi-index which runs over both the survey images containing the object and over the two dimensions (x and y) in each image.

In this notation, the metric from Eq. 1 takes the form

$$\Delta s^2 = \eta_{ij} \Delta x^i \Delta x^j, \quad (3)$$

where

$$\eta_{ij} = \frac{1}{N} \frac{1}{b_i^2} \delta_{ij} = \frac{1}{N} \text{Diag} \left(\frac{1}{b_i^2} \right). \quad (4)$$

Equation 3 and subsequent equations use the Einstein summation convention, where a repeated upper and lower index implies summation over all values the index can take (i.e. $\eta_{ij} \Delta x^i \Delta x^j$ means $\sum_i \sum_j \eta_{ij} \Delta x^i \Delta x^j$). In Eq. 4, $\delta_{ij} = 1$ when $i = j$ and 0 otherwise. It is straightforward to generalize Eq. 4 for an elliptical PSF, but below we only consider isotropic PSFs with a single b_i for the both the x and y directions in each image.

Now we consider the space of possible orbits. In linear digital tracking, phase space is four-dimensional: an orbit is specified by an initial position and angular velocity in the plane of the sky. In the general case of moving bodies in the solar system, the parameter space is six-dimensional. For concreteness, we use the Keplerian elements (semi-major axis, eccentricity, inclination, argument of perihelion, longitude of the ascending node, and mean anomaly at a given epoch), but the formalism equally applies to any other element sets, e.g. equinoctial elements, the moving object's 3-d position and velocity at a fixed moment in time, or the angular position plus parallax parameterization of G. Bernstein & B. Khusha-

lani (2000). As a geometric object, the metric we will derive is fundamentally independent of the choice of phase space coordinates.

Coordinates in orbital parameter space will be grouped into the vector θ^μ , where $\mu = 1, 2, \dots, 6$ (e.g. for Kepler elements θ^1 is semi-major axis, θ^2 is eccentricity, etc).

The forward model from orbital parameters to image space has two steps. First, the orbital parameters are used to determine where in the sky the solar system body was during each exposure of the survey. This yields a list of intersections, i.e. survey images that should contain an object on that orbit. Next, for each intersection there is a mapping from sky coordinates to image coordinates (also called pixel coordinates). The second step is done with either a basic tangent plane projection or with the FITS World Coordinate System (WCS, M. R. Calabretta & E. W. Greisen 2002) solution for the image.

The forward model is a mapping from θ^μ to x^i . In other words, a given set of orbital elements maps to a set of pixel coordinates in our survey images. The key to the metric is the Jacobian of this mapping,

$$J_\mu^i(\theta) \equiv \frac{\partial x^i}{\partial \theta^\mu}. \quad (5)$$

If the orbit θ under consideration intersects N survey images, the Jacobian J_μ^i can be thought of as a $2N \times 6$ matrix J (the 2 coming from the x and y coordinates for each image). The Jacobian is straightforward to compute numerically, either through finite differencing (what we do) or automatic differentiation (which requires coding the forward model using a framework like PyTorch or JAX).

The Jacobian describes how a small change in a solar system body's orbital elements $\Delta \theta^\mu$ maps to a small change in image coordinates Δx^i ,

$$\Delta x^i = J_\mu^i \Delta \theta^\mu. \quad (6)$$

Therefore, the distinguishability metric between the two orbits θ^μ and $\theta^\mu + \Delta \theta^\mu$ is

$$\begin{aligned} \Delta s^2 &= \eta_{ij} \Delta x^i \Delta x^j \\ &= g_{\mu\nu} \Delta \theta^\mu \Delta \theta^\nu, \end{aligned} \quad (7)$$

where $g_{\mu\nu}$ is defined by

$$g_{\mu\nu} = J_\mu^i \eta_{ij} J_\nu^j. \quad (8)$$

The object $g_{\mu\nu}$ is a metric tensor on orbital parameter space. It can be thought of as the 6×6 matrix g ,

$$g = J^T \eta J, \quad (9)$$

³ In this paper we will refer to the space of possible orbits as the parameter space, orbital space, or phase space.

where η is the diagonal matrix defined in Eq. 4.

Importantly, the metric is not constant but is a function of location within orbital parameter space. The space of possible orbits has become a Riemannian manifold with curvature. This is an example of information geometry (e.g. S.-i. Amari 1985): the sum of squared offsets in Eq. 1 is essentially a chi-square or negative log-likelihood function and the metric $g_{\mu\nu}$ is the Fisher information matrix of second derivatives of this log-likelihood with respect to the model parameters.

Figure 1 illustrates the transformation from orbital parameters to sky coordinates as seen by the Zwicky Transient Facility survey (ZTF; E. C. Bellm et al. 2019a,b). In each panel, the black curve shows the path of the dwarf planet Sedna during 2019. The dots mark the times that Sedna is present in a ZTF survey image, i.e. each dot denotes an intersection with the survey at position x^i . The gray curves and points show how the x^i 's would change under a change $\Delta\theta^\mu$ in two of the orbital parameters (upper: semi-major axis, lower: longitude of the ascending node). The arrows illustrate the Jacobian of the transformation $\partial x^i / \partial \theta^\mu$. The figure shows very large shifts in the orbital parameters for visu-




Figure 1. Illustration of the change in sky coordinates of an orbit under changes $\Delta\theta^\mu$ in the orbital parameters. Black curves show the path of Sedna during 2019. Points mark the intersections with ZTF, i.e. times when the object would be present in a survey image. The gray curves show how the orbital path shifts under a change in semi-major axis a (top) or longitude of the ascending node Ω (bottom). The arrows illustrate the Jacobian of the transformation. The changes in orbital parameters are exaggerated for illustration. In practice the shifts used in digital tracking will be at the arcsecond level.

alization. In practice, we will be concerned with changes in orbital elements $\Delta\theta^\mu$ that change the position on the sky by arcseconds.

3. THE METRIC DERIVED FROM EXPECTED DETECTION SIGNIFICANCE

The previous section provided an intuitive motivation for the image space metric in Eq. 1. The shift of the orbit in image space is only a proxy for what we really care about in digital tracking: how detection sensitivity falls off as the trial orbit deviates from an object's true orbit. Here, we derive a metric based directly on expected detection significance. This is the geometric approach introduced by B. J. Owen (1996) to construct template banks for gravitational wave searches, building on the work of B. S. Sathyaprakash & S. V. Dhurandhar (1991) and S. V. Dhurandhar & B. S. Sathyaprakash (1994).

3.1. Optimal image stacking

The optimal linear method of combining observations to detect a faint source is matched filtering (also called pixel weighting, Weiner filtering, or linear filtering; see, e.g. G. B. Rybicki & W. H. Press 1992; W. H. Press et al. 2007). In digital tracking, after a trial orbit predicts the location of a moving object in each image, the nearby pixels in all images are combined into a scalar test statistic via a weighted sum. The expected detection significance is (see A. Geringer-Sameth et al. 2015, Sec. IV)

$$\text{SNR} = \frac{\sum_i w_i s_i}{\sqrt{\sum_i w_i^2 \sigma_i^2}}, \quad (10)$$

where w_i is the weight assigned to pixel i , s_i is the expected flux in pixel i due to the object, and σ_i^2 is the variance in the pixel due to all background processes (this assumes that the noise in different pixels is uncorrelated). The sums are over all pixels in all the images the trial orbit intersects.

As a measure of detection significance, the SNR tells us how many standard deviations above background level the source will be detected at. It is a dimensionless quantity, i.e., $\text{SNR} = 5$ is a “5 sigma” detection.

The optimal pixel weights for source discovery are (A. Geringer-Sameth et al. 2015, Eq. 16),

$$w_i = \frac{s_i}{\sigma_i^2}. \quad (11)$$

Note that the expected significance in Eq. 10 is unchanged if all weights are scaled by a constant factor.

This is important since it implies the absolute magnitude of a solar system body need not be known in advance. One can choose any absolute magnitude when calculating s_i for the weights in Eq. 11 and the search remains optimal.

3.2. Analytic form for expected significance

In practice, it is impossible to carry out optimal weighting since we do not know the true orbit of a solar system body before we search for it. The pixel weights (Eq. 11) are instead calculated using a trial orbit. The relevant question is how the sensitivity decreases as a function of the difference between the true orbit and the trial orbit. To answer this, we start from a closed-form expression for the SNR (Eq. 10) in terms of the image-space offsets Δx^i between the trial and true orbits.

The theory is worked out in Appendix A; we summarize the results here. If a point source's true position in image i is x^i , but the images are stacked using trial positions $\tilde{x}^i = x^i + \Delta x^i$ (Δx^i being a 2-d vector for each image), the expected detection significance is

$$\text{SNR}(\Delta x) = \text{SNR}_{\max} \sum_{i=1}^N c_i \frac{\bar{\phi}_{\sqrt{2}b_i}(\Delta x^i)}{\bar{\phi}_{\sqrt{2}b_i}(0)}. \quad (12)$$

The sum is over the images i that the orbits intersect (Δx^i will always be small enough that the trial and true positions are in the same images). The PSF in image i is modeled as a 2-d Gaussian ϕ_b (Eq. A10) with standard deviation b_i , and $\bar{\phi}_b$ is the Gaussian density averaged over an image pixel (defined in Eq. A18; approximately equal to ϕ_b for small pixel sizes).

The c_i in Eq. 12 are a set of image weights that sum to 1 and $\text{SNR}_{\max}$ is the expected significance when the trial positions are equal to the true positions, i.e. $\text{SNR}_{\max}$ is maximum SNR that could be achieved with an infinite density of trial orbits. These are given by

$$\text{SNR}_{\max} = \frac{\sum_i \frac{A_i^2 q_i^{(0)}}{4\pi\sigma_i^2 (b_i/\varepsilon_i)^2}}{\sqrt{\sum_i \frac{A_i^2}{4\pi\sigma_i^2 (b_i/\varepsilon_i)^2}}} \quad (13)$$

and

$$c_i = \frac{c'_i}{\sum_i c'_i}, \quad \text{where} \quad c'_i = \frac{\tilde{A}_i^2 q_i^{(0)}}{\sigma_i^2 (b_i/\varepsilon_i)^2}. \quad (14)$$

In these equations, A_i is the total flux of the source in image i (with the same units as σ_i , e.g. ADU), while $\tilde{A}_i$ is the predicted flux at the trial positions (only relative flux differences between images matter for $\tilde{A}_i$, not their absolute scaling). As above, σ_i^2 is the pixel noise variance (now assumed the same for all pixels near the

source in image i), ε_i is the pixel scale (so that b_i/ε_i is the PSF size in pixel units), and $q_i^{(0)}$ (shorthand for $q_{\varepsilon_i/\sqrt{2}b_i}^{(0)}$, defined in Eq. B24) is a finite pixel size correction factor that converges to 1 for small pixels (i.e. when $\varepsilon_i \ll b_i$). The above equations can be found in Appendix A.4.

As a function of i , the sequence A_i is the solar system body's light curve. The light curve is a function of the Sun-body-observer geometry as well as intrinsic properties of the body, traditionally parameterized by an absolute magnitude H and slope parameter G in the "H, G system" (E. [Bowell et al. 1989](#)). Absolute magnitude does not need to be searched over as it is an overall scaling in $\tilde{A}_i$. It would be straightforward to augment orbital parameter space to include G , but for efficiency we fix it at the conventional value $G = 0.15$, leading to a negligible loss in sensitivity except in very peculiar circumstances involving the opposition surge effect. In what follows, the difference between the trial and true orbits will always be so small that the trial light curve is proportional to the true light curve ($\tilde{A}_i \propto A_i$). The constant of proportionality depends on the object's (unknown) absolute magnitude.

3.3. From expected significance to metric

Equation 12 describes how the expected detection significance falls off as the location of the trial object deviates from the true location, i.e. as a function of Δx^i . We now proceed as in Sec. 2.3 to relate this falloff to the separation of the two orbits in orbital parameter space, $\Delta\theta^\mu = \tilde{\theta}^\mu - \theta^\mu$, where $\tilde{\theta}$ and θ denote the trial and true orbits.

When the trial and true orbits are similar, the separations Δx^i in the individual images are small. Expanding Eq. 12 to lowest order in Δx^i we have

$$\text{SNR}(\Delta x) = \text{SNR}_{\max} \left(1 - \sum_i \frac{1}{4} c_i q_i^{(2)} \left| \frac{\Delta x^i}{b_i} \right|^2 + \mathcal{O} \left| \frac{\Delta x^i}{b_i} \right|^4 \right), \quad (15)$$

where we used the Taylor series from Eq. B26, the finite pixel size correction factor $q_i^{(2)} = q_{\varepsilon_i/\sqrt{2}b_i}^{(2)}$ defined in Eq. B25, and the fact that c_i sum to 1 (Eq. 14).

There is no linear Δx^i term because the test statistic (i.e. set of optimal weights) is derived from a minimization procedure (A. [Geringer-Sameth et al. 2015](#)). This means that first order changes in the weights (i.e. changes in the trial orbit) lead to second-order changes in detection significance. This is a general feature of optimal linear (Weiner) estimators (e.g. [G. B. Rybicki &](#)

W. H. Press 1992) and is what allows us to construct a metric, which must be quadratic in the orbital parameters $\Delta\theta^\mu$.

