AFTERGLOW

NASA SPACE APPS CHALLENGE

TECHNICAL REPORT AND METHODOLOGICAL FRAMEWORK

Afterglow: A Calibrated Search for Change in the SPHEREx Infrared Sky

Two epoch spectrophotometry of 2,310 stars, proper motions verified against Gaia DR3, an evaluated and rejected AI vetter, and the requirements of a SPHEREx search for a distant Planet X

Team Afterglow
NASA Space Apps Challenge 2026
Challenge: Planet X and SPHEREx

November 2026
Data through October 2, 2026
Costa Rica

Abstract

NASA's SPHEREx observatory surveys the full sky every six months in 102 spectral channels between 0.75 and 5.0 microns. Repeated coverage makes SPHEREx a time domain instrument, and the Planet X and SPHEREx challenge asks for a public tool that shows how the infrared sky changes. We present Afterglow, a calibrated pipeline and web tool built on SPHEREx Level 2 images from Data Release 1 and Quick Release 3.

In a one degree field near the north ecliptic pole we measured 2,310 AllWISE stars by forced aperture photometry in 168 exposures, for a total of 345,952 measurements over two epochs separated by about sixteen months. After a channel by channel cross calibration between epochs, the rejection of 808 stars with contaminating neighbors, and a threshold of 5 sigma in at least three channels, one of 1,500 compared stars survived. It coincides within 0.12 arcsec with ZTF J175232.01+653228.6, an eclipsing binary already catalogued. The pipeline recovered it without prior knowledge of its position.

For any position on the sky, Afterglow stacks each survey pass per detector and measures the displacement of every source. The measured motion of Barnard's Star is 10.24 arcsec between passes, against 10.27 arcsec predicted by Gaia DR3. Across the five test objects of the challenge, the method flags as moving exactly the two objects whose Gaia proper motion predicts a measurable shift, and none of the other twelve sources in those fields. On 47 blind test positions, an untuned version of the measurement classified 43 correctly, while a general purpose AI vision model classified 37. We therefore did not ship the AI component. The tool also extracts sources present in only one pass, the signature of a body that moved beyond the matching radius, and on 47 real fields this extraction raised no false detections. Visitors can mark a source in several frames, and the tool fits its motion.

We close with a methodological framework for a SPHEREx search for a distant planet. A body at 400 to 800 AU would move 57 to 162 arcsec per year between same season passes, beyond the 25 arcsec search window of the present method. A dedicated search therefore requires single pass source extraction over large areas and orbit based linking. The extraction now exists in Afterglow for fields of 1 to 10 arcmin; the linking does not. We describe which components already provide that foundation and which do not. Finally, we injected 5,040 fake planets, displaced by the parallax computed from the Earth's position on the real dates of the passes, into the stacks of nine real fields: a single pass recovers half of them down to AB 17.4 to 17.8 in quiet fields, more than four magnitudes brighter than the $22 < V < 25$ predicted for Planet Nine.

Keywords: SPHEREx; infrared spectrophotometry; time domain astronomy; proper motion; Gaia DR3; difference imaging; Planet Nine; citizen science; reproducibility.

Contents

Abstract 2

1. Introduction 4

    1.1 The Planet X question 4

    1.2 SPHEREx as a time domain instrument 4

    1.3 Objectives and scope 4

2. Research Questions and Hypotheses 5

3. Data 5

4. Methods 6

    4.1 Field and epochs 7

    4.2 Forced photometry 7

    4.3 Spectral binning and inter epoch calibration 7

    4.4 Isolation, detection and classification 7

    4.5 Motion measurement for any position 8

    4.6 Known small bodies 9

    4.7 Single pass sources 9

    4.8 Citizen reports 10

5. Evaluation of an AI Vetter 10

6. Results 13

    6.1 Detection funnel and calibration 13

    6.2 The surviving candidate 14

    6.3 Proper motions against Gaia DR3 15

    6.4 Known asteroids 16

    6.5 Single pass sources on real fields 17

7. Framework for a SPHEREx Planet X Search 17

8. Limitations 22

9. Software, Reproducibility and Public Tool 23

10. Conclusions and Future Work 24

Acknowledgments and Use of AI 24

References 26

Appendix A. Pipeline Parameters 27

Appendix B. Blind Test Protocol 28

1. Introduction

1.1 The Planet X question

The possibility of an undiscovered planet beyond Neptune is one of the oldest open questions of solar system astronomy. Its modern form, often called Planet Nine, rests on the observed clustering of the orbits of distant trans Neptunian objects, which Batygin and Brown (2016) attributed to the gravitational influence of a planet of several Earth masses on a distant, eccentric orbit. A later review places the planet at a semimajor axis of roughly 400 to 800 AU (Batygin et al. 2019). At such distances a planet is faint, and it moves slowly against the background stars, so a search requires repeated, wide area, sensitive imaging.

Searches in the infrared have already been carried out with WISE and NEOWISE, which together observed the full sky at 3.4 and 4.6 microns over several years (Wright et al. 2010; Meisner et al. 2018). The citizen science project Backyard Worlds: Planet 9 showed that volunteers who flip between epochs of WISE images can find nearby brown dwarfs that automated pipelines missed (Kuchner et al. 2017). These efforts share one principle: an object that changes position or brightness between epochs reveals itself in a comparison of images taken at different times.

1.2 SPHEREx as a time domain instrument

SPHEREx, the Spectro Photometer for the History of the Universe, Epoch of Reionization and Ices Explorer, observes the full sky every six months in 102 spectral channels between 0.75 and 5.0 microns (Doré et al. 2014). Each Level 2 image covers a field of about 3.5 by 3.5 degrees with pixels of 6.15 arcsec, and a linear variable filter makes the wavelength change across the detector. A source is therefore observed at a different wavelength in each exposure, and its spectrum is assembled from many exposures taken on different dates. Data Release 1 and Quick Release 3, distributed by NASA/IPAC IRSA, already contain two or more passes over most of the sky.

This structure gives SPHEREx two properties that matter for the Planet X question: repeated coverage of the same position, and a spectrum for every source. It also introduces two difficulties. Exposures of the same position are rotated and differ in wavelength, so they cannot be compared directly, and the two releases differ slightly in calibration.

1.3 Objectives and scope

The challenge asks for a public facing web tool that makes it easy for anyone to see how the SPHEREx sky changes over time. We set four objectives:

1. Build a measurement that compares two SPHEREx epochs with a false positive rate low enough that every reported change can be defended.

2. Verify that measurement against independent catalogs, rather than by visual inspection alone.
3. Expose the same measurement in a public tool for any position on the sky.
4. State quantitatively what the method can and cannot detect, including what a dedicated Planet X search would require.

This report describes the data, the methods, the evaluation of an AI component, the results and the limitations. All numbers come from our own runs on the public data.

2. Research Questions and Hypotheses

We formulated three questions before analyzing the results, each with a testable hypothesis and a criterion for rejecting it.

Table 1. Research questions, hypotheses and the outcome reported in this document.