Next we rewrite the quadratic term in tensor notation as in Eqs. 3 and 4 (dropping the $\mathcal{O}|\Delta x/b|^4$ terms from now on),

$$\begin{aligned}\text{SNR}(\Delta x) &\approx \text{SNR}_{\text{max}} (1 - \eta_{ij} \Delta x^i \Delta x^j) \\ &= \text{SNR}_{\text{max}} (1 - \Delta s^2),\end{aligned}\quad (16)$$

where

$$\eta_{ij} = \left(\frac{1}{4} c_i q_i^{(2)} \frac{1}{b_i^2} \right) \delta_{ij}.\quad (17)$$

As in Sec. 2.3, the indices i and j are multi-indices that index both images and the two dimensions within each image (though c_i , $q_i^{(2)}$, and b_i are equal for the x and y directions within an image).

From Eq. 16 we see that metric distance

$$\Delta s^2 = \eta_{ij} \Delta x^i \Delta x^j\quad (18)$$

represents the fractional loss in expected detection significance when using a trial orbit that deviates from the true orbit by Δx^i in each image.

The new metric is very similar to the heuristic one from Sec. 2.2. In Eq. 1, each image intersection contributes to the metric with an effective weight of $1/N$. In the new metric, the contribution of image i is c_i (Eq. 14), which captures the relative importance of the image, taking into account the noise level, light curve, and sharpness of the PSF in that image. If all intersections are of similar importance, then the c_i 's will all be approximately $1/N$.

We also have the same criterion (Eq. 2) for two orbits to be independent or distinguishable. From Eq. 16 we see that when $\Delta s^2 < 1$ the expected significance is close to its maximum but drops quickly once $\Delta s^2 \gtrsim 1$. Thus the condition $\Delta s^2 = 1$ sets the scale at which we expect the trial positions to become suboptimal for detecting a source.

Of course, SNR does not actually go negative when Δs^2 reaches 1 as Eq. 16 seems to imply. Instead, Eq. 12 (or, more transparently, Eq. A15) shows that the falloff will be roughly squared exponential. Therefore, the better target for quadratic approximation is probably $\log[\text{SNR}(\Delta x)]$ rather than $\text{SNR}(\Delta x)$. From Eq. 15 or Eq. 16 one sees that

$$\log \text{SNR}(\Delta x) = \log \text{SNR}_{\text{max}} - \Delta s^2 + \mathcal{O} \left| \frac{\Delta x^i}{b_i} \right|^4, \quad (19)$$

so that

$$\text{SNR}(\Delta x) \approx \text{SNR}_{\text{max}} \exp(-\Delta s^2). \quad (20)$$

We can use this exponential form to predict that the detection significance drops to 50% of its maximum when $\Delta s^2 \approx 0.7$ and to 10% when $\Delta s^2 \approx 2.3$. Jensen's inequality applied to Eq. A15 shows that this approximation is conservative: the actual expected detection significance declines more gradually than Eq. 20 predicts.

Finally, we retrace our steps from Sec. 2.3 to transfer the metric to orbital parameter space. For small deviations of the orbital elements, we have $\Delta x^i = J_\mu^i \Delta \theta^\mu$ (Eqs. 5 and 6), which leads to

$$\Delta s^2 = g_{\mu\nu} \Delta \theta^\mu \Delta \theta^\nu, \quad (21)$$

where

$$g_{\mu\nu} = J_\mu^i \eta_{ij} J_\nu^j, \quad (22)$$

with η_{ij} now given by Eq. 17. The behavior of the expected significance for small orbital separations $\Delta \theta$ is given by Eq. 20.

3.4. Application to Sedna

We demonstrate these results for the minor planet Sedna as observed by ZTF. We query Sedna's orbital elements from NASA/JPL Horizons⁴ (J. D. Giorgini & JPL Solar System Dynamics Group; J. D. Giorgini et al. 1996), propagate the orbit to the observation times of all ZTF r and g -band difference images, and identify those that should contain Sedna. There are 971 intersections, excluding those where Sedna was in a masked part of the image. For each image we extract the PSF width and zero-point magnitude from the FITS header, estimate the pixel noise level, and calculate Sedna's predicted flux in the image (see Sec. 7.1 for details). These are used to form the pixel weights (Eq. A4) and evaluate the observed matched filter detection significance SNR_{obs} . The stacked detection significance for Sedna in ZTF is $\text{SNR}_{\text{obs}} = 127$, while its median SNR in individual images is 3.

Next, for each intersection we generate an ideal, zero-noise version of the difference image, where the value in each pixel is the Gaussian PSF integrated over the pixel and scaled by the predicted flux in the image. Since the predicted fluxes were calculated with an arbitrary overall scaling (arbitrary absolute magnitude H), the ideal images are scaled by an overall factor such that the analytic SNR_{max} from Eq. 13 is the same as SNR_{obs} . This scaling factor corresponds to an absolute magnitude for Sedna of $H = 1.87$, consistent⁵ with 1.829 ± 0.048 measured by D. L. Rabinowitz et al. (2007).

⁴ Data retrieved 2025-Mar-14.

⁵ A fractional uncertainty in flux of $1/\text{SNR}_{\text{obs}}$ gives an uncertainty of $\Delta H \approx 0.01$. The agreement with D. L. Rabinowitz et al. (2007) is perhaps coincidental since we use Makemake's




Figure 2. Detection significance in stacked images as the semi-major axis of the trial orbit deviates from Sedna’s true orbit. The black curve is for “perfect” images simulated without noise while white points correspond to actual ZTF data. The red dotted curve is the analytic form of expected significance (Eq. 12) and the dashed gray shows the metric approximation (Eq. 20). The predicted significance matches the simulation and observations for all trial orbits, while the metric approximation is valid near the true orbit.

We then repeatedly perform the stacked image search for different trial orbits, varying the semi-major axis while holding the other orbital elements fixed. Figure 2 shows the results for both the simulated images (black curve) and the actual ZTF images (white circles). The analytic form for the expected detection significance (Eq. 12) is plotted as the red dotted curve, showing excellent agreement with the simulated images. The significance curve for the actual ZTF data also falls off as predicted, verifying that our assumptions are safe, in particular that the PSF can be approximated as Gaussian. We note that the only free parameter that has been adjusted here is an overall scaling (the absolute magnitude of Sedna). The fact that significance in ZTF peaks at $\Delta a = 0$ verifies our orbit propagation as well as the astrometric accuracy of the ZTF pipeline.

Finally, the approximate significance using the metric (Eqs. 20 and 21) shows good agreement with the real and simulated data near the peak (gray dashed curve). The metric tensor $g_{\mu\nu}$ (Eq. 22) is evaluated once for the true orbit. The varying trial orbit only enters through $\Delta\theta^\mu = (\Delta a, 0, 0, 0, 0)$ in Eq. 21. The falloff in significance is notably non-Gaussian away from the peak and cannot be described globally by any single Gaussian. However, the metric serves its purpose of predicting how far the trial orbit can deviate from the true orbit before suffering a loss of significance. When the metric approximation

(Eq. 20) predicts that the significance has dropped by 10% ($\Delta s^2 = 0.11$), the actual significance has decreased by 9.3%. For predicted 50% and 90% drops ($\Delta s^2 = 0.69$ and 2.3), the actual decreases are 37% and 60%.

4. LOCAL COORDINATES FOR ORBITAL ELEMENTS

The metric $g_{\mu\nu}$ leads to a natural coordinate system on orbital phase space analogous to locally inertial coordinates in general relativity (GR). According to the equivalence principle, the local experience of a freely falling observer in a gravitational field is indistinguishable from flat Minkowski spacetime. Mathematically, this means that at any point in spacetime there exists a local coordinate system in which the GR metric is flat ($g_{\mu\nu} = \text{Diag}(-1, 1, 1, 1)$), and that it stays flat in the nearby vicinity ($\partial g_{\mu\nu}/\partial x^\sigma = 0$), so that the coordinates represent an inertial frame in the absence of gravity (S. Weinberg 1972, Ch. 3). In differential geometry such coordinates are known as (Riemann) normal coordinates.

For digital tracking, “locally inertial orbital elements” make the relation between orbits and detection significance flat (i.e. Euclidean): a unit step in any direction in local coordinates changes the detection SNR by a fixed fraction. In this work, we will stop short of constructing true normal coordinates, but carry out the first step of Euclideanizing the metric at a point in phase space by a linear transformation. The next step, a quadratic transformation that zeros out the first derivatives of the metric, would be an interesting topic to explore.

One use of these local coordinates is to construct a trial orbit lattice such that every possible orbit of interest is at a maximum metric distance $\Delta s^2 < \Delta s_{\text{max}}^2$ from a lattice point. This ensures that a digital tracking search completely covers a region of orbital space in the sense that every potentially detectable solar system body will be found. An optimal lattice fills the space with a minimum number of points. The metric also allows for a simple determination of the density of this optimal trial orbit lattice (Sec. 5), which quantifies the computational difficulty of the search in various regions of orbital parameter space.

4.1. Coordinate transformation

Normal coordinates, denoted ξ^μ , are an alternate set of orbital elements that parameterize the possible orbits in the neighborhood of a particular orbit θ_0^μ . In other words, there is an invertible mapping between the two sets of coordinates, $\theta^\mu = \theta^\mu(\xi^\nu)$, where $\xi^\mu = 0$ maps to $\theta^\mu = \theta_0^\mu$. In the new coordinates, the metric at the point ξ^μ has components $h_{\mu\nu}(\xi)$ given by

$$h_{\mu\nu}(\xi) = g_{\alpha\beta}(\theta) \frac{\partial \theta^\alpha}{\partial \xi^\mu} \frac{\partial \theta^\beta}{\partial \xi^\nu}, \quad (23)$$

colors rather than Sedna’s when predicting image fluxes (see Sec. 7.1).

which is just the statement that the infinitesimal metric distance between the two orbits ξ^μ and $\xi^\mu + d\xi^\mu$ is $ds^2 = h_{\mu\nu} d\xi^\mu d\xi^\nu$.

The actual transformation is built by examining Eq. 23 order by order in ξ^μ . Specifically, the mapping is written (e.g. L. Brewin 1996)

$$\theta^\mu = \theta_0^\mu + \Lambda^\mu{}_\nu \xi^\nu + \frac{1}{2} B^\mu{}_{\nu\rho} \xi^\nu \xi^\rho + \dots, \quad (24)$$

and the constant coefficients $\Lambda^\mu{}_\nu$ and $B^\mu{}_{\nu\rho}$ are chosen so that $h_{\mu\nu}$ is diagonal and has first derivatives equal to zero. We only consider the linear transformation $\Lambda^\mu{}_\nu$ to keep things simple, and set all higher coefficients to zero.

4.2. Linear transformation of the orbital elements

The goal here is to choose the linear part of the transformation $\Lambda^\mu{}_\nu$ so that at $\xi^\mu = 0$ the metric takes the Euclidean form $h_{\mu\nu}(0) = \delta_{\mu\nu}$. From Eqs. 23 and 24, we have

$$h_{\mu\nu}(0) = g_{\alpha\beta}(\theta_0) \Lambda^\alpha{}_\mu \Lambda^\beta{}_\nu = \delta_{\mu\nu}. \quad (25)$$

Finding $\Lambda^\mu{}_\nu$ means solving the matrix equation

$$\Lambda^T g \Lambda = I, \quad (26)$$

where g is the metric in θ -coordinates evaluated at θ_0 , I is the identity matrix, and $\Lambda^\mu{}_\nu$ is the row μ column ν entry of the square matrix Λ .

Equation 26 specifies Λ only up to an overall rotation of coordinates⁶ (if Λ solves Eq. 26 so does ΛO for any orthogonal matrix O). A particular solution is to let each column of Λ be a normalized eigenvector of g scaled by the inverse square root of the corresponding eigenvalue. This solution always exists (i.e. the eigenvalues are positive and eigenvectors mutually orthogonal) since the metric g (Eq. 22) is positive definite (which in turn comes from the diagonal image-space metric $\eta_{\mu\nu}$ (Eq. 17) having all positive entries).

Numerically, one can diagonalize g directly (which is efficient for a 6×6 matrix), but it is more revealing to arrive at the transformation by singular value decomposition (SVD; e.g. W. H. Press et al. 2007, Ch. 2.6). The metric g is first factorized (see Eq. 9 or 22) as

$$g = J^T \eta J = (\sqrt{\eta} J)^T (\sqrt{\eta} J),$$

where the square root of the positive diagonal matrix η (Eq. 17) is taken elementwise. At this point $\sqrt{\eta} J$ has units, which makes matrix decomposition awkward.

The values in column μ have units of $1/[\text{units of } \theta^\mu]$ so a choice of length scale for each orbital element can be used to remove them. A. Milani et al. (2005) suggest some simple scalings for different sets of orbital elements (e.g. for Kepler elements, dividing semi-major axis by a_0 , inclination by 180° , and the other angular elements by 360°). We find it natural to use the metric itself to define a length scale ℓ^μ for each coordinate,

$$\ell^\mu = (g_{\mu\mu})^{-1/2}. \quad (27)$$

Taking a step in an orbital element $\theta^\mu \rightarrow \theta^\mu + \ell^\mu$ and holding the others constant gives a metric change of $\Delta s^2 = 1$, meaning ℓ^μ is a step size that is just distinguishable in the data.

Let ℓ be the diagonal matrix $\text{Diag}(\ell^\mu)$. Then $\sqrt{\eta} J \ell$ is dimensionless and the SVD is performed:

$$\sqrt{\eta} J \ell = U D V^T. \quad (28)$$

The diagonal matrix D (6×6) contains the square roots of the eigenvalues of $\ell g \ell$. The columns of U ($2N \times 6$, for N image intersections) are an orthonormal basis for the possible shifts in image-space (scaled by $\sqrt{\eta}$) corresponding to any small change in orbital elements. The columns of V (6×6) are the orthonormal eigenvectors of $\ell g \ell$. A solution for Λ is then

$$\Lambda = \ell V D^{-1}.$$

The coordinate transformation between the original orbital parameters θ^μ and local coordinates ξ^μ is given in matrix form by

$$\theta = \theta_0 + \ell V D^{-1} \xi \quad (29)$$

$$\xi = D V^T \ell^{-1} (\theta - \theta_0), \quad (30)$$

where θ and ξ are vectors with components θ^μ and ξ^μ .

The linear transformation is already quite useful on its own, even without the higher-order terms in Eq. 24. The dimensionless ξ -coordinates make it trivial to calculate the drop in detection significance as a trial orbit departs from a true orbit at θ_0 . The fractional reduction in significance is just the squared Euclidean distance from the origin to the point ξ corresponding to the trial orbit, i.e.

$$\Delta s^2 = (\xi^1)^2 + \dots + (\xi^6)^2. \quad (31)$$

Had we carried on to the second order transformation using the $B^\mu{}_{\nu\rho}$ coefficients in Eq. 24, the above result would extend to a local region. In other words, throughout a neighborhood surrounding θ_0 , the expected loss in significance would be simply $\Delta s^2 = (\Delta \xi^1)^2 + \dots + (\Delta \xi^6)^2$, where $\Delta \xi^\mu$ are the normal coordinate differences between the trial and true orbits, and this would hold to

⁶ This is likely related to the observation by A. Milani et al. (2005) that the Line of Variations and “weak direction” depend on the choice of orbital elements.

first order anywhere within the neighborhood, not just at θ_0 .

In practice, the local coordinates remain flat to a good approximation out to different distances along the different ξ^μ directions. The coordinate ξ^6 , corresponding to the smallest eigenvalue of g , can be identified with Milani’s “weak direction” (A. Milani 1999; A. Milani et al. 2005). Taking a unit step in ξ^6 , although it is only predicted to change the relative SNR by $\Delta s^2 = 1$, is a very large step in the original orbital elements, as can be seen from Eq. 29 (D_{66}^{-1} is large). Physically, what is happening is that ξ^6 is a nearly degenerate direction in the original orbital parameters, where moving along it barely changes the sky position in the intersected images at all. The change in θ^μ for $\Delta \xi^6 = 1$ is often large enough that the linear theory breaks down and the image coordinates Δx^i change by much more than a PSF size. In contrast, for the ZTF data, it is usually possible to move by thousands of unit steps along ξ^1 and ξ^2 and the shifts Δx^i remain confined to a compact region of the sky.

Figure 3 illustrates the local coordinates for the numbered minor planet 612911, a high-inclination outer solar system object currently 57 au from the Sun. The helical path shows its sky position over the duration of the ZTF survey. The small dots with attached arrows correspond to times when 612911 was present in a ZTF exposure. The arrows show the shift in sky coordinates corresponding to a change in each of the local coordinates by $\Delta \xi^\mu = 1$, holding the others constant. The arrows are magnified by a factor of 1000 for visibility. In the linear approximation, $\Delta \xi^\mu = 1$ leads to $\Delta s^2 = 1$, which corresponds to a shift in the sky of about a PSF size in every image.

Each set of arrows corresponds to a column of the U matrix in the SVD decomposition of the Jacobian (Eq. 28). These columns are a basis for describing all possible small changes in the orbit of this object. The various “eigenorbits” have nicely distinct behaviors. For example, ξ^1 and ξ^2 are simple translations of the orbital path, ξ^4 rotates the path, and ξ^3 and ξ^5 are stretches roughly parallel and perpendicular to the long-term sky motion. In contrast, a similar plot of the sky shifts caused by changing each Kepler element separately shows that most of the Kepler elements are highly degenerate with one another. Except for semi-major axis, changing the other Kepler elements appears to simply translate the orbit on the sky (similar to ξ^1 and ξ^2 above). The local coordinates isolate the subtle ways that these translations are actually not uniform over the duration of the observations. In our experience, the translational and rotational “modes” always occur,

but the stretching modes for some orbits are not always obviously separated as parallel and perpendicular to the long-term motion.

The inset in each panel of Fig. 3 shows the linear combination of the different Kepler elements corresponding to the local coordinate (Eq. 29). In the ξ^1 panel, for example, the bar heights show the absolute value of the Kepler element shifts $\Delta \theta^\mu$ corresponding to $\Delta \xi^1 = 1$. The $\Delta \theta^\mu$ are normalized by their natural length scales ℓ^μ (Eq. 27) and the bars are plotted in logarithmic scaling. The horizontal black line corresponds to $|\Delta \theta^\mu / \ell^\mu| = 1$ and the gray lines are steps by factors of 10. The first three local coordinates correspond to Kepler element shifts that would be expected to shift the sky position by about 1 pixel in each image ($\Delta \theta^\mu \approx \ell^\mu$). The next three local coordinates correspond to much larger changes in Kepler elements that nonetheless conspire to keep the on-sky changes small. The last local coordinate ξ^6 corresponds to a near-exact degeneracy in orbital elements: $\Delta \xi^6 = 1$ is expected to be a barely distinguishable change in sky coordinates but involves changing the Kepler elements by many thousands of times the scale that would cause a shift by one PSF. This is a weak direction where the linear theory breaks down and the actual change in sky coordinates will cause more than a $\Delta s^2 = 1$ drop in detection significance, as seen by the longer arrows.

4.3. Use of the local coordinates to construct a lattice of trial orbits

The local coordinates can be used to build a lattice of trial orbits for digital tracking. The spacing between points in ξ -coordinates should be such that the maximum squared distance from any point to the closest grid point is $\Delta s_{\max}^2$. This ensures that the worst-case relative reduction in SNR due to mismatch between the trial and true orbits is $\exp(-\Delta s_{\max}^2)$.

We find it is possible to grow local coordinate patches around a seed orbit θ_0 by moving along each ξ^μ coordinate until the metric distance Δs^2 reaches some value (e.g. $\Delta s^2 = 3000$). However, if a direction ξ^μ is a weak direction, the extent of the patch is just the $\Delta \xi^\mu$ (which will be less than 1) such that $\Delta s^2 = 1$. The result is a rectangular region in ξ -coordinates (corresponding to a hyperparallelogram in the original orbital elements) that can be either filled with a lattice of trial points, or filled with self-avoiding random points sampled according to the density described in the next section. Details on piecing together local coordinate patches to cover a large region of orbital parameter space will be explored in future work.




Figure 3. Illustration of local coordinates for orbital parameter space ξ^μ near the orbit of numbered minor planet 612911, as seen over six years of ZTF imaging. Each small dot and attached arrow correspond to a time when the object was present in a ZTF image. The arrow shows the shift in the orbital path corresponding to a change $\Delta\xi^\mu = 1$ holding the other local coordinates fixed (for visibility, arrows are magnified by a factor of 1000 and are colored according to time of year). The sets of arrows for the different ξ^μ form a basis for exploring the local region of orbital parameter space. The inset in each panel shows the combination of Kepler element shifts corresponding to a change $\Delta\xi^\mu = 1$ in the local coordinate. The height of a bar is the logarithm of the absolute value of the Keplerian element shift $\Delta\theta^\nu$ in units of the Kepler element’s natural length scale ℓ^ν (see Eqs. 29 and 27). The horizontal black line corresponds to $|\Delta\theta^\nu/\ell^\nu| = 1$ and each gray line is a factor of 10.

5. THE DENSITY OF TRIAL ORBITS

The metric leads to a direct answer to the question: how many trial orbits must we test to fully cover a region of orbital parameter space? A region is “covered” when every orbit in the region has a squared metric distance at most $\Delta s_{\max}^2$ from the nearest trial orbit. Thus we can be sure that all solar system bodies in the region will be detected with SNR within a fraction $\Delta s_{\max}^2$ of the maximum possible SNR.

In the neighborhood of each orbit θ_0 there is a local coordinate system ξ^μ in which the metric is Euclidean (Sec. 4). The placement of trial orbits is a packing problem in these local coordinates. There must be roughly one trial orbit per sphere of squared radius $\Delta s_{\max}^2$ in ξ -space. An infinitesimal volume dV_ξ in ξ -space maps to the infinitesimal volume dV_θ in the original orbital elements, where $dV_\theta = |g(\theta_0)|^{-1/2} dV_\xi$ and $|g(\theta_0)|$ is the determinant of the metric tensor in θ -coordinates (this can be derived from Eq. 26).

The volume of a d -dimensional Euclidean ball with radius r is $V_d(r) = c_d r^d$, where c_d is a constant prefactor (in six dimensions $V_6(r) = \frac{\pi^3}{6} r^6$; the prefactor is roughly 5 for $d = 4$ to 7). Therefore, in θ -space we need one trial orbit per volume $dV_\theta = |g(\theta_0)|^{-1/2} c_d (\Delta s_{\max}^2)^{d/2}$. Equivalently, the required density of trial orbits in or-

bital parameter space is

$$\rho(\theta) = \frac{|g(\theta)|^{1/2}}{c_d (\Delta s_{\max}^2)^{d/2}}. \quad (32)$$

We define

$$\rho_{\text{trial}}(\theta) = |g(\theta)|^{1/2}, \quad (33)$$

to be the benchmark trial orbit density since $\Delta s_{\max}^2$ will usually be chosen to be of order 1. This density is a generally covariant quantity. Calculating the number of trial orbits by integrating $\rho(\theta)$ over a finite volume of phase space will give the same result for any equivalent parameterization of orbits (e.g. Keplerian elements, equinoctial elements, Cartesian position and velocity vectors at an epoch, etc).

5.1. Validation of trial orbit density in a blind search

We tested the density formalism by performing blind digital tracking searches for two known minor planets in six years of ZTF data (see Sec. 7.1 for a description of the data). One of the authors selected distant objects from the Minor Planet Center, queried Horizons for their “true” orbital elements (both had uncertainty code 1), and then defined search regions in phase space enclosing the true orbits. Search regions were 6-d rectangular boxes in Keplerian elements where the width in each element θ^μ was chosen to be ~ 600 times the element’s

natural length scale ℓ^μ (Eq. 27). This box size is such that the orbits in the search region collectively intersect several thousand ZTF images, which fit into memory on a single compute node. Each search box was then shifted so that the true orbit ended up at a uniformly random location within the box. The density ρ (Eq. 32) was calculated and used to predict the number of trial orbits required to discover the hidden object. Specifically, density was calculated at the 3^6 combinations of lower bound, upper bound, and midpoint along each dimension and then multiplied by the volume of the box. The standard deviation among the density samples was around 15% and we used the mean to estimate the required number of orbits. The box boundaries and number of trial orbits were then passed to other members of the team who performed the search with our high-performance computing (HPC) digital tracking software (N. Golovich et al. 2025). The code draws trial orbits using a space-filling Sobol sampling method, which ensures fairly uniform coverage of the search region.

Both of our experiments were successful in detecting the hidden minor planets. We detail the results for the second one, (341520) Mors-Somnus. This object intersects 780 ZTF images with a predicted median magnitude of $V = 21.5$, just fainter than the limiting magnitude for a single image (the Minor Planet Center has no reported observations of Mars-Somnus from ZTF or Palomar). It was predicted that 2.3×10^9 trial orbits would be needed to detect it at $\Delta s^2 < 1$ and twice this number were trialed by the HPC code, meaning we expect about 2 trial orbits to have SNR within a factor $\exp(-1)$ of $\text{SNR}_{\max}$ corresponding to the true orbit. After the search, the SNR corresponding to the true orbit was measured to be $\text{SNR}_{\max} = 60.3$, and there were 4 trial orbits with $\text{SNR} > 22.2$ with the highest being 46.1.

Figure 4 confirms that Mors-Somnus was indeed recovered in the blind search. The large circles show the offset in pixel coordinates between the orbit with the highest SNR and the true orbit. The top trial orbit stays within about 1 pixel over the course of the six-year observations. The other sets of points in this figure correspond to the other top 3 trial orbits. Numerically optimizing the SNR starting at the best trial orbit converges to the true orbit. There are many trial orbits which come close to the true orbit in a subset of intersections. This is reflected in the falloff of significance with Δs^2 (see Fig. 2 and the discussion after Eq. 20). In this experiment there were 345 trial orbits with $\text{SNR} > 10$ and 1800 with $\text{SNR} > 7$. We did not explore whether optimizing SNR from these starting points converges to the true orbit, but a two-step procedure of trial-and-




Figure 4. Recovery of (341520) Mors-Somnus in a blind search experiment. The large circles show the offsets in image coordinates between the trial orbit with the highest observed SNR and Mors-Somnus’s true orbit. Each circle corresponds to one of the 780 intersections with ZTF (color shows the date). This trial orbit gave a detected at $\text{SNR} = 46$, whereas stacking along the true orbit gives $\text{SNR}_{\max} = 60$. The small sets of dots show the next three top trial orbits having $\text{SNR} = 31$ (blue), 28 (orange), and 24 (green).

error search at a lower density, followed by optimization of high-SNR orbits may be efficient.

6. APPLICATION TO LINEAR MOTION

In this section we apply the metric, local coordinates, and trial density to the simple case of objects moving with constant velocity in two dimensions. This is the motion model used for most digital tracking and can be used for stacking observations over short time baselines.