Question	Hypothesis and test	Outcome
Q1. Can two SPHEREx epochs from different releases be compared with a negligible rate of false changes?	H1. After channel by channel calibration and neighbor rejection, the number of stars passing a 5 sigma, three channel threshold is of order unity, and each survivor corresponds to a physical change. Rejected if the survivors cannot be explained.	Supported. One survivor among 1,500 stars, coincident with a catalogued eclipsing binary (Section 6.2).
Q2. Do proper motions measured on SPHEREx stacks agree with Gaia DR3?	H2. Measured displacements agree with Gaia predictions to within one SPHEREx pixel, and no static source is flagged. Rejected if any target disagrees by more than one pixel.	Supported for the detectable targets. Barnard's Star 10.24 against 10.27 arcsec; Epsilon Indi B 6.2 against 6.4 arcsec (Section 6.3).
Q3. Does an AI vision model improve on the photometric measurement as a vetter of moving objects?	H3. On blind cases with known answers, the AI model achieves at least the accuracy of the measurement with fewer false alarms. Rejected otherwise.	Rejected. 37 of 47 against 43 of 47, with nine false alarms against three (Section 5).

3. Data

All data are public. Table 2 lists each product and its role. SPHEREx images were accessed through the IRSA Simple Image Access service, the IRSA cutout service and the public cloud bucket of the mission.

Table 2. Data products used in this work.

Product	Provider	Role
SPHEREx Level 2 spectral images, Data Release 1	NASA/IPAC IRSA	First epoch images; photometry, stacks and cutouts
SPHEREx Quick Release 3	NASA/IPAC IRSA	Most recent passes, through October 2, 2026 at our test positions
SPHEREx All Sky HiPS cube	NASA/IPAC IRSA	Visualization of the 102 channels in the public atlas
2MASS, DSS2 and AllWISE HiPS surveys	CDS, mirrored at NASA/IPAC IRSA	Background images of the atlas; reference epochs for the drawn motion paths
AllWISE Source Catalog (Cutri et al. 2013)	NASA/IPAC IRSA; VizieR II/328	Positions and W1 magnitudes of the 2,310 stars
Gaia DR3 (Gaia Collaboration 2023)	ESA; VizieR I/355	Independent proper motions for verification
SIMBAD (Wenger et al. 2000)	CDS Strasbourg	Known objects near candidates; name resolution
Zwicky Transient Facility (Bellm et al. 2019)	Through SIMBAD	Identification of the surviving candidate
Small Body Identification API	NASA JPL Solar System Dynamics	Known asteroids in each frame

Each SPHEREx Level 2 file contains the image in MJy/sr, a variance plane, a bit mask of quality flags, a model of the zodiacal light and the spectral world coordinate system that gives the wavelength at every pixel.

4. Methods

Figure 1 summarizes the change detection pipeline. The parameters are listed in Appendix A. Sections 4.1 to 4.4 describe the measurement on a fixed field; Sections 4.5 and 4.6 describe the measurements available for any position.



Figure 1. Change detection pipeline, with the number of stars that remain after each step in the NEP_off field.

4.1 Field and epochs

The study field, which we call NEP_off, is a circle of radius 0.5 degrees centered at right ascension 268.0 and declination 65.5 degrees, near the north ecliptic pole, where SPHEREx coverage is dense. Because coverage near the pole is nearly continuous, gaps between exposures do not separate natural epochs. We therefore defined each epoch as a 30 day window at the two ends of the available data: epoch A from April 25 to May 25, 2025, with 102 exposures, and epoch B from September 4 to October 2, 2026, with 66 exposures. This maximizes the baseline while keeping each epoch compact.

4.2 Forced photometry

We selected the 2,310 AllWISE sources in the field with W1 brighter than 15.5 and measured each of them in every exposure at its catalog position. The flux is the sum inside a circular aperture of 2 pixels minus the median of an annulus between 4 and 7 pixels, scaled to the aperture area. The uncertainty comes from the variance plane. A measurement is flagged when any pixel of the aperture carries a quality flag; the flag that marks pixels on known sources is excluded from this test, because every star would otherwise be flagged. The wavelength of each measurement is read from the spectral coordinate system of the image at the position of the star.

4.3 Spectral binning and inter epoch calibration

Measurements without flags are combined by inverse variance weighting into 102 channels of equal width between 0.75 and 5.0 microns, separately for each epoch. The two epochs come from different data releases, so we calibrate them against each other using the field as a reference. For each channel with at least 20 stars of signal to noise ratio above 10 in both epochs, the scale s is the median of the ratio between epoch B and epoch A. A systematic floor f is the part of the fractional scatter of stars with signal to noise ratio above 50 that the statistical uncertainties do not explain. The change of a star in a channel is then expressed as

$$z = (F_B / s - F_A) / [\sigma_A{}^2 + (\sigma_B / s)^2 + (f F_A)^2]^{1/2} / n,$$

where F and σ are the binned flux and its uncertainty, and n, never below 1, is the robust standard deviation of z over all stars in the channel. With this normalization a static star has z distributed close to a unit normal.

4.4 Isolation, detection and classification

A star is compared only if no other catalog source lies within 15 arcsec, and no source at least one magnitude brighter lies within 45 arcsec. A neighbor inside the aperture or the background annulus changes the measured flux when the exposures differ in orientation and wavelength, and this was the dominant source of false changes in early runs.

A star is a candidate when at least three channels reach |z| of 5 or more. Its score is the square root of the sum of z squared over those channels. Candidates are classified in order:

1. **Discarded** when more than half of its measurements are flagged.
2. **Appeared or vanished** when its signal to noise ratio is above 5 in one epoch and below 2 in the other, the signature of a moving object.
3. **Brightness change** when the significant channels span at least 30 percent of the spectral range and at least 80 percent of them change in the same direction.
4. **Unexplained change** otherwise.

Every candidate is then matched against SIMBAD within 6 arcsec, about one SPHEREx pixel.

4.5 Motion measurement for any position

For a position chosen by the user, Afterglow queries the IRSA archive for every SPHEREx image that contains it, groups the images into survey passes, and builds for each pass and each of the six detectors an average of up to eight cutouts. Each cutout is first corrected for zodiacal light and local background, its flagged pixels are masked, and it is resampled onto a common grid with north up and a pixel of 6.15 arcsec.

In the stack of each detector, a source is a connected region above 5 times the robust noise, with a flux weighted centroid. Each source of the first pass is matched to the nearest source of the last pass within 4 pixels and with a flux ratio between one third and three. The displacement of a source is the median of its displacements over the detectors in which it is detected, and a source must be detected in at least two detectors when two or more are available. A displacement of at least 3 arcsec, half a pixel, is reported as motion.