There are four orbital elements in this case: the position (x_0, y_0) at the epoch t_0 and the (constant) velocities (v_x, v_y) in the two directions, i.e. $\theta^\mu = (x_0, v_x, y_0, v_y)$. For motion in the sky tangent plane, the velocities are the components of proper motion (e.g. in arcsec/day). Observations take place at the times $t_1, \dots, t_N$ measured from the epoch. At these times the object is located at

$$\begin{pmatrix} x_1 \\ y_1 \\ \vdots \\ x_N \\ y_N \end{pmatrix} = \begin{pmatrix} x_0 + v_x t_1 \\ y_0 + v_y t_1 \\ \vdots \\ x_0 + v_x t_N \\ y_0 + v_y t_N \end{pmatrix}.$$

Therefore, the Jacobian of the transformation from orbital elements to sky positions is

$$J_{\mu}^i = \begin{pmatrix} \frac{\partial x_1}{\partial x_0} & \frac{\partial x_1}{\partial v_x} & \frac{\partial x_1}{\partial y_0} & \frac{\partial x_1}{\partial v_y} \\ \vdots & \vdots & \vdots & \vdots \\ \frac{\partial y_N}{\partial x_0} & \frac{\partial y_N}{\partial v_x} & \frac{\partial y_N}{\partial y_0} & \frac{\partial y_N}{\partial v_y} \end{pmatrix} = \begin{pmatrix} 1 & t_1 & 0 & 0 \\ 0 & 0 & 1 & t_1 \\ \vdots & \vdots & \vdots & \vdots \\ 1 & t_N & 0 & 0 \\ 0 & 0 & 1 & t_N \end{pmatrix}.$$

For simplicity, assume each image has the same pixel noise level σ^2 , a Gaussian PSF with standard deviation b , small pixels, and that objects maintain constant brightness over time. Then the c_i weights in the image-space metric η_{ij} (Eq. 17) are all $1/N$ and η_{ij} is simply a constant diagonal matrix with entries $1/(4Nb^2)$.

The metric in orbital parameter space (Eq. 22) is

$$g_{\mu\nu} = J_{\mu}^i \eta_{ij} J_{\nu}^j = \frac{1}{4b^2} \begin{pmatrix} 1 & \langle t \rangle & 0 & 0 \\ \langle t \rangle & \langle t^2 \rangle & 0 & 0 \\ 0 & 0 & 1 & \langle t \rangle \\ 0 & 0 & \langle t \rangle & \langle t^2 \rangle \end{pmatrix},$$

where $\langle t \rangle = \frac{1}{N} \sum_i t_i$ and $\langle t^2 \rangle = \frac{1}{N} \sum_i t_i^2$. If we choose the epoch t_0 to be the mean observation time, then $\langle t \rangle = 0$ and $\langle t^2 \rangle$ is the variance of the distribution of observation times, a measure of the survey baseline. The metric is then simply

$$g_{\mu\nu} = \frac{1}{4b^2} \text{Diag}(1, \langle t^2 \rangle, 1, \langle t^2 \rangle).$$

The metric does not depend on location in orbital elements space. Phase space is flat and there is a single set of normal coordinates $(\xi^1, \dots, \xi^4)$ for which the metric is everywhere Euclidean. The transformation is simply

$$\begin{aligned} x_0 &= 2b\xi^1 & v_x &= \frac{2b}{\sqrt{\langle t^2 \rangle}} \xi^2 \\ y_0 &= 2b\xi^3 & v_y &= \frac{2b}{\sqrt{\langle t^2 \rangle}} \xi^4. \end{aligned}$$

In a digital tracking search, the trial orbit lattice could be a square grid in the 4-d ξ -space. The grid spacing should be such that the maximum squared distance from any point to the closest grid point is Δs_{max}^2 . This ensures that the worst-case relative reduction in SNR due to mismatch between the trial and true orbits is $\exp(-\Delta s_{\text{max}}^2)$ (Eq. 20). If the ξ -grid has spacing δ in all four dimensions, the worst-case scenario is when the true orbit is at the center of a grid cell, giving $\Delta s_{\text{max}}^2 = 4(\delta/2)^2 = \delta^2$. For $\Delta s_{\text{max}}^2 = 1$ (i.e. the tolerable reduction of significance is a few 10s of percent) we should set the grid spacing to be $\Delta \xi^{\mu} = \delta = 1$ in every dimension.

Translating back to the original orbital elements, the spacing between trial initial positions would then be $\Delta x_0 = \Delta y_0 = 2b$ and the spacing between trial velocities would be $\Delta v_x = \Delta v_y = 2b/\sqrt{\langle t^2 \rangle}$. This is similar to the spacing heuristically derived by A. N. Heinze et al. (2015, Eq. 4), who set the velocity spacing to be $\sqrt{2}$ times the maximum acceptable blur divided by the time spanned by the observations. D. Lang et al. (2009) also note that the relevant timescale for measuring proper motion is the variance of the observation times rather than the total time span.

The benchmark density of trial orbits (Eq. 33) needed for digital tracking to discover linear motion is simply

$$\rho_{\text{trial}} = \frac{\langle t^2 \rangle}{64b^4}. \quad (34)$$

The required number of trial motions is just the product of density and the volume of phase space to be searched. If the sky area surveyed is A_{sky} and the goal is to find sources moving in any direction with proper motion up to v_{max} , the phase space volume is $A_{\text{sky}} \pi v_{\text{max}}^2$. Combining these factors and using Eq. 32 with $c_4 = \pi^2/4$, the total number of trials needed is

$$N_{\text{trial}} = \frac{A_{\text{sky}} v_{\text{max}}^2 \langle t^2 \rangle}{16\pi b^4 (\Delta s_{\text{max}}^2)^2}.$$

For the simple assumption of constant survey cadence, the time variance is $\langle t^2 \rangle = T^2/12$, where T is the total survey time. Putting in units, the number of trials needed per sky area surveyed is (P. McGill et al. 2025)

$$\frac{N_{\text{trial}}}{A_{\text{sky}}} = \frac{6.6 \times 10^5}{\text{deg}^2} \frac{\left(\frac{v_{\text{max}}}{1 \text{ arcsec/day}}\right)^2 \left(\frac{T}{\text{day}}\right)^2}{\left(\frac{\text{PSF FWHM}}{\text{arcsec}}\right)^4 (\Delta s_{\text{max}}^2)^2}. \quad (35)$$

7. APPLICATION TO THE ZWICKY TRANSIENT FACILITY

The results in this paper are in preparation for a digital tracking search using the Zwicky Transient Facility survey (E. C. Bellm et al. 2019a,b). In this section we calculate the sensitivity limit when stacking images across the multi-year survey and the density of trial orbits needed to carry out such a search.

7.1. Data preparation and orbit sampling

We obtained all r and g -band difference imaging produced by ZTF (F. J. Masci et al. 2019; IRSA 2022) over approximately six years from March 20, 2018 to February 29, 2024. This comprises 849,750 30-second exposures, stored in about 50 million ccd-quadrant FITS files, each containing a 3072×3080 -pixel image covering

a sky area of $\approx 0.75 \text{ deg}^2$. The survey covers the entire sky above declination -30° . The limiting magnitude for a 5σ point source detection in individual r -band images is 20 to 21.3 (16th to 84th percentile) and similar in g band. The median seeing is 2.0 and 2.2 arcsec in r and g images (FWHM) using a fiducial pixel scale of 1.012 arcsec/pix. ZTF observations are ongoing but we will refer to the “ZTF survey” to mean the 6-year span we have on disk.

To assess sensitivity to a diversity of minor planet populations we generated 2^{27} (about 130 million) orbits over the 6-d space of Keplerian orbital elements using a space-filling Sobol sampler⁷. Sampling was log-uniform in semi-major axis between 5 and 1000 au, uniform in eccentricity between 0 and 1, uniform in $\cos$ (inclination) between 0° and 180° , and uniform between 0° and 360° in argument of perihelion, longitude of ascending node, and mean anomaly. We do not explore semi-major axes less than 5 au because, as shown below, the density of required trial orbits increases dramatically in the inner solar system.

For each sample orbit θ we identified the ZTF intersections (the images that would contain an object on that orbit) and computed the pixel coordinates within those images using the WCS parameters for the TAN projection (M. R. Calabretta & E. W. Greisen 2002) in the FITS header (ignoring the TPV distortion coefficients). Orbits are modeled as Kepler ellipses around an inertial Sun, neglecting perturbations from the planets⁸. We obtained the Jacobian J_μ^i of the transformation from θ to pixel coordinates (Eq. 5) by numerical differentiation⁹.

The density of trial orbits and the expected detection significance require the (relative) object flux in each image, the image PSFs, and pixel noise variances. The PSF is assumed Gaussian (Eq. A10) with b_i calculated from the SEEING parameter in the FITS header. A robust pixel noise variance σ_i^2 is calculated with an iteratively sigma-clipped MAD (median absolute deviation




Figure 5. The fraction of orbital parameter space intersecting the ZTF survey in various numbers of images. The boundary of the gray shaded region shows the fraction of orbits that will be captured in at least one ZTF image (orbits are sampled as described in Sec. 7.1 and binned in semi-major axis). The black curves show the distribution of the number of ZTF intersections for the orbits that have at least one intersection. Specifically, of all orbits that are present in ZTF, the curves show the fraction that have greater than 100, 500, 800, etc intersections. For example, the median outer solar system orbit observed by ZTF will be present in a stack of about 1100 images.

from the median) estimator applied to the image’s pixel values. The predicted flux $\hat{A}_i$ is the light curve computed in the H,G magnitude system for V band, color-corrected to r or g using colors we derived for Makemake¹⁰, and then converted to image units using the zero point magnitude in the FITS header¹¹. We use Makemake’s color to optimize for the outer solar system but do not expect the color correction to have a strong effect on the sensitivity of the search.

Figure 5 shows the distribution of the number of ZTF images intersected by the population of orbits as a function of semi-major axis. The boundary of the gray shaded region gives the fraction of sampled orbits that have at least one intersection with ZTF. For high semi-major axes, a solar system object’s sky position is roughly constant over 6 years and the fraction of orbits with a ZTF intersection just reflects the three quarters of the sky covered by the survey. At low semi-major axes, nearly all bodies are moving fast enough that they

⁷ We used the `scipy` (v1.15.1) implementation `scipy.stats.qmc.Sobol` with default settings.

⁸ This simplification does not affect the results of this paper on limiting magnitude and the density of distinguishable orbits. For the outer solar system, Keplerian motion is also sufficient for the actual stacking search, maintaining arcsec precision over several years. However, a more complete force model must be used to perform a stacking search in the inner solar system.

⁹ We used the finite difference package `numdifftools` (P. A. Brodtkorb 2025; J. D’Errico 2006), the `numdifftools.Jacobian` function with `order=4`, `num_steps=5`, and an adaptive procedure for setting `base_step`. The estimated maximum relative error in any element of the Jacobian was $\sim 10^{-7}$ for the median orbit, and larger than 1% in about 1 in 10^5 sample orbits.

¹⁰ Makemake’s normalized reflectance spectrum is taken from A. Alvarez-Candal et al. (2020, Fig. 1) and multiplied by the solar spectrum at zero airmass (ASTM International 2019) to obtain an unnormalized spectrum. Its AB magnitudes are then found for the Pan-STARRS1 bandpasses (J. L. Tonry et al. 2012) used by ZTF. The resulting colors are $V - r = 0.24$ and $V - g = -0.37$.

¹¹ MAGZP in the header is adjusted using the CLRCOEFF coefficient and Makemake’s $g - r$ color.

will enter the northern hemisphere at some point over the 6-year baseline and be captured by ZTF. The black curves in the figure are conditional on orbits that have at least one intersection. They show the fraction of such orbits with greater than 100, 500, etc intersections in ZTF. Nearly all moving objects observable with ZTF will typically be seen in hundreds to a few thousand images. Image stacking sensitivity, scaling as $\text{SNR} \propto \sqrt{N}$ for N intersections, means there is a substantial benefit from stacking in years-long, wide-area surveys like ZTF — assuming the computational demands can be met.

7.2. Sensitivity limit of a stacking search

For each sample orbit, the absolute magnitude corresponding to an $\text{SNR}_{\text{max}} = 10$ detection, denoted H_{10} , is computed with Eq. 13. This represents the faintest object (on the given orbit) a stacking search can be expected to discover. For this limiting magnitude we also compute the expected SNR and apparent V magnitude in the individual exposures making up the stack. We choose $\text{SNR}_{\text{max}} = 10$ instead of, say, 5 because the naive significance (sometimes called local significance) must be reduced due to the huge number of orbits that are tested. However, it is easy to scale our sensitivities to other thresholds: expected significance is proportional to an overall scaling in the source flux so the absolute magnitude corresponding to, say, $\text{SNR}_{\text{max}} = 60$ would be $H_{60} = H_{10} - 2.5 \log_{10}(60/10)$.

An object's distance is the dominant factor in H_{10} . Figure 6 shows how H_{10} scales with heliocentric distance at the midpoint of the ZTF survey, which we denote $d_{\odot}$. For each bin in $d_{\odot}$ the thin black lines show the 5th percentile, median (dashed line), and 95th percentile of H_{10} among all orbits in that bin which have at least one ZTF intersection. The jagged line corresponds to the orbit with the largest value of H_{10} , i.e. the most favorably observed orbit.

The variation in depth among orbits at equal distance is mainly due to differing numbers of intersections. For a moving object with absolute magnitude H that appears in N images, the detection significance scales approximately as (see Eq. A16)

$$\text{SNR} \propto \frac{10^{-H/2.5}}{d_{\odot}^4} \sqrt{N},$$

which implies

$$H_{10} = -10 \log_{10} d_{\odot} + 1.25 \log_{10} N + \text{const.}$$

We fit a linear relation between H_{10} and either $\log_{10} d_{\odot}$ alone or $\log_{10} d_{\odot}$ and $\log_{10} N$ to the orbits with $d_{\odot} > 20$ au (half the sample is used for fitting and the other




Figure 6. Depth of a ZTF stacking search as a function of a solar system body's heliocentric distance at the midpoint of the survey. The thin black lines are upper limits on an object's absolute magnitude (H_{10}) that yield an $\text{SNR} \geq 10$ stacked detection. From top to bottom they show the 5th percentile, median (dashed), 95th percentile, and maximum (jagged) H_{10} among all orbits that have at least one intersection with ZTF. In general, objects with $H < H_{10}$ will be detectable by stacking. The gray points show the currently known population of distant objects, and the blue markers highlight the population used to infer the orbit of the hypothesized Planet 9. The three regions in the upper right show the possibilities for Planet 9 itself, based on the reference models of M. E. Brown & K. Batygin (2022). The thin black contour surrounds 99% of the sample Planet 9's that would have already been detected by ZTF, Pan-STARRS, or DES (78% of the reference population). The green 99%-contour shows candidates that have been previously undetectable but could be seen at $\text{SNR} > 10$ with ZTF stacking (18% of the population). The red heat map shows the remaining 1% of orbits that have at least 10 ZTF intersections but are too faint to be detected by stacking (an additional 2.7% of the reference population is undetectable because it falls outside the ZTF footprint).

half to compute the root mean square of the residuals). The results are

$$H_{10} = -9.91 \log_{10} \frac{d_{\odot}}{\text{au}} + 23.83 \quad \text{rms} = 0.60$$

$$H_{10} = -9.91 \log_{10} \frac{d_{\odot}}{\text{au}} + 1.64 \log_{10} N + 18.97 \quad \text{rms} = 0.17.$$

Including the number of intersections soaks up most of the variation. The remainder is due to variation in seeing, noise level, and zero point among the individual images. Interestingly, the best fit coefficient for $\log_{10} N$ is 1.64 rather than the expected 1.25, meaning that depth seems to scale somewhat more favorably than $\sqrt{N}$, perhaps because N is correlated with sky position and therefore image quality in some way.

For context and motivation, Fig. 6 also highlights several solar system populations. The gray points show the Minor Planet Center’s distant object list¹² with the distant Kuiper Belt objects shown as blue circles.

The distant KBOs are selected according to the same criteria as in M. E. Brown & K. Batygin (2021) ($a > 150$ au and perihelion > 42 au). Apparent clustering in this population has been used to hypothesize the existence of a large planet (Planet 9) at many hundreds of au (K. Batygin & M. E. Brown 2016; K. Batygin et al. 2019) (but see also C. Shankman et al. (2017) for a cautionary view). These outer KBOs pile up near the faint end of what ZTF stacking can detect. The benefit of stacking will be in finding these objects over an enormously larger sky area than has been explored at this depth (see Fig. 9). Enlarging this population is important for the Planet 9 hypothesis and would alleviate potential selection bias of previous narrow-field observations (see, e.g. C. Shankman et al. 2017).

Finally, the top right of Fig. 6 shows possibilities for Planet 9 itself. This reference population is described by M. E. Brown & K. Batygin (2021), who sampled a posterior distribution of Planet 9 orbits and masses given the orbits of the known distant KBOs (A. Siraj et al. (2025) obtained a somewhat different distribution). The reference population also indicates whether each sample Planet 9 would have been detected by searches for linked, single-epoch detections performed in ZTF, the Dark Energy Survey (DES), and Pan-STARRS1 (M. E. Brown & K. Batygin 2022; M. Belyakov et al. 2022; M. E. Brown et al. 2024). The H mag of each Planet 9 candidate is computed from the V mag listed in the table and its heliocentric distance (assuming zero phase angle and equal geocentric and heliocentric distances). Then, for each Planet 9 orbit, the limiting H_{10} for ZTF stacking is found as described above. The black contour encloses 99% of the Planet 9 samples that are already ruled out by ZTF, DES, and Pan-STARRS1. This constitutes 78% of the reference population. The green contour encloses 99% of the remaining Planet 9 samples that would be detectable at $\text{SNR} = 10$ in a stacked ZTF search (18% of the reference population). Stacking pushes to both fainter fluxes and to further distances compared to the parameter space already explored. The red heat map shows the 1% of Planet 9 reference orbits that intersect at least 10 ZTF images but at magnitudes too faint to be detected by stacking. A remaining 2.7% of allowed orbits intersect fewer than 10 images and are not shown. Stacking in ZTF, therefore, has the ability




Figure 7. The density at which Keplerian orbital elements must be tested to discover solar system bodies in a 6-year ZTF stacking search. The curves show the 5th, median, and 95th percentiles over orbital parameter space conditioned on $d_{\odot}$, an object’s heliocentric distance at the midpoint of ZTF.




Figure 8. The number of trial orbits required to fully search the 6-year ZTF data. The curves show the integral of the required orbit density $\rho(\theta)$ over all orbits for which the object is beyond heliocentric distance $d_{\odot}$ at the midpoint of ZTF. The solid curve is for the full range of bound orbits, dashed is restricted to inclinations less than 10° , while dot-dashed gives the number of orbits needed to saturate each square degree patch of sky.

to detect or rule out the large majority (80%) of the remaining parameter space for Planet 9.

7.3. Density of trial orbits in ZTF

Equations 17, 22, and 33 are used to compute the density of trial orbits $\rho_{\text{trial}}(\theta)$. Figure 7 shows the distribution of trial orbit density as a function of heliocentric distance of the orbit $d_{\odot}$. As with H_{10} , density is determined primarily by $d_{\odot}$ and the curves show the 5th percentile, median, and 95th percentile of our sampled orbits within each $d_{\odot}$ bin. The units of density are $1/(\text{volume of Keplerian orbital elements space})$, i.e. $\text{au}^{-1} \text{deg}^{-4}$.

¹² 6,441 objects as of July 18, 2025.

The integral of ρ_{trial} over a finite orbital region is the total number of distinguishable orbits that must be tested to fully search that region. While the trial orbit density depends on the parameterization of orbital phase space, its integral is independent of phase space coordinates and directly quantifies the computational scale of the digital tracking problem.

We compute $\int \rho_{\text{trial}}(\theta) d^6\theta$ by quasi-Monte Carlo integration using our 2^{27} sample orbits (including a factor $a/\sin(i)$ in the integrand to adjust for the non-uniform orbit sampling in semi-major axis and inclination). Figure 8 shows the total number of trial orbits needed to fully search various regions of orbital phase space using the 6-year ZTF data (using Eq. 32 with $\Delta s_{\text{max}}^2 = 1$, $c_6 = \pi^3/6$). For each value of $d_{\odot}$, the curves give the number of trial orbits needed to find objects at heliocentric distance $d_{\odot}$ and beyond. The solid line includes all bound orbits (i.e. the full range for each Keplerian element), while the dashed line shows orbits with inclinations less than 10° . The dot-dashed line gives the number of trial orbits needed to “fill” a square degree patch of sky. It is computed by restricting the integration region to orbits for which the object is in a particular patch of sky at the midpoint of ZTF. The number of required orbits grows linearly with the area of the patch and there is minimal dependence on the location of the patch within the ZTF footprint. We checked the accuracy of the integration by dividing our sample orbits into 32 subsamples and re-calculating the results, finding negligible variation except at the lowest value of $d_{\odot}$.

The number of required trial orbits scales extremely rapidly with inverse distance, approximately as $\rho_{\text{trial}} \propto d_{\odot}^{-7}$. This harsh scaling makes a fully coherent 6-year blind stacking search for inner solar system objects prohibitively expensive. Even the outer solar system will be a challenge for the world’s largest computers, at least with the brute-force, trial-and-error approach. We discuss possible methods to ease the computational burden in Sec. 8.1.

8. DISCUSSION

Multi-year digital tracking can increase the sensitivity of a moderately sized survey telescope to that of a world-class instrument. With ZTF, for example, stacking $N \approx 10^3$ images essentially improves the detection threshold by 3 to 4 magnitudes over the nominal $r = 21$ single-exposure limit. This stacked depth is on par with individual exposures of the Rubin Observatory’s Legacy Survey of Space and Time (LSST; [Ž. Ivezić et al. 2019](#)), which will detect point sources at $r \approx 24.6$.




Figure 9. Median V magnitude in individual ZTF exposures leading to an $\text{SNR} = 10$ detection in a stacked search. For comparison, the median 5σ magnitude limit for source detection in single ZTF images is around 21 (similar in r and g), while for Rubin LSST the single-image depth will be $r \approx 24.6$. The solid and dashed black lines are the ecliptic and galactic planes. This figure includes all sampled orbits that would appear in at least one ZTF image. Declinations below -30° are masked, which erases the very small fraction of orbits that enter the ZTF footprint after, or exit before, the midpoint of the survey.

Figure 9 shows the effective depth ZTF achieves over the wide swath of sky it surveys. Solar system bodies with apparent V magnitude brighter than the limit shown will be detected at $\text{SNR} = 10$ in a stacked search. The depth varies somewhat across the sky because of ZTF’s observing pattern. To produce the figure, for each of our sample orbits, we scale its absolute magnitude so that the stacked SNR is 10 and then find the median V magnitude such an object would have in the ZTF images it intersects. Orbits are then binned spatially into $1^\circ \times 1^\circ$ regions according to their sky location at the midpoint of the survey. For each spatial bin, the color shows the median over the orbits in that bin. We use our full orbit sample, but the figure is unchanged if only orbits with semi-major axes greater than 10, 20, or 100 au are included.

The implication is that surveys like ZTF or Pan-STARRS ([N. Kaiser et al. 2002](#)) can provide a northern hemisphere capability that compliments LSST. This might be used, for example, in pre-discovery searches of LSST alerts ([S. Li et al. 2025](#)). While the required density of trial orbits for near-Earth objects is enormous, the volume of search space would be much reduced if an approximate initial orbit is determined by LSST. It should also be possible to do digital tracking within LSST itself, which would reach unprecedented sensitivity to faint outer solar system objects. The deep drilling fields, in particular, will reach a stacking depth




Figure 10. Scaling of depth and trial orbit density of digital tracking with ZTF as a function of the time span T over which images are stacked. The black line shows the detection significance, scaling roughly as $\text{SNR} \propto \sqrt{T}$ as expected for stacking. The colored lines are the median trial orbit density ρ_{trial} for sample orbits in 5 bins of heliocentric distance $d_{\odot}$ (log-spaced from 10 to 1000 au; distance increases up the plot from green to purple). The density scaling depends on the orbit (thin red lines): distant orbits scale as $\rho \propto T^3$ while those closer in scale as $\rho \propto T^4$. The difference in scaling between SNR and density means that large computational savings can be had at relatively minimal loss of sensitivity.

of $r \gtrsim 28.5^{13}$ (F. B. Bianco et al. 2022). Section 8.2 considers digital tracking with Rubin over short time spans.

8.1. Tradeoff between processing effort and depth

The number of trial orbits that must be tested in wide-area, multi-year digital tracking is clearly enormous (Sec. 7.3 and Fig. 8): a six-year northern hemisphere search for all solar system bodies further than 100 au would require 10^{19} trial orbits.

One move toward feasibility would be to divide the survey into shorter time segments and perform stacking separately on the images from each segment. To quantify the cost-benefit, we take 2^{15} of our sample orbits (Sec. 7.1) and recalculate the density and expected detection significance as the time span of ZTF is reduced. Specifically, for each time span T we ignore images taken more than $T/2$ from the midpoint of ZTF. Figure 10 shows the results. The black curve is the scaling of expected detection significance SNR_{max} with T (the median over all sample orbits at each value of T is shown; there is little variation with orbit). The scaling is roughly $\sqrt{T}$, as expected for image stacking. The colored curves show the scaling of the required trial orbit

density ρ_{trial} . The sample orbits are binned by heliocentric distance at the midpoint of ZTF ($d_{\odot}$ used above) into 5 log-spaced bins between 10 and 1000 au and the density ratio compared to $T = 6$ yrs is found for each orbit and each T . The curves show, at each value of T , the median ratio for the orbits in a bin. There is some variability in the scaling relation, with orbits with larger $d_{\odot}$ having a shallower density slope (T^3 for $d_{\odot} \sim 600$ au vs. T^4 for $d_{\odot} \sim 15$ au). The reason for the apparent break at 500 days is also unclear, but perhaps the variations in slope are caused by a change in the effective dimensionality of the motion model.

In any case, the very different scaling between SNR and density means that as the time span of the sub-stacks decreases, the search gets easier “faster” than it loses sensitivity. This suggests a hierarchical algorithm for years-long digital tracking. The first step would be to stack over short time spans but with a relatively low detection threshold. Trial orbits or regions of phase space which show promise can then be searched more densely in stacks over longer time spans. This is somewhat reminiscent of the semi-coherent algorithm K. M. Smith (2016) proposed to search for radio pulsars in long time series.

There are other ways to trade sensitivity for computational cost. Equation 34 shows that for a linear motion model the number of trial orbits increases as the fourth power of the inverse PSF size. For nonlinear digital tracking the scaling is likely to be even steeper: from Eqs. 17 and 22, for a general 6-parameter motion model the metric g is b^{-2} times a 6×6 matrix, so that $\rho_{\text{trial}} = |\det g|^{1/2} \propto b^{-6}$. This suggests that increasing the effective PSF size will drastically reduce the number of trial orbits. This can be done, for example, by combining pixels to produce lower-resolution images or convolving images with blurring kernels. Artificially increasing the PSF will decrease significance, but at a much slower rate than b^{-6} ($\text{SNR}_{\text{max}} \propto b^{-1}$ according to Eq. 13, though that equation does not account for the pixel noise becoming correlated).

8.2. Digital tracking with Rubin

With a single night of dedicated observation, the Rubin Observatory can conduct an exceptionally deep search for outer solar system objects. Over the course of $T = 8$ hrs, Rubin can take $N \approx 800$ 30-second exposures of a 9.6 deg^2 field of view¹⁴. The single-exposure

¹³ LSST Baseline Strategy, <https://survey-strategy.lsst.io/baseline/index.html>, accessed 2025-09-10.