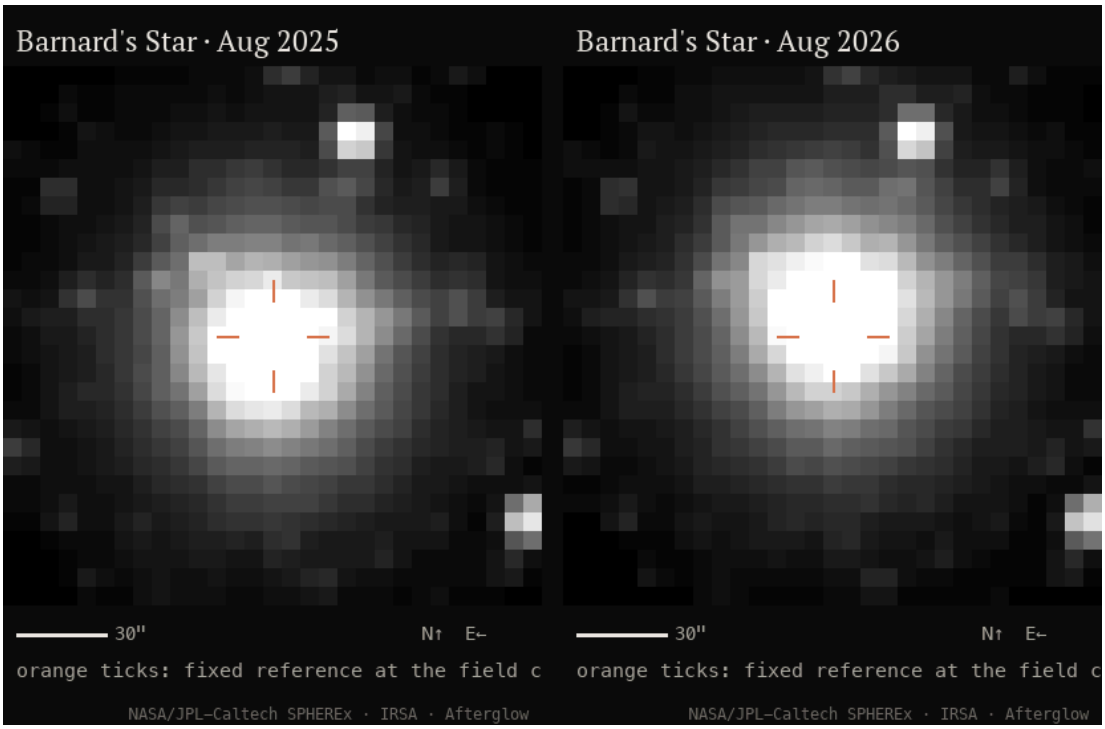


Figure 2. Stacks of Barnard's Star in the first and last pass (detector 2, 1.1 to 1.6 microns), resampled north up on the common grid and shown with the same intensity scale. The orange ticks mark a fixed reference at the field center.

Each measurement is compared with Gaia DR3. Because Gaia positions refer to epoch 2016.0, every Gaia star within the field and a margin of 150 arcsec is propagated with its proper motion to the date of the first pass before matching within 6 arcsec. The expected displacement is the Gaia proper motion multiplied by the interval between passes.

4.6 Known small bodies

For the frames displayed by the tool, Afterglow queries the NASA JPL Small Body Identification API once per window of 12 hours of exposures, with a geocentric observer, and extrapolates each returned position linearly to the time of each exposure using the rates provided by the service. The service is called in its first pass mode, which returns positions with uncertainties of about 20 arcsec. Bodies with an uncertainty above 120 arcsec are ignored. Each query takes one to two minutes, so results are computed in the background and cached.

4.7 Single pass sources

A body that moves farther than the matching radius of Section 4.5, or a source that appears or vanishes, leaves a source in one pass with no counterpart in the other. For any position, Afterglow extracts the sources of each detector stack and keeps a source of one pass as a single pass source when no source of the other pass lies within 2 pixels, its peak exceeds 7 times the robust noise, and the other pass has valid pixels at that position with no local maximum

above 3 times the noise. The last condition prevents a source just below the detection threshold, or the wing of a neighbor, from being taken as an absence. Saturated cores are skipped, because their shape changes between passes.

Detections of the same source in different detectors are merged within 1.5 pixels, and a single pass source must be detected in at least two detectors when two or more are available. A moving body is detected over a broad range of wavelengths, while instrumental artifacts usually affect a single detector. Each single pass source is then matched against Gaia DR3, propagated to the date of the first pass as in Section 4.5, and against the known bodies of Section 4.6.

4.8 Citizen reports

Following the practice of Backyard Worlds: Planet 9 (Kuchner et al. 2017), a visitor can report a source by marking it in one or more frames of the time lapse, with an optional note of at most 280 characters. When two to twelve frames are marked, each mark is converted to an offset on the sky, and the motion is fitted with a straight line in each axis against the date of each frame, which gives a displacement, a position angle and a rate in arcsec per year. In the automated tests, three marks placed on Barnard's Star recover a displacement of 10.2 arcsec to the north. Reports are public and require no account. They are limited to ten per connection in ten minutes, notes are stored as plain text, and the address of the connection is stored only as a salted hash.

5. Evaluation of an AI Vetter

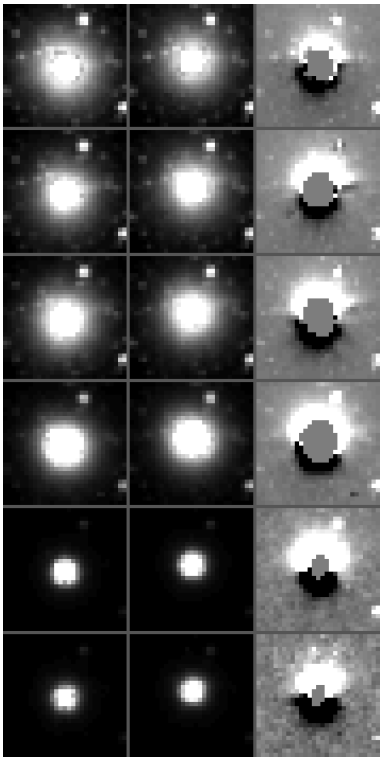
Volunteer projects detect moving objects by eye, by flipping between epochs. We asked whether a general purpose vision model could play that role on SPHEREx data and add value beyond our photometric measurement. The question matters because an AI component would change what the public tool reports, and its errors would reach users.

Design

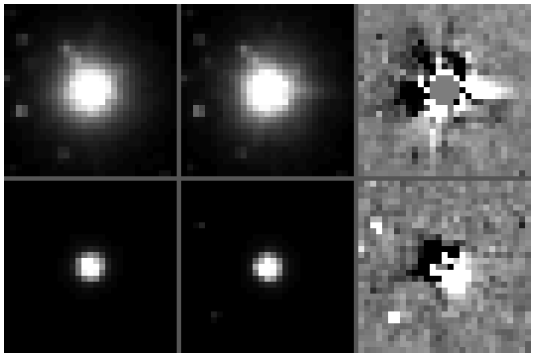
We built 47 test positions with a known answer. Twenty one are stars with large catalogued proper motions retrieved from SIMBAD, 25 are isolated static stars from the NEP_off field, and some are bright static stars that leave residuals in difference images, a realistic source of false alarms. For each position we stacked four exposures per detector in the first and last pass and rendered a single image with one row per detector and three columns: epoch A, epoch B and the difference. The ground truth was fixed before any evaluation: a position counts as moving when its catalogued proper motion multiplied by the interval between passes reaches 3 arcsec. Eight positions met that condition.

Blind protocol

The images were copied to a separate folder under random file names. The model saw only the images and a fixed instruction, without names, coordinates, catalog data or the interval between passes, because the interval alone would have identified the static field. It answered moving, static or unclear, with a confidence and a short justification. The comparison baseline was a measurement written before the evaluation: the median shift of the flux weighted centroid of the central source between the two stacks, with the same 3 arcsec threshold.



(a) Barnard's Star, moving 10.3 arcsec



(b) Static star of W1 5.3, called moving by the AI model

Figure 3. Two of the 47 blind cases exactly as the evaluator saw them. Rows are detectors; columns are the first pass, the last pass and their difference. In (a) the paired signature of motion is consistent across the six detectors. In (b) the residual ring of a saturated static star produces a similar pattern, the main source of false alarms of the AI model. The static field has two detectors with coverage.