¹⁴ Declinations below -8.3° are amenable to 8-hour continuous viewing from Cerro Pachón at airmass less than 2, which includes the ecliptic between 201° and 339° ecliptic longitude. See Appendix C for details.

5σ magnitude limit in r band is around 24.6, with an effective PSF FWHM of 0.7 arcsec. This means that a stacking search can make an $\text{SNR}_{\text{max}} = 10$ detection of a moving object at $r = 27.5$.

The feasibility of stacking these observations can be found via Eq. 35, which is applicable to the outer solar system at this timescale. The number of trial orbits required to test all proper motions less than 21.6 arcsec/day (the mean motion of Neptune) is $N_{\text{trial}} = 1.4 \times 10^9$ for $\Delta s_{\text{max}}^2 = 1$. Our current searches for nonlinear motion in ZTF are already at this scale and linear stacking is considerably more efficient (e.g. P. J. Whidden et al. 2019; N. Lifset et al. 2021).

For $\Delta s_{\text{max}}^2 = 1$, objects will likely be detected at $\exp(-1)$ of their maximum SNR (Eq. 20). Equivalently, objects found at $\text{SNR}_{\text{obs}} = 10$ will need to be 1.1 mag brighter than the magnitude corresponding to $\text{SNR}_{\text{max}} = 10$. The single-night search would thus be complete to $r = 26.4$. A depth arbitrarily close to 27.5 can be achieved by increasing the number of trials. This is about 2 magnitudes shy of the unprecedented depth achieved by G. M. Bernstein et al. (2004) using HST, but covers a sky area around 500 times the size.

9. CONCLUSIONS

In this paper we assessed the possibilities and challenges of carrying out multi-year, nonlinear digital tracking searches for moving objects with wide-field astronomical surveys. Here we summarize our findings:

- The sensitivity and computational difficulty of digital tracking for an arbitrary observing strategy can be quantified with a differential geometry formalism on orbital parameter space. The framework leads to an optimal sampling of trial orbits.
- A fully coherent digital tracking search on multi-year, all-sky data has the capability to turn current telescope surveys into the equivalent of the Rubin Observatory’s single exposures. This is an enticing opportunity to extract vastly more detections from ecliptic plane surveys (e.g. Pan-STARRS) and synoptic surveys like ZTF (Fig. 9).
- However, doing so requires truly enormous computation that scales rapidly as the distance to the targeted populations decreases (Fig. 8). Nonetheless, the distant solar system, including most

of Planet 9’s untested parameter space, can be probed with existing ZTF survey data (Sec. 7.3).

- There are several tradeoffs that can enable tractable searches across wide portions of parameter space. We highlighted a few, such as hierarchical searches over short time spans and smoothing or binning pixels (see Sec. 8.1).
- Targeted searches within all-sky surveys, such as the LSST deep drilling fields, dense single or few-night campaigns, or searches in restricted portions of parameter space, are feasible (though often still requiring world-class HPC) and offer avenues for high-impact research, particularly for outer solar system population studies.

ACKNOWLEDGEMENTS

The authors thank Peter McGill, Sage Li, Lila Braff, Bob Armstrong, Scott Perkins, and Michael Schneider for their input over the course of this study.

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. This work was supported by the Lawrence Livermore National Laboratory LDRD Program under Project 2023-ERD-044. The document release number is LLNL-JRNL-2011617.

Based on observations obtained with the Samuel Oschin Telescope 48-inch and the 60-inch Telescope at the Palomar Observatory as part of the Zwicky Transient Facility project. ZTF is supported by the National Science Foundation under Grants No. AST-1440341 and AST-2034437 and a collaboration including current partners Caltech, IPAC, the Oskar Klein Center at Stockholm University, the University of Maryland, University of California, Berkeley, the University of Wisconsin at Milwaukee, University of Warwick, Ruhr University, Cornell University, Northwestern University and Drexel University. Operations are conducted by COO, IPAC, and UW.

Software: Astropy (Astropy Collaboration et al. 2013, 2018, 2022), Astroquery (A. Ginsburg et al. 2019), Matplotlib (J. D. Hunter 2007), numdifftools (P. A. Brodtkorb 2025; J. D’Errico 2006), NumPy (C. R. Harris et al. 2020), pandas (Wes McKinney 2010; The pandas development team 2020), SciPy (P. Virtanen et al. 2020), SSAPy (J. E. Meyers et al. 2025).

APPENDIX

A. EXPECTED DETECTION SIGNIFICANCE IN MATCHED FILTERING

In this appendix we derive the expected detection significance for optimal image stacking when there is a mismatch between the model used for the matched filter and the true source.

A.1. Pixel sums as convolutions

The expected detection significance using an arbitrary set of pixel weights is given by Eq. 10,

$$\text{SNR} = \frac{\sum_i w_i s_i}{\sqrt{\sum_i w_i^2 \sigma_i^2}}, \quad (\text{A1})$$

where s_i is the object's true flux in pixel i , w_i is the weight assigned, and σ_i^2 is the noise variance in the pixel (noise in different pixels is assumed to be independent; σ_i has the same units as s_i , e.g. ADU). Each sum is over all images and over the pixels within each image. For now consider the inner sum, so that the index i just runs over the pixels in a single image.

The true flux s_i in pixel i can be written in terms of a continuous, normalized flux density $f(x)$. The flux density is scaled by the object's brightness A , shifted to its location within the image x_0 , and then integrated over pixels. Mathematically,

$$\begin{aligned} s_i &= \int_{\text{pixel } i} A f(x - x_0) d^2x \\ &= \int_{\mathbb{R}^2} A f(x - x_0) \varepsilon^2 W_\varepsilon(x - x_i) d^2x \\ &= (A f * \varepsilon^2 W_\varepsilon^-)(x_i - x_0), \end{aligned} \quad (\text{A2})$$

where x_i is the location of the center of pixel i , ε is the pixel side length, and we introduced the square top hat pixel density $W_\varepsilon(x)$, which equals $1/\varepsilon^2$ when x is in the range $-\varepsilon/2$ to $\varepsilon/2$ in both dimensions and equals 0 otherwise. In the last line, the symbol $*$ represents 2-d convolution and the minus sign superscript denotes the mirror image of a function, i.e. $W_\varepsilon^-(x) = W_\varepsilon(-x)$, though in this case $W_\varepsilon^- = W_\varepsilon$ since the top hat function is symmetric.

The optimal weight assigned to pixel i is $w_i = \tilde{s}_i / \sigma_i^2$ (Eq. 11), where $\tilde{s}_i$ is the predicted flux according to the model used for the matched filter (quantities with tildes are associated with the trial model used for the weights). Similar to s_i above,

$$w_i = \frac{1}{\sigma^2} \int_{\text{pixel } i} \tilde{A} \tilde{f}(x - \tilde{x}_0) d^2x, \quad (\text{A3})$$

$$\approx \frac{1}{\sigma^2} \tilde{A} \tilde{f}(x_i - \tilde{x}_0) \varepsilon^2, \quad (\text{A4})$$

where $\tilde{f}(x - \tilde{x}_0)$ is the normalized flux density of the trial object located at position $\tilde{x}_0$ in the image. In the second line we use a modified weight that slightly trades optimality for speed of computation by approximating the integral over pixel i . In practice these modified weights remain very close to optimal for the PSFs and pixel scales of typical imaging systems¹⁵. We also assume that the noise level is constant in the vicinity of the source so that there is a single value σ^2 for all relevant pixels in the image.

We can find closed form expressions for the inner pixel sums in Eq. A1 by taking averages with respect to the location of the source within a pixel, i.e. the true position of the source is equally likely to fall anywhere relative to the pixel boundaries. This averaging is highly apposite since when searching for unknown objects, source positions within the images are completely independent from the pixel alignment. Stacking many images on a target that moves over time is a further averaging of source position relative to the pixel boundaries.

To do the averaging, first fix the separation between the trial position $\tilde{x}_0$ and true position x_0 ,

$$\Delta x \equiv \tilde{x}_0 - x_0 = \text{constant}, \quad (\text{A5})$$

¹⁵ We numerically checked the reduction in expected significance using these suboptimal weights when f and $\tilde{f}$ are 2-d Gaussian densities. The loss is negligible when the pixel scale is small compared to the PSF width. For an extreme case where the PSF full width at half max (FWHM) is 0.9 pixels, SNR decreases by 3% when using weights from Eq. A4 instead of Eq. A3. For the more realistic case of a FWHM of 2 pixels the loss in SNR is less than 0.2%.

and then average over the subpixel location x_0 . Concretely, the expected value of the arbitrary function $F(x_0)$ with respect to the subpixel position of x_0 is

$$E_{x_0} [F(x_0)] = \int_{\text{pixel } 0} F(x_0) \frac{d^2 x_0}{\varepsilon^2}, \quad (\text{A6})$$

where the integral is over the “central pixel”, i.e. the integration limits for x_0 are from $-\varepsilon/2$ to $+\varepsilon/2$ in both dimensions. There is no loss of generality in considering the source to be located in the central pixel as opposed to an arbitrary pixel, or even randomly located over the entire plane.

From Eqs. A2 and A4, the numerator of Eq. A1 has the form $\sum_i g(x_i - x_0)h(x_i - \tilde{x}_0)$. Averaging over the subpixel location x_0 gives

$$\begin{aligned} E_{x_0} \left[\sum_i g(x_i - x_0)h(x_i - \tilde{x}_0) \right] &= \frac{1}{\varepsilon^2} \int_{\text{pixel } 0} d^2 x_0 \left(\sum_i g(x_i - x_0)h(x_i - x_0 - \Delta x) \right) \\ &= \frac{1}{\varepsilon^2} \sum_i \int_{\text{pixel } i} d^2 x' g(x')h(x' - \Delta x) \\ &= \frac{1}{\varepsilon^2} \int_{\mathbb{R}^2} d^2 x' g(x')h(x' - \Delta x) \\ &= \frac{1}{\varepsilon^2} (g * h^-)(\Delta x), \end{aligned} \quad (\text{A7})$$

where in the second line we exchanged the sum and integral and then changed variables to $x' = x_i - x_0$.

Using Eqs. A2 and A4 for g and h , we have

$$E_{x_0} \left[\sum_i w_i s_i \right] = \frac{A\tilde{A}}{\sigma^2/\varepsilon^2} (f * \tilde{f}^- * W_\varepsilon^-)(\Delta x). \quad (\text{A8})$$

Similarly, the sum in the denominator of Eq. A1 is

$$E_{x_0} \left[\sum_i w_i^2 \sigma^2 \right] = \frac{\tilde{A}^2}{\sigma^2/\varepsilon^2} (\tilde{f} * \tilde{f}^-)(0). \quad (\text{A9})$$

So far we have only calculated sums over the pixels within a single image. In Eq. A1 there is an additional outer sum over all the images which contain the source. The numerator and denominator in the expected significance will therefore consist of sums of many terms. We can invoke the law of large numbers to conclude that these outer sums will be close to their means, converging with increasing number of stacked images. In other words, we can use Eqs. A8 and A9 to compute the expected detection significance and introduce little error compared to what we would have obtained using Eq. A1 without averaging over subpixel location.

A.2. Gaussian intensity profile

Most solar system bodies will appear as point sources in images. In fact, the original motivation for digital tracking was to use short exposure times to prevent so-called trailing loss where a moving object’s flux is spread across multiple pixels in a streak.

The normalized flux density $f(x)$ in this case simply traces the telescope point spread function (PSF) ϕ_b , and we set $f(x) = \tilde{f}(x) = \phi_b(x)$. We model the PSF as a 2-d Gaussian with standard deviation b ,

$$\phi_b(x) = \frac{1}{2\pi b^2} \exp\left(-\frac{|x|^2}{2b^2}\right). \quad (\text{A10})$$

We do not consider streak profiles for $f(x)$ but it would be relatively straightforward to do so (perhaps at the expense of closed form expressions).

A.3. Expected significance for small pixels

Equations A8 and A9 involve the convolution of the Gaussian PSF ϕ_b with itself. This is simply another Gaussian with twice the variance: $\phi_b * \phi_b^- = \phi_{\sqrt{2}b}$. The doubled variance exactly corresponds to the standard practice of convolving an image with its PSF to create a significance map, which means a point source is convolved with the PSF twice, first by the atmosphere and then in software.

Convolution with the pixel density W_ε^- causes a further blur by the pixel size. If the pixels are much smaller than the PSF width ($\varepsilon \ll b$) this blurring is negligible compared to the width of $\phi_{\sqrt{2}b}$ and we can ignore the convolution with W_ε^- :

$$\mathbb{E}_{x_0} \left[\sum_i w_i s_i \right] = \frac{A\tilde{A}}{\sigma^2/\varepsilon^2} \phi_{\sqrt{2}b}(\Delta x), \quad (\text{A11})$$

$$\mathbb{E}_{x_0} \left[\sum_i w_i^2 \sigma^2 \right] = \frac{\tilde{A}^2}{\sigma^2/\varepsilon^2} \phi_{\sqrt{2}b}(0). \quad (\text{A12})$$

Plugging these into Eq. A1 and restoring the outer sums over all the images which contain the source we obtain

$$\text{SNR}(\Delta x) = \frac{\sum_i \frac{\tilde{A}_i A_i}{\sigma_i^2/\varepsilon_i^2} \phi_i(\Delta x^i)}{\sqrt{\sum_i \frac{\tilde{A}_i^2}{\sigma_i^2/\varepsilon_i^2} \phi_i(0)}}. \quad (\text{A13})$$

The index i now runs over all images so that b_i , ε_i , and σ_i^2 , are the PSF width, pixel scale, and noise variance for image i . The amplitudes $\tilde{A}_i$ and A_i are the fluxes of the trial and true source, and Δx^i is the separation between the trial and true source positions in image i . For clarity we use ϕ_i to stand for $\phi_{\sqrt{2}b_i}$.