Results

Table 3. Classification of the 47 blind cases, of which eight are moving.

Method	Correct	Movers detected	Movers missed	False alarms
Centroid measurement, threshold 3 arcsec	43 of 47	7	1	3
AI vision model, blind	37 of 47	7	1	9
Measurement confirmed by the AI model	44 of 47	7	1	2
Measurement with threshold 3.5 arcsec, no AI	44 of 47	7	1	2

Both methods missed WISE 0855, the faintest mover. The model also returned nine answers of unclear, counted here as not moving.

The model detected every mover that the measurement detected and identified the direction of motion correctly in those cases. Its errors fell into two groups. All four bright static stars were called moving, because their residual rings in the difference resemble the paired signature of motion. Five stars with real but sub threshold motion, between 1.4 and 2.6 arcsec, were also called moving, often with the correct direction. The model perceives motion below half a pixel but does not judge whether it is large enough to assert, and its stated confidence, between 0.45 and 0.75 in every case, did not separate correct from incorrect answers.

Combining the two methods did not help in a way that can be demonstrated. Using the model to confirm the measurement removed one false alarm, which raising the threshold of the measurement to 3.5 arcsec also removes, without cost, latency or an external dependency. Both differences amount to one case out of 47 and are within statistical noise. We rejected hypothesis H3 and kept the measurement.

The experiment also improved the measurement itself. A second version, based on connected regions instead of a peak search and on matching by proximity and flux ratio, classified 46 of the 47 cases correctly and raised one false alarm among all sources of the 26 static fields. Because this version was developed while looking at the same cases, 46 of 47 is an optimistic estimate, and we report 43 of 47 as the unbiased figure until the measurement is tested on a new set.

6. Results

6.1 Detection funnel and calibration

Table 4. Number of stars remaining after each step in the NEP_off field.

Step	Stars
AllWISE sources with W1 brighter than 15.5	2,310
Measured without flags in both epochs	2,308
Compared after neighbor rejection (808 removed)	1,500
Candidates with at least three channels at 5 sigma or more	1
Candidates coincident with a known object in SIMBAD	1

The field produced 345,952 measurements, of which 12.5 percent carried a quality flag. Measured wavelengths range from 1.08 to 4.43 microns, so the channels between 0.75 and 1.08 microns are empty in this field. The calibration was determined in 29 channels, with a median of 420 reference stars per channel. The median scale between epochs is 0.986, between 0.973 and 1.000 across channels, which corresponds to a systematic difference of about 1.4 percent between the two releases. The median systematic floor is 1.79 percent, and the median noise normalization is 1.03, close to the value of 1 expected from correctly estimated uncertainties. In three channels near 3.8 and 4.2 microns the statistical uncertainties already explain the observed scatter, and the floor is zero.

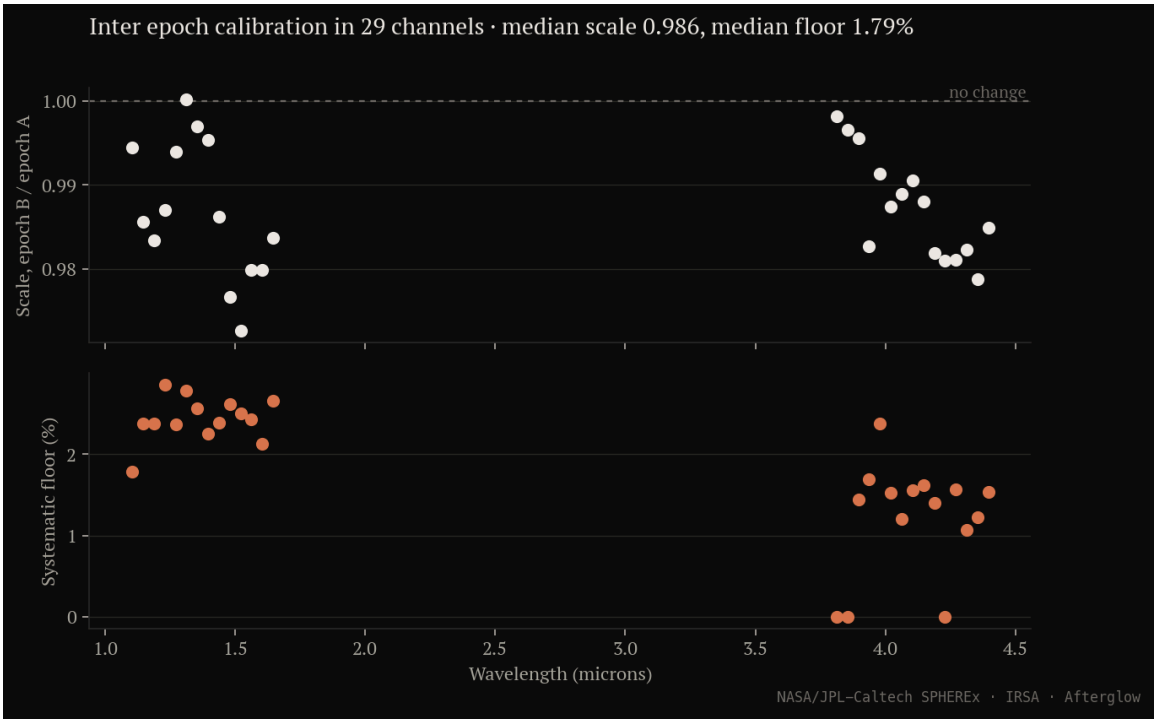


Figure 4. Inter epoch calibration of the NEP_off field. Top: scale between epoch B and epoch A in each channel, determined from a median of 420 reference stars. Bottom: systematic floor added in quadrature to the uncertainties.

6.2 The surviving candidate

The only star that passed every step is AllWISE J175232.00+653228.5, with W1 of 12.08, at right ascension 268.13334 and declination 65.54126 degrees. Its median signal to noise ratio per channel is 62 in epoch A and 46 in epoch B, and 11 percent of its measurements are flagged. Of 25 channels measured in both epochs, 6 change by 5 sigma or more, with a maximum |z| of 7.7 and a score of 16.3. The significant channels span 74 percent of the spectral range, but they do not change in the same direction: four channels between 1.27 and 1.40 microns are fainter in epoch B, while channels at 1.56 and 4.40 microns are brighter. The pipeline therefore classified the change as unexplained.

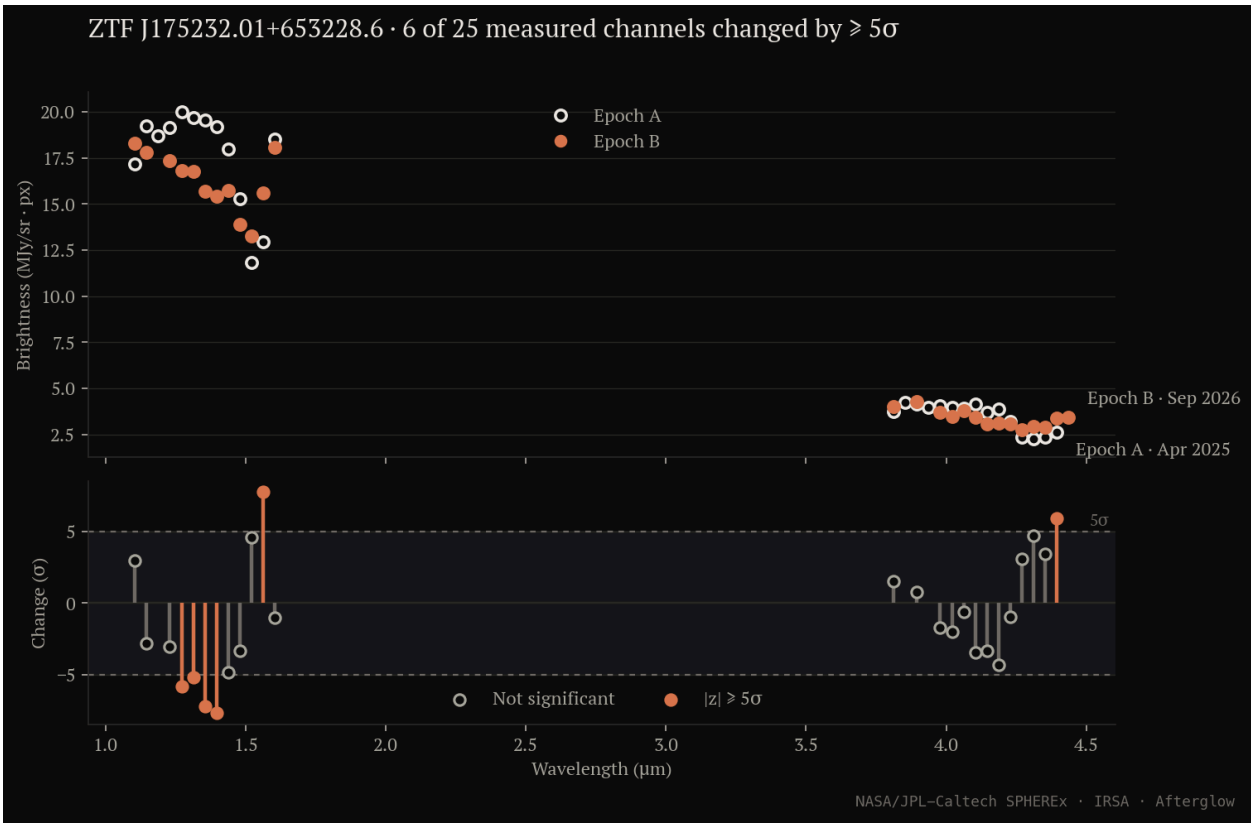


Figure 5. Spectrum of the surviving candidate in epoch A (open markers) and epoch B (filled markers), and the significance of the change in each measured channel. The gap between 1.6 and 3.8 microns reflects the channels without measurements of this star.

A search in SIMBAD within 6 arcsec returned a single object at 0.12 arcsec: ZTF J175232.01+653228.6, classified as an eclipsing binary. The pipeline had no knowledge of this object, so its recovery is an independent validation of hypothesis H1. The classification as unexplained is also consistent with the nature of the source. Each SPHEREx channel of a star is measured on a different date, so some channels of epoch B can fall inside an eclipse while others do not, which produces significant changes of opposite sign. We consider this the most likely interpretation, but we have not confirmed it against the ephemeris of the binary.

6.3 Proper motions against Gaia DR3

Table 5 lists the measurement for the five test objects named in the challenge, at positions propagated to epoch 2026.0.

Table 5. Displacement between the first and last survey pass, measured on SPHEREx stacks and predicted from Gaia DR3.

Object	Interval (yr)	Measured (arcsec)	Gaia (arcsec)	Result
Barnard's Star	0.99	10.24, north	10.27	Moving, confirmed
Epsilon Indi B	1.36	6.2, southeast	6.38	Moving, confirmed
WISE 0855	0.99	not detected	not in Gaia	Too faint to measure
RW Aurigae	1.00	0.56	below 3	Static, as expected
R Draconis	0.96	2.17	0.02	Static, as expected

Fourteen sources were measured in the five fields. The two flagged as moving are exactly the two for which Gaia predicts a shift of at least 3 arcsec. For WISE 0855 the source nearest the center is a background star; the brown dwarf itself is too faint in eight exposures per pass.

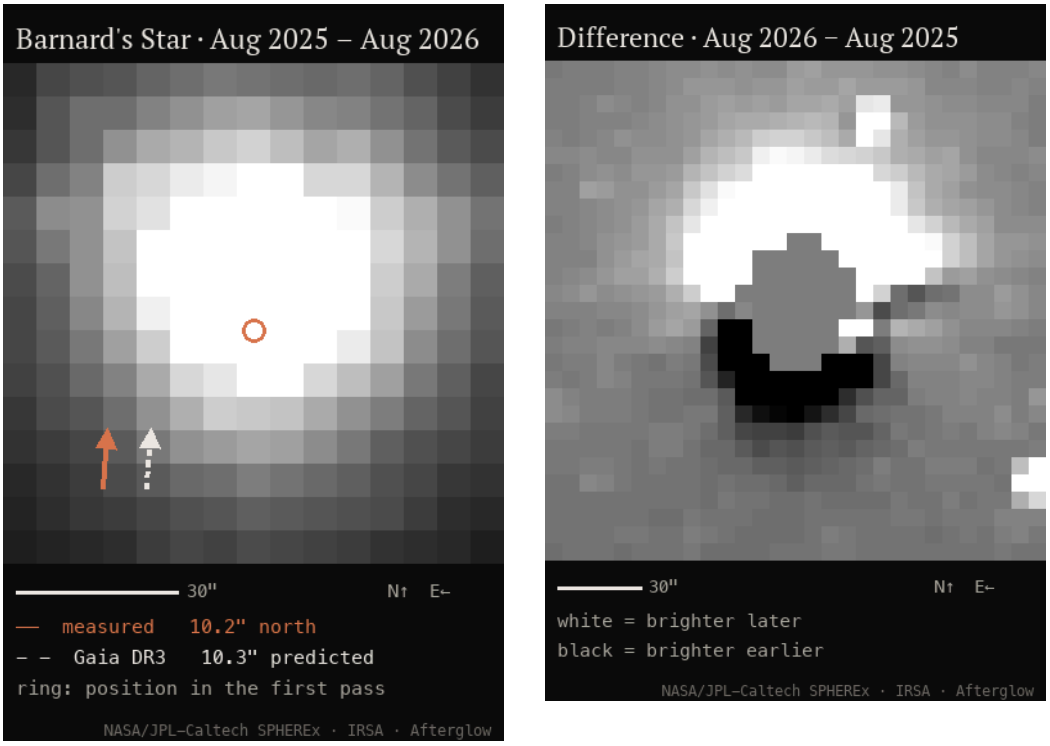


Figure 6. Barnard's Star. Left: the last pass, with a ring at the position measured in the first pass, the measured displacement as a solid vector and the displacement predicted by Gaia DR3 as a dashed vector, at the same scale. Right: difference between the last and the first pass; a source that moves north leaves a white region to the north and a black region to the south.

6.4 Known asteroids

As a test of the small body module we chose a position on the ecliptic, at right ascension 180.0 and declination 0.0 degrees, and checked 12 frames from three passes. The 12 queries to the JPL service took 174 seconds in total and none failed. Three catalogued asteroids crossed the 3 arcmin field during the exposures: 144501 (2004 EF69), with visual magnitude 21.3, (2013 BJ2), with 23.0, and 722821 (2006 CU64), with 23.5. Their positions are marked in the corre-

sponding frames. At these magnitudes the asteroids are not visible in the SPHEREx images, so the module identifies where known bodies passed but does not yet confirm them photometrically.