Note that SNR is not affected by the overall scaling of $\tilde{A}_i$ (see the discussion after Eq. 11). What is important is the relative flux of the source between different images, i.e. the light curve. For our purposes, minor planet light curves are determined by the Sun-object-Earth geometry of their orbit. We will always consider trial orbits that are quite close to the true orbit. Therefore, the trial light curve is proportional to the true light curve ($\tilde{A}_i \propto A_i$), which we now use to further simplify Eq. A13.

The maximum detection significance occurs when the predicted position of the source coincides with its true position in all images ($\Delta x^i = 0$ for all i). We can factor out this maximum significance to rewrite Eq. A13 as

$$\text{SNR}(\Delta x) = \text{SNR}_{\max} \sum_i c_i \frac{\phi_i(\Delta x^i)}{\phi_i(0)} \quad (\text{A14})$$

$$= \text{SNR}_{\max} \sum_i c_i \exp\left(-\frac{|\Delta x^i|^2}{4b_i^2}\right), \quad (\text{A15})$$

where

$$\text{SNR}_{\max} = \sqrt{\sum_i \frac{A_i^2}{4\pi\sigma_i^2(b_i/\varepsilon_i)^2}}, \quad (\text{A16})$$

and the normalized image weights c_i are defined by

$$c_i = \frac{c'_i}{\sum_i c'_i} \quad \text{with} \quad c'_i = \frac{\tilde{A}_i^2}{\sigma_i^2(b_i/\varepsilon_i)^2}. \quad (\text{A17})$$

Equation A16 is well known in astronomy as the inverse of the relative uncertainty of source's flux in background-dominated imaging (e.g. I. R. King 1983). It is a convenient formula that can be used to find, for example, the completeness of a stacking search. One simply scales the A_i 's until $\text{SNR}_{\max}$ equals, say, 10 and then converts the scaling into a magnitude.

A.4. Expected significance for finite pixels

Most astronomical imaging is designed so that the PSF is comparable to the pixel scale. In this case, to retain full accuracy, we should include the convolution with W_ε^- in Eq. A8. The convolution of a Gaussian with the top hat pixel density is just the average of the Gaussian density over a pixel, which we call $\bar{\phi}_b(x)$,

$$\bar{\phi}_b(x) \equiv (\phi_b * W_\varepsilon^-)(x) = \frac{1}{\varepsilon^2} \int_{\text{pixel } 0} d^2x' \phi_b(x - x'), \quad (\text{A18})$$

where the integration is from $-\varepsilon/2$ to $\varepsilon/2$ in both the x and y directions (the dependence on pixel size ε is implicit in $\bar{\phi}_b$).

For finite size pixels Eq. A8 is

$$E_{x_0} \left[\sum_i w_i s_i \right] = \frac{A\tilde{A}}{\sigma^2/\varepsilon^2} \bar{\phi}_{\sqrt{2}b}(\Delta x), \quad (\text{A19})$$

which is just Eq. A11 averaged over a pixel. Equation A12 is unchanged.

We therefore obtain slight modifications of Eqs. A14, A16, and A17,

$$\text{SNR}(\Delta x) = \text{SNR}_{\max} \sum_i c_i \frac{\bar{\phi}_i(\Delta x^i)}{\bar{\phi}_i(0)}, \quad (\text{A20})$$

$$\text{SNR}_{\max} = \frac{\sum_i \frac{A_i^2 q_i^{(0)}}{4\pi\sigma_i^2 (b_i/\varepsilon_i)^2}}{\sqrt{\sum_i \frac{A_i^2}{4\pi\sigma_i^2 (b_i/\varepsilon_i)^2}}}, \quad (\text{A21})$$

$$c_i = \frac{c'_i}{\sum_i c'_i}, \quad \text{where} \quad c'_i = \frac{\tilde{A}_i^2 q_i^{(0)}}{\sigma_i^2 (b_i/\varepsilon_i)^2}. \quad (\text{A22})$$

In these equations $\bar{\phi}_i$ is shorthand for $\bar{\phi}_{\sqrt{2}b_i}$ and $q_i^{(0)}$ is shorthand for $q_{\varepsilon_i/\sqrt{2}b_i}^{(0)}$ (defined in Eq. B24). The latter is a pixel size correction factor that converges to 1 for small pixels, meaning that Eqs. A20, A21, and A22 reduce to Eqs. A14, A16, and A17 in that limit.

B. PROPERTIES OF $\bar{\phi}_b(x)$

In Sec. 3.3 we use the Taylor series of the average 2-d Gaussian density $\bar{\phi}_b(x)$ around $x = 0$. Since $\bar{\phi}_b(x)$ is an even function of x , the series only contains even powers of x . The lowest order terms are

$$\bar{\phi}_b(x) = \bar{\phi}_b(0) - \frac{1}{2} \sqrt{\bar{\phi}_b(0)} \frac{\exp\left(-\frac{(\varepsilon/2)^2}{2b^2}\right)}{\sqrt{2\pi b^2}} \left| \frac{x}{b} \right|^2 + \mathcal{O} \left| \frac{x}{b} \right|^4.$$

We introduce two pixel size correction factors $q_{\varepsilon/b}^{(0)}$ and $q_{\varepsilon/b}^{(2)}$ to isolate the effects of the finite pixel size. First,

$$\bar{\phi}_b(0) \equiv \phi_b(0) q_{\varepsilon/b}^{(0)}, \quad (\text{B23})$$

where $\phi_b(0) = 1/(2\pi b^2)$ is the $\varepsilon \rightarrow 0$ limit so that

$$q_{\varepsilon/b}^{(0)} = \left[\frac{\sqrt{2\pi}}{\varepsilon/b} \text{erf} \left(\frac{\varepsilon/b}{2\sqrt{2}} \right) \right]^2, \quad (\text{B24})$$

which converges to 1 when $\varepsilon \lesssim b$ (i.e. when the pixel size is negligible).

Next, we define

$$q_{\varepsilon/b}^{(2)} \equiv \frac{\exp\left(-\frac{1}{8}(\varepsilon/b)^2\right)}{\sqrt{q_{\varepsilon/b}^{(0)}}} = \frac{\exp\left(-\frac{1}{8}(\varepsilon/b)^2\right)}{\frac{\sqrt{2\pi}}{\varepsilon/b} \text{erf} \left(\frac{\varepsilon/b}{2\sqrt{2}} \right)}, \quad (\text{B25})$$




Figure 11. The pixel size correction factors $q_{\epsilon/b}^{(0)}$ (Eq. B24) and $q_{\epsilon/b}^{(2)}$ (Eq. B25) for a range of PSF sizes.

where both the numerator and denominator are ≈ 1 when $\epsilon \lesssim b$.

Using Eq. B25, the Taylor series for $\bar{\phi}_b(x)$ can be written

$$\frac{\bar{\phi}_b(x)}{\bar{\phi}_b(0)} = 1 - \frac{1}{2} q_{\epsilon/b}^{(2)} \left| \frac{x}{b} \right|^2 + \mathcal{O} \left| \frac{x}{b} \right|^4, \quad (\text{B26})$$

which is needed in Sec. 3.3.

Figure 11 shows the two pixel size correction factors as a function of b/ϵ and the PSF full width at half max (FWHM = $(2\sqrt{2\log 2})b$ for a Gaussian PSF). The correction factors are $\gtrsim 0.9$ when the FWHM $\gtrsim 2$ pixels but can be significantly less than one for narrower PSFs.

C. FEASIBILITY OF EIGHT-HOUR CONTINUOUS TRACKING WITH AIRMASS ≤ 2 AT CERRO PACHÓN

This appendix derives the condition under which a fixed field can be tracked from the ground for at least 8 hours while remaining at airmass ≤ 2 . Refraction and local horizon obstructions are neglected; using a more accurate airmass model only shifts the altitude threshold slightly and does not change the algebra below.

Altitude h is elevation above the horizon. The zenith angle is $z = 90^\circ - h$. The hour angle H , local sidereal time, and right ascension of the field α are related by $H = \text{LST} - \alpha$ (H is zero at upper transit and increases westward at 15° per hour). The declination of the field is δ . Assuming the airmass is $X \approx \sec z$, the requirement $X \leq 2$ gives $z \leq 60^\circ$, i.e. the elevation must be $h \geq h_0 \equiv 30^\circ$.

The standard spherical relation is

$$\sin h = \sin \phi \sin \delta + \cos \phi \cos \delta \cos H,$$

where ϕ is the latitude of the observatory. The airmass requirement is $\sin h \geq \sin h_0$, which implies

$$\cos H \geq C(\delta) \equiv \frac{\sin h_0 - \sin \phi \sin \delta}{\cos \phi \cos \delta}. \quad (\text{C27})$$

If $|C(\delta)| \leq 1$, the allowed hour angles are $|H| \leq H_0$ where $H_0 \equiv \arccos(C(\delta))$. The continuous observation period meeting the airmass requirement is therefore

$$t(\delta; h_0, \phi) = \frac{2H_0}{15^\circ/\text{hr}}.$$

Requiring $t \geq 8$ hours gives $H_0 \geq 60^\circ$, i.e. $\cos H_0 = C(\delta) \leq \frac{1}{2}$. Using the definition of $C(\delta)$ (Eq. C27), this yields a constraint on the declination of the field,

$$\sin \phi \sin \delta + \frac{1}{2} \cos \phi \cos \delta \geq \sin h_0.$$

Inserting the latitude of the Rubin Observatory at Cerro Pachón ($\phi = -30.24^\circ$) and solving for the allowed declination gives

$$\delta \leq -8.3^\circ.$$

This means that a field at $\delta = -8.3^\circ$ is visible at airmass less than 2 for 8 hours; fields further south can be viewed for longer.

If we are interested in observing along the ecliptic plane, we can solve for the viable ecliptic longitudes λ using $\sin(\delta(\lambda)) = \sin \varepsilon \sin \lambda$, where $\varepsilon = 23.44^\circ$ is Earth's obliquity. The result is

$$\lambda \in [201.2^\circ, 338.8^\circ].$$

REFERENCES

- Allen, R. L., Bernstein, G. M., & Malhotra, R. 2001, *ApJL*, 549, L241, doi: [10.1086/319165](https://doi.org/10.1086/319165)
- Alvarez-Candal, A., Souza-Feliciano, A. C., Martins-Filho, W., Pinilla-Alonso, N., & Ortiz, J. L. 2020, *MNRAS*, 497, 5473, doi: [10.1093/mnras/staa2329](https://doi.org/10.1093/mnras/staa2329)
- Amari, S.-i. 1985, *Lecture Notes in Statistics*, Vol. 28, Differential-Geometrical Methods in Statistics, 1st edn. (New York, NY: Springer), doi: [10.1007/978-1-4612-5056-2](https://doi.org/10.1007/978-1-4612-5056-2)
- ASTM International. 2019, *Standard Solar Constant and Zero Air Mass Solar Spectral Irradiance Tables*, ASTM International, doi: [10.1520/E0490-00AR19](https://doi.org/10.1520/E0490-00AR19)
- Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, *A&A*, 558, A33, doi: [10.1051/0004-6361/201322068](https://doi.org/10.1051/0004-6361/201322068)
- Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, *AJ*, 156, 123, doi: [10.3847/1538-3881/aabc4f](https://doi.org/10.3847/1538-3881/aabc4f)
- Astropy Collaboration, Price-Whelan, A. M., Lim, P. L., et al. 2022, *ApJ*, 935, 167, doi: [10.3847/1538-4357/ac7c74](https://doi.org/10.3847/1538-4357/ac7c74)
- Batygin, K., Adams, F. C., Brown, M. E., & Becker, J. C. 2019, *PhR*, 805, 1, doi: [10.1016/j.physrep.2019.01.009](https://doi.org/10.1016/j.physrep.2019.01.009)
- Batygin, K., & Brown, M. E. 2016, *AJ*, 151, 22, doi: [10.3847/0004-6256/151/2/22](https://doi.org/10.3847/0004-6256/151/2/22)
- Bellm, E. C., Kulkarni, S. R., Graham, M. J., et al. 2019a, *PASP*, 131, 018002, doi: [10.1088/1538-3873/aaecbe](https://doi.org/10.1088/1538-3873/aaecbe)
- Bellm, E. C., Kulkarni, S. R., Barlow, T., et al. 2019b, *PASP*, 131, 068003, doi: [10.1088/1538-3873/ab0c2a](https://doi.org/10.1088/1538-3873/ab0c2a)
- Belyakov, M., Bernardinelli, P. H., & Brown, M. E. 2022, *AJ*, 163, 216, doi: [10.3847/1538-3881/ac5c56](https://doi.org/10.3847/1538-3881/ac5c56)
- Bernstein, G., & Khushalani, B. 2000, *AJ*, 120, 3323, doi: [10.1086/316868](https://doi.org/10.1086/316868)
- Bernstein, G. M., Trilling, D. E., Allen, R. L., et al. 2004, *AJ*, 128, 1364, doi: [10.1086/422919](https://doi.org/10.1086/422919)
- Bianco, F. B., Ivezić, Ž., Jones, R. L., et al. 2022, *ApJS*, 258, 1, doi: [10.3847/1538-4365/ac3e72](https://doi.org/10.3847/1538-4365/ac3e72)
- Bowell, E., Hapke, B., Domingue, D., et al. 1989, in *Asteroids II*, ed. R. P. Binzel, T. Gehrels, & M. S. Matthews, 524–556
- Brewin, L. 1996, *Riemann Normal Coordinates*, Department of Mathematics Preprint, Monash University, Clayton, Vic. 3168. <https://users.monash.edu.au/~leo/research/papers/files/lcb96-01.pdf>
- Brodtkorb, P. A. 2025, *Numdifftools: Numerical differentiation tools*, 0.9.41, <https://github.com/pbrod/numdifftools>
- Brown, M. E., & Batygin, K. 2021, *AJ*, 162, 219, doi: [10.3847/1538-3881/ac2056](https://doi.org/10.3847/1538-3881/ac2056)
- Brown, M. E., & Batygin, K. 2022, *AJ*, 163, 102, doi: [10.3847/1538-3881/ac32dd](https://doi.org/10.3847/1538-3881/ac32dd)
- Brown, M. E., Holman, M. J., & Batygin, K. 2024, *AJ*, 167, 146, doi: [10.3847/1538-3881/ad24e9](https://doi.org/10.3847/1538-3881/ad24e9)
- Burdanov, A. Y., Hasler, S. N., & de Wit, J. 2023, *MNRAS*, 521, 4568, doi: [10.1093/mnras/stad808](https://doi.org/10.1093/mnras/stad808)
- Calabretta, M. R., & Greisen, E. W. 2002, *A&A*, 395, 1077, doi: [10.1051/0004-6361:20021327](https://doi.org/10.1051/0004-6361:20021327)
- Chiang, E. I., & Brown, M. E. 1999, *AJ*, 118, 1411, doi: [10.1086/301005](https://doi.org/10.1086/301005)
- Cochran, A. L., Levison, H. F., Stern, S. A., & Duncan, M. J. 1995, *ApJ*, 455, 342, doi: [10.1086/176581](https://doi.org/10.1086/176581)
- D'Errico, J. 2006, *Adaptive Robust Numerical Differentiation*, <https://www.mathworks.com/matlabcentral/fileexchange/13490-adaptive-robust-numerical-differentiation>
- Dhurandhar, S. V., & Sathyaprakash, B. S. 1994, *PhRvD*, 49, 1707, doi: [10.1103/PhysRevD.49.1707](https://doi.org/10.1103/PhysRevD.49.1707)
- Fraser, W. C., Brown, M. E., Morbidelli, A., Parker, A., & Batygin, K. 2014, *ApJ*, 782, 100, doi: [10.1088/0004-637X/782/2/100](https://doi.org/10.1088/0004-637X/782/2/100)
- Fraser, W. C., Kavelaars, J. J., Holman, M. J., et al. 2008, *Icarus*, 195, 827, doi: [10.1016/j.icarus.2008.01.014](https://doi.org/10.1016/j.icarus.2008.01.014)
- Fraser, W. C., Porter, S. B., Peltier, L., et al. 2024, *PSJ*, 5, 227, doi: [10.3847/PSJ/ad6f9e](https://doi.org/10.3847/PSJ/ad6f9e)