6.5 Single pass sources on real fields

We ran the extraction of Section 4.7 on 47 real fields. Without the requirement of two detectors it returned eleven single pass sources, all detected in one detector only and all instrumental on inspection. With the requirement it returned none. Barnard's Star, which moves about 1.7 pixels between passes, stays within the 2 pixel matching radius and appears as a pair in the difference rather than as a single pass source, as expected. This is a negative validation: it shows that the extraction does not produce false single pass sources on these fields, and its sensitivity to a real single pass source is measured by injection into these real stacks in Section 7.

7. Framework for a SPHEREx Planet X Search

The results above establish a calibrated measurement of change and motion. Here we ask what it would take to use SPHEREx data in a search for a distant planet, and where the present method stands. The numbers in this section are geometric estimates, not observations.

Expected apparent motion

For a body on a circular orbit of semimajor axis a, the orbital period is $a^{1.5}$ years with a in AU, so its mean angular motion seen from the Sun is about

$$\omega \approx 1{,}296{,}000 / a^{1.5} \text{ arcsec per year.}$$

When two passes are taken one year apart, the observer returns to nearly the same position, the annual parallax largely cancels, and the displacement is close to ω. When the passes are half a year apart, the parallax dominates, with an amplitude of up to about 412,500 / a arcsec between the two sides of the Earth's orbit.

Table 6. Approximate displacement of a distant body between SPHEREx passes, in arcsec and in SPHEREx pixels of 6.15 arcsec.

Semimajor axis (AU)	One year apart	Pixels	Half year apart, parallax
400	162	26	1,031 (17.2 arcmin)
500	116	19	825 (13.8 arcmin)
800	57	9	516 (8.6 arcmin)
1,400	25	4	295 (4.9 arcmin)
5,000	3.7	0.6	83 (1.4 arcmin)

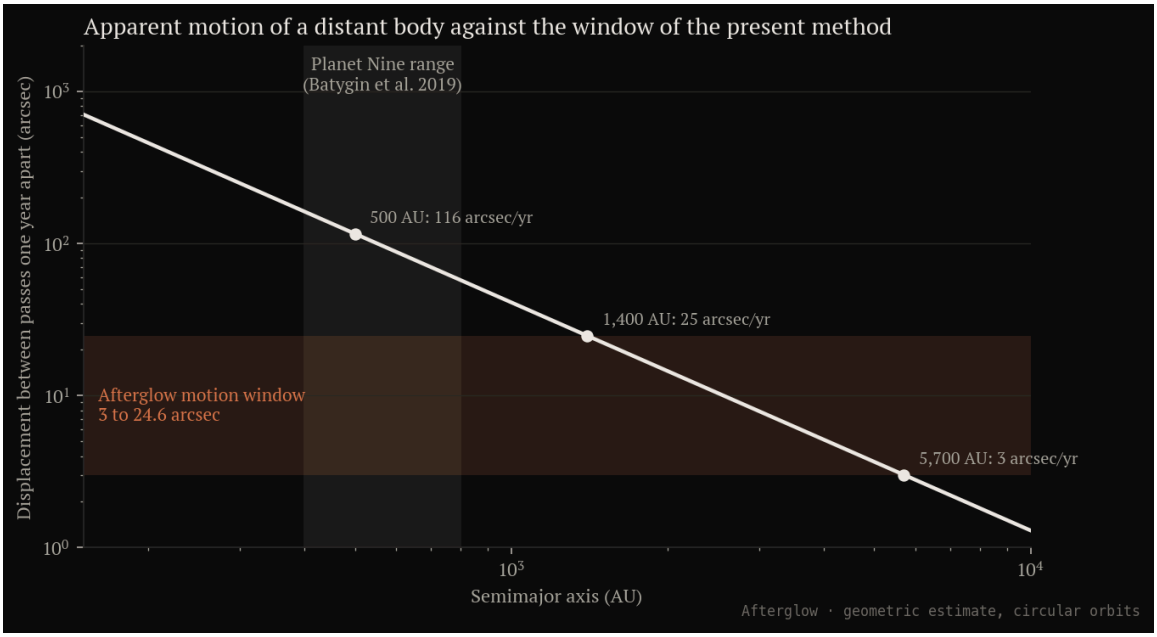


Figure 7. Displacement of a body on a circular orbit between two passes one year apart, as a function of its semimajor axis, compared with the range in which the present motion measurement operates. The shaded column marks the range proposed for Planet Nine.

Where the present method stands

The motion measurement of Section 4.5 searches up to 4 pixels, 24.6 arcsec, and asserts motion from 3 arcsec. For passes one year apart this window corresponds to bodies between about 1,400 and 5,700 AU. A planet at the 400 to 800 AU range proposed for Planet Nine would move 9 to 26 pixels between same season passes, outside the window, and would leave the 3 arcmin field between passes half a year apart. The change detection of Sections 4.2 to 4.4 measures sources at catalog positions, so a planet that is not in the catalog at the epoch of observation would not be measured at all. The present pipeline therefore cannot detect Planet Nine, and we do not claim that it does.

To make these numbers visible, the public tool includes a Planet X mode. For a hypothetical body at 300 to 1,000 AU at the center of the field in the first frame, it computes the position of the Earth on the date of every frame and draws where the body would appear, with a circle for its unknown orbital motion. At 500 AU it reproduces the displacement of Table 6, 825 arcsec between dates half a year apart. It is an illustration of the expected motion, not a detection.

Injection and recovery: could one pass see Planet Nine?

The geometry above says where a distant planet would be; it does not say whether Afterglow would notice it. We measured that directly. For each field we choose the pair of passes in which a body at 500 AU would jump the most, compute the Earth's barycentric position at the middle of each pass with the analytic ephemeris of astropy, and predict the apparent displacement of a body seen at the field center in the first pass, including its own motion on a circular prograde orbit around the ecliptic pole. We then add a fake planet to the stacks of the first pass at a random position and to the stacks of the second pass at the displaced position, or nowhere if the displacement leaves the field, with the same AB magnitude in every detector. The fake planet has the point spread function measured on 170 isolated stars in 26 real stacks (a full width at half maximum of about 11 arcsec in detectors 1, 2, 4, 5 and 6 and 27 arcsec in detector 3). The unmodified single pass extraction of Section 4.7 then searches the field without knowledge of the injected position, and a fake planet counts as recovered when a single pass source lies within 2 pixels of it.

We ran 40 injections per half magnitude from AB 14 to 20.5 in each of the 9 fields with complete stacks, 5,040 in total. Between passes about six months apart the predicted jump at 500 AU is 12.7 to 14.7 arcmin, much larger than the 3 arcmin field, so a planet appears as a single pass source in one pass only. Table 7 and Figure 8 give the result. In the four quiet fields half of the fake planets are recovered down to AB 17.4 to 17.8. Near a bright star the noise of the stack rises by an order of magnitude and the limit falls to AB 14 to 14.5; in three fields half is never reached. Even bright fake planets are lost 7 to 30 percent of the time in quiet fields, when they fall within about 20 arcsec of a real star, where the extraction merges them, or on masked pixels.

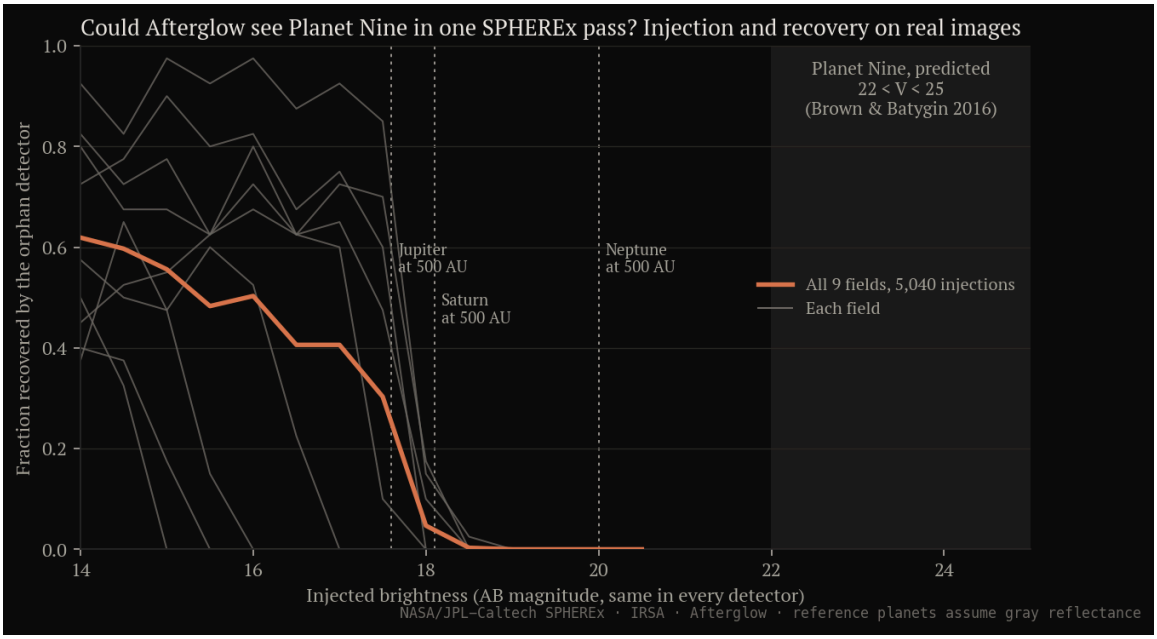


Figure 8. Fraction of fake planets recovered by the single pass extraction as a function of their brightness, for each field (gray) and for all 5,040 injections (orange). The shaded band is the brightness predicted for Planet Nine. Dotted lines mark the brightness of Jupiter and Saturn placed at 500 AU, under the approximation of gray reflectance.

Table 7. Injection and recovery per field: pass pair, predicted jump of a body at 500 AU, recovery of fake planets brighter than AB 16, and the faintest magnitude at which half are recovered.

RA (deg)	Dec (deg)	Passes	Jump at 500 AU (arcmin)	Recovered, AB ≤ 16	50% limit (AB)
217.76	+7.47	2025-06 / 2026-01	14.6	92%	17.8
124.06	+24.62	2025-10 / 2026-03	12.7	75%	17.7
133.75	-7.25	2025-05 / 2025-10	14.6	80%	17.6
307.95	+70.42	2025-12 / 2026-06	13.9	70%	17.4
77.99	-45.07	2026-01 / 2026-08	14.1	34%	14.5
269.45	+4.77	2025-08 / 2026-03	14.2	16%	14.0
43.27	+16.85	2026-01 / 2026-07	14.7	56%	not reached
1.41	-37.37	2025-11 / 2026-05	14.5	19%	not reached
164.09	+7.00	2025-04 / 2025-11	14.6	52%	not reached

Planet Nine is predicted at 22 < V < 25 (Brown and Batygin 2016), and the most recent orbit fit places it at a semimajor axis of about 380 AU (Brown and Batygin 2021); it is more than four magnitudes, or about fifty times, fainter than the faintest fake planet a single SPHEREx pass recovers in our best fields. Under the approximation that its brightness in the SPHEREx bands is close to V, a single pass would reach a body as bright as Jupiter at 500 AU (V 17.6) and miss one as bright as Neptune (V 20.0). This is the quantitative form of our negative statement: the present tool cannot see Planet Nine in one pass, and the number above measures how far it is from doing so. Stacking the passes along the predicted parallax track, the strategy of shift and stack searches, would gain about 0.38 magnitudes per doubling of the passes combined, which is not enough with the two to four passes available today. The same test is available in the public tool for any position, where it runs with 8 injections per magnitude in about one second once the stacks exist.

Requirements and existing components

A SPHEREx search for a distant planet would need four capabilities.

1. **Difference images over large areas.** Afterglow already builds, for any position, calibrated stacks per pass and detector on a common north up grid, with zodiacal light removed, and their differences. Extending them from a 3 to 10 arcmin cutout to full survey tiles is a matter of scale, not of method.
2. **Source extraction on the differences.** A moving planet appears as an isolated positive source in one pass without a counterpart at the same position in the other. The appeared or vanished class of Section 4.4 encodes this signature at catalog positions. The single pass extraction of Section 4.7 now applies it to sources extracted directly from the stacks of any field of 1 to 10 arcmin. It remains to be run over full survey tiles and validated on real moving bodies.
3. **Orbit based linking.** Detections in different passes must be linked through the displacements of Table 6, including parallax, as in the searches carried out with WISE (Meisner et al. 2018). This component does not exist in Afterglow.
4. **Rejection of known bodies and of artifacts.** The small body module of Section 4.6 removes known asteroids, and the quality flags, isolation and calibration of Sections 4.2 to 4.4 remove most instrumental changes. The false alarms on bright stars found in Section 5 indicate that residual rings around bright stars must also be masked.

A spectrum in 102 channels would then add information that earlier infrared searches lacked: a candidate planet would have to show a spectrum consistent with a cold body and inconsistent with a background star. The injection test above sets the depth such a search would start from in a single pass.

8. Limitations

1. **Saturation.** Bright stars saturate the detector. Barnard's Star, RW Aurigae and R Draconis lose their spectrum, because the affected measurements are flagged, and their displacement is marked as less precise.
2. **Motion threshold.** Displacements below 3 arcsec, half a pixel, are not reported, which is about 3 arcsec per year for passes one year apart. The upper limit is the 24.6 arcsec search window.
3. **Faint movers.** WISE 0855, a brown dwarf with a proper motion of about 8 arcsec per year, is not detected with eight exposures per pass.
4. **Partial spectra.** In the NEP_off field each star is measured in about 25 of the 102 channels, and channels below 1.08 microns are empty. Channels are measured on different dates, so a short eclipse can appear as an unexplained change.
5. **Cadence.** Two passes per year sample changes on time scales of months; faster variability is sampled, not followed.
6. **Single field.** The change detection was run on one field of one square degree near the ecliptic pole. The rate of false changes elsewhere, for example in crowded fields of the

Galactic plane, has not been measured.