- Fuentes, C. I., George, M. R., & Holman, M. J. 2009, *ApJ*, 696, 91, doi: [10.1088/0004-637X/696/1/91](https://doi.org/10.1088/0004-637X/696/1/91)
- Geringer-Sameth, A., & Koushiappas, S. M. 2012, *MNRAS*, 425, 862, doi: [10.1111/j.1365-2966.2012.21139.x](https://doi.org/10.1111/j.1365-2966.2012.21139.x)
- Geringer-Sameth, A., Koushiappas, S. M., & Walker, M. G. 2015, *PhRvD*, 91, 083535, doi: [10.1103/PhysRevD.91.083535](https://doi.org/10.1103/PhysRevD.91.083535)
- Ginsburg, A., Sipőcz, B. M., Brasseur, C. E., et al. 2019, *AJ*, 157, 98, doi: [10.3847/1538-3881/aafc33](https://doi.org/10.3847/1538-3881/aafc33)
- Giorgini, J. D., & JPL Solar System Dynamics Group. , NASA/JPL Horizons On-Line Ephemeris System, <https://ssd.jpl.nasa.gov/horizons>
- Giorgini, J. D., Yeomans, D. K., Chamberlin, A. B., et al. 1996, in *AAS/Division for Planetary Sciences Meeting Abstracts*, Vol. 28, AAS/Division for Planetary Sciences Meeting Abstracts #28, 25.04
- Gladman, B., Kavelaars, J. J., Nicholson, P. D., Lored, T. J., & Burns, J. A. 1998, *AJ*, 116, 2042, doi: [10.1086/300573](https://doi.org/10.1086/300573)
- Golovich, N., Steil, T., Geringer-Sameth, A., et al. 2025, *Astronomy and Computing*, 53, 100987, doi: [10.1016/j.ascom.2025.100987](https://doi.org/10.1016/j.ascom.2025.100987)
- Granvik, M., Morbidelli, A., Jedicke, R., et al. 2018, *Icarus*, 312, 181, doi: [10.1016/j.icarus.2018.04.018](https://doi.org/10.1016/j.icarus.2018.04.018)
- Gural, P. S., Otto, P. R., & Tedesco, E. F. 2018, *PASP*, 130, 074504, doi: [10.1088/1538-3873/aac1ff](https://doi.org/10.1088/1538-3873/aac1ff)
- Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, *Nature*, 585, 357, doi: [10.1038/s41586-020-2649-2](https://doi.org/10.1038/s41586-020-2649-2)
- Heinze, A. N., Metchev, S., & Trollo, J. 2015, *AJ*, 150, 125, doi: [10.1088/0004-6256/150/4/125](https://doi.org/10.1088/0004-6256/150/4/125)
- Heinze, A. N., Trollo, J., & Metchev, S. 2019, *AJ*, 158, 232, doi: [10.3847/1538-3881/ab48fa](https://doi.org/10.3847/1538-3881/ab48fa)
- Holman, M. J., Kavelaars, J. J., Grav, T., et al. 2004, *Nature*, 430, 865, doi: [10.1038/nature02832](https://doi.org/10.1038/nature02832)
- Hunter, J. D. 2007, *Computing in Science & Engineering*, 9, 90, doi: [10.1109/MCSE.2007.55](https://doi.org/10.1109/MCSE.2007.55)
- IRSA. 2022, Zwicky Transient Facility Image Service, IPAC, doi: [10.26131/IRSA539](https://doi.org/10.26131/IRSA539)
- Ivezić, Ž., Kahn, S. M., Tyson, J. A., et al. 2019, *ApJ*, 873, 111, doi: [10.3847/1538-4357/ab042c](https://doi.org/10.3847/1538-4357/ab042c)
- Jones, R. L., Slater, C. T., Moeyens, J., et al. 2018, *Icarus*, 303, 181, doi: [10.1016/j.icarus.2017.11.033](https://doi.org/10.1016/j.icarus.2017.11.033)
- Kaiser, N., Aussel, H., Burke, B. E., et al. 2002, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 4836, *Survey and Other Telescope Technologies and Discoveries*, ed. J. A. Tyson & S. Wolff, 154–164, doi: [10.1117/12.457365](https://doi.org/10.1117/12.457365)
- King, I. R. 1983, *PASP*, 95, 163, doi: [10.1086/131139](https://doi.org/10.1086/131139)
- Lang, D., Hogg, D. W., Jester, S., & Rix, H.-W. 2009, *AJ*, 137, 4400, doi: [10.1088/0004-6256/137/5/4400](https://doi.org/10.1088/0004-6256/137/5/4400)
- Li, S., et al. 2025, In preparation
- Lifset, N., Golovich, N., Green, E., Armstrong, R., & Yeager, T. 2021, *AJ*, 161, 282, doi: [10.3847/1538-3881/abf7af](https://doi.org/10.3847/1538-3881/abf7af)
- Masci, F. J., Laher, R. R., Rusholme, B., et al. 2019, *PASP*, 131, 018003, doi: [10.1088/1538-3873/aae8ac](https://doi.org/10.1088/1538-3873/aae8ac)
- McGill, P., et al. 2025, In preparation
- Meyers, J. E., Schneider, M. D., Ebert, J. T., et al. 2025, *Journal of Open Source Software*, 10, 8147, doi: [10.21105/joss.08147](https://doi.org/10.21105/joss.08147)
- Milani, A. 1999, *Icarus*, 137, 269, doi: [10.1006/icar.1999.6045](https://doi.org/10.1006/icar.1999.6045)
- Milani, A., Sansaturio, M. E., Tommei, G., Arratia, O., & Chesley, S. R. 2005, *A&A*, 431, 729, doi: [10.1051/0004-6361:20041737](https://doi.org/10.1051/0004-6361:20041737)
- Muironen, K. 1996, *MNRAS*, 280, 1235, doi: [10.1093/mnras/280.4.1235](https://doi.org/10.1093/mnras/280.4.1235)
- Muironen, K., & Bowell, E. 1993, *Icarus*, 104, 255, doi: [10.1006/icar.1993.1100](https://doi.org/10.1006/icar.1993.1100)
- Muironen, K., Virtanen, J., Granvik, M., & Laakso, T. 2006, *MNRAS*, 368, 809, doi: [10.1111/j.1365-2966.2006.10168.x](https://doi.org/10.1111/j.1365-2966.2006.10168.x)
- Owen, B. J. 1996, *PhRvD*, 53, 6749, doi: [10.1103/PhysRevD.53.6749](https://doi.org/10.1103/PhysRevD.53.6749)
- Parker, A. H., & Kavelaars, J. J. 2010, *PASP*, 122, 549, doi: [10.1086/652424](https://doi.org/10.1086/652424)
- Press, W. H., Teukolsky, S. A., Vetterling, W. T., & Flannery, B. P. 2007, *Numerical Recipes 3rd Edition: The Art of Scientific Computing*, 3rd edn. (Cambridge University Press)
- Rabinowitz, D. L., Schaefer, B. E., & Tourtellotte, S. W. 2007, *AJ*, 133, 26, doi: [10.1086/508931](https://doi.org/10.1086/508931)
- Rice, M., & Laughlin, G. 2020, *PSJ*, 1, 81, doi: [10.3847/PSJ/abc42c](https://doi.org/10.3847/PSJ/abc42c)
- Rybicki, G. B., & Press, W. H. 1992, *ApJ*, 398, 169, doi: [10.1086/171845](https://doi.org/10.1086/171845)
- Sathyaprakash, B. S., & Dhurandhar, S. V. 1991, *PhRvD*, 44, 3819, doi: [10.1103/PhysRevD.44.3819](https://doi.org/10.1103/PhysRevD.44.3819)
- Shankman, C., Kavelaars, J. J., Bannister, M. T., et al. 2017, *AJ*, 154, 50, doi: [10.3847/1538-3881/aa7aed](https://doi.org/10.3847/1538-3881/aa7aed)
- Shao, M., Nemati, B., Zhai, C., et al. 2014, *ApJ*, 782, 1, doi: [10.1088/0004-637X/782/1/1](https://doi.org/10.1088/0004-637X/782/1/1)
- Siraj, A., Chyba, C. F., & Tremaine, S. 2025, *ApJ*, 978, 139, doi: [10.3847/1538-4357/ad98f6](https://doi.org/10.3847/1538-4357/ad98f6)
- Smith, K. M. 2016, arXiv e-prints, arXiv:1610.06831, doi: [10.48550/arXiv.1610.06831](https://doi.org/10.48550/arXiv.1610.06831)
- Smotherman, H., Bernardinelli, P. H., Portillo, S. K. N., et al. 2024, *AJ*, 167, 136, doi: [10.3847/1538-3881/ad1524](https://doi.org/10.3847/1538-3881/ad1524)
- The pandas development team. 2020, pandas-dev/pandas: Pandas, 2.2.1 Zenodo, doi: [10.5281/zenodo.3509134](https://doi.org/10.5281/zenodo.3509134)

- Tonry, J. L., Stubbs, C. W., Lykke, K. R., et al. 2012, ApJ, 750, 99, doi: [10.1088/0004-637X/750/2/99](https://doi.org/10.1088/0004-637X/750/2/99)
- Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, Nature Methods, 17, 261, doi: [10.1038/s41592-019-0686-2](https://doi.org/10.1038/s41592-019-0686-2)
- Wang, B., Zhao, H.-b., & Li, B. 2018, ChA&A, 42, 433, doi: [10.1016/j.chinastron.2018.09.007](https://doi.org/10.1016/j.chinastron.2018.09.007)
- Weinberg, S. 1972, Gravitation and Cosmology: Principles and Applications of the General Theory of Relativity (John Wiley & Sons, Inc.)
- Wes McKinney. 2010, in Proceedings of the 9th Python in Science Conference, ed. Stéfan van der Walt & Jarrod Millman, 56 – 61, doi: [10.25080/Majora-92bf1922-00a](https://doi.org/10.25080/Majora-92bf1922-00a)
- Whidden, P. J., Bryce Kalmbach, J., Connolly, A. J., et al. 2019, AJ, 157, 119, doi: [10.3847/1538-3881/aafd2d](https://doi.org/10.3847/1538-3881/aafd2d)
- Zhai, C., Shao, M., Nemati, B., et al. 2014, ApJ, 792, 60, doi: [10.1088/0004-637X/792/1/60](https://doi.org/10.1088/0004-637X/792/1/60)
- Zhai, C., Shao, M., Saini, N. S., et al. 2018, AJ, 156, 65, doi: [10.3847/1538-3881/aacb28](https://doi.org/10.3847/1538-3881/aacb28)
- Zhai, C., Ye, Q., Shao, M., et al. 2020, PASP, 132, 064502, doi: [10.1088/1538-3873/ab828b](https://doi.org/10.1088/1538-3873/ab828b)