7. **Asteroids.** First pass positions from JPL are accurate to about 20 arcsec, and most catalogued asteroids in a field are fainter than SPHEREx can show.
8. **Validation sample.** The blind test contains 47 cases, eight of them moving. Differences of one or two cases between methods are within statistical noise, and the tuned motion measurement has not yet been tested on an independent set.
9. **Single pass sources.** The extraction of Section 4.7 has no false detections on 47 real fields, and its sensitivity has been measured by injecting fake planets into real stacks (Section 7), not on a real moving body. The injections assume a flat spectrum in f_ν, and the comparison with reference planets assumes gray reflectance.
10. **Coverage of the all sky cube.** The SPHEREx HiPS cube used by the atlas is incomplete in some fields. In NEP_off only detectors 2 and 5 are covered, so the tool shows the 102 channel sweep on the Rho Ophiuchi cloud, where 89 to 100 percent of the tiles of every detector contain data.
11. **Citizen reports.** Reports are not reviewed before they appear in the public list, and the motion fitted from marks depends on where each visitor places them.

9. Software, Reproducibility and Public Tool

The pipeline is written in Python with astropy, astroquery and photutils (Astropy Collaboration 2022; Ginsburg et al. 2019; Bradley et al. 2024), organized so that the scientific computations are independent of databases, networks and the web interface. A suite of 73 automated tests runs without network access on simulated data with injected changes, injected motions and simulated archive failures. Every step of the pipeline can be repeated from the command line, and the blind test of Section 5 is reproducible from its own folder, including the random file names and the scoring script.

The same measurements are exposed in a public web tool. A user enters coordinates or the name of an object and sees every SPHEREx image of that position as a time lapse aligned north up, the stack of each pass, the difference between passes, the measured motion next to the Gaia prediction, the known asteroids in each frame and the spectrum measured on the fly. The atlas, built with Aladin Lite (Boch and Fernique 2014), plays the 102 channels of the SPHEREx all sky cube. The candidate of Section 6.2 and every field of the five test objects can be opened directly from a link.

Several features follow the practice of citizen science projects. A visitor who does not know where to look can ask for a patch of sky: the tool alternates between 25 selected fields with large proper motion or strong variability, each verified to have three or four SPHEREx passes, and random positions with at least two passes. On the atlas, the selected object's path is drawn from its position in the background survey (2MASS, about 2000) through Gaia (2016) to the SPHEREx passes. A field guide shows real SPHEREx examples of real changes and of arti-

facts, such as saturated stars and the residual rings of stars that did not move, and a five step tutorial introduces the tool on the first visit. Single pass sources (Section 4.7), citizen reports (Section 4.8) and the Planet X mode (Section 7) complete the tool. The application is distributed as a single container that includes the results database and the cached data of the five test objects and of further fields, so the demonstration does not depend on the availability of the archive.

10. Conclusions and Future Work

We compared two SPHEREx epochs from different data releases for 2,310 stars and showed that, after channel by channel calibration and the rejection of contaminated stars, the comparison yields one significant change among 1,500 stars, and that this change belongs to a catalogued eclipsing binary that the method found without prior knowledge. Proper motions measured on SPHEREx stacks agree with Gaia DR3 to within 0.03 arcsec for Barnard's Star and 0.2 arcsec for Epsilon Indi B, and no static source in the five test fields was flagged. A general purpose AI vision model performed worse than a simple measurement as a vetter of moving objects and was not adopted.

These results answer the challenge with a public tool whose every claim can be checked against independent data. They do not amount to a search for Planet X. Section 7 sets out what such a search requires and measures, by injection into real images, that a single SPHEREx pass falls more than four magnitudes short of the predicted brightness of Planet Nine. We plan four next steps: to test the motion measurement on an independent set of positions; to extend the change detection to fields of different stellar density; to run the single pass extraction over full survey tiles and link its sources across passes; and to confirm the interpretation of the surviving candidate against the ephemeris of the binary.

Acknowledgments and Use of AI

This publication makes use of data products from the Spectro Photometer for the History of the Universe, Epoch of Reionization and Ices Explorer (SPHEREx), which is a joint project of the Jet Propulsion Laboratory and the California Institute of Technology, and is funded by the National Aeronautics and Space Administration. This research has made use of the NASA/IPAC Infrared Science Archive, which is funded by the National Aeronautics and Space Administration and operated by the California Institute of Technology. This publication makes use of data products from the Wide field Infrared Survey Explorer. This work has made use of data from the European Space Agency mission Gaia, processed by the Gaia Data Processing and Analysis Consortium. This research has made use of the SIMBAD database and the VizieR catalogue access tool, operated at CDS, Strasbourg, France, and of the Small Body Identification API of the JPL Solar System Dynamics group.

Use of AI. The project idea and the approach were defined by the team. Code was written with the help of AI coding assistants (Anthropic's Claude, through Claude Code), which helped write and debug the software and suggested implementation details that the team reviewed. The logo was drawn as vector code with AI assistance. No data, measurements or scientific results were generated by AI; every sky image in this report is SPHEREx data rendered by our software. The AI vision model evaluated in Section 5 was rejected on the evidence and is not part of the tool.

References

Astropy Collaboration, Price Whelan, A. M., et al. 2022, ApJ, 935, 167. The Astropy Project: Sustaining and Growing a Community oriented Open source Project and the Latest Major Release (v5.0) of the Core Package.

Batygin, K., and Brown, M. E. 2016, AJ, 151, 22. Evidence for a Distant Giant Planet in the Solar System.

Batygin, K., Adams, F. C., Brown, M. E., and Becker, J. C. 2019, Physics Reports, 805, 1. The Planet Nine Hypothesis.

Bellm, E. C., Kulkarni, S. R., Graham, M. J., et al. 2019, PASP, 131, 018002. The Zwicky Transient Facility: System Overview, Performance, and First Results.

Boch, T., and Fernique, P. 2014, ASP Conference Series, 485, 277. Aladin Lite: Embed your Sky in the Browser.

Bradley, L., Sipőcz, B., Robitaille, T., et al. 2024, Photutils: an Astropy package for detection and photometry of astronomical sources, Zenodo.

Brown, M. E., and Batygin, K. 2016, ApJL, 824, L23. Observational Constraints on the Orbit and Location of Planet Nine in the Outer Solar System.

Brown, M. E., and Batygin, K. 2021, AJ, 162, 219. The Orbit of Planet Nine.

Cutri, R. M., Wright, E. L., Conrow, T., et al. 2013, Explanatory Supplement to the AllWISE Data Release Products; VizieR catalog II/328.

Doré, O., Bock, J., Ashby, M., et al. 2014, arXiv:1412.4872. Cosmology with the SPHEREX All Sky Spectral Survey.

Gaia Collaboration, Vallenari, A., Brown, A. G. A., et al. 2023, A&A, 674, A1. Gaia Data Release 3: Summary of the content and survey properties.

Ginsburg, A., Sipőcz, B. M., Brasseur, C. E., et al. 2019, AJ, 157, 98. astroquery: An Astronomical Web querying Package in Python.

Kuchner, M. J., Faherty, J. K., Schneider, A. C., et al. 2017, ApJL, 841, L19. The First Brown Dwarf Discovered by the Backyard Worlds: Planet 9 Citizen science Project.

Luhman, K. L. 2014, ApJL, 786, L18. Discovery of a about 250 K Brown Dwarf at 2 pc from the Sun.

Meisner, A. M., Bromley, B. C., Schlafly, E. F., and Finkbeiner, D. P. 2018, AJ, 155, 166. A 3pi Search for Planet Nine at 3.4 microns with WISE and NEOWISE.

NASA/IPAC Infrared Science Archive. SPHEREx Data Release 1 and Quick Release 3. https://irsa.ipac.caltech.edu/Missions/spherex.html

NASA JPL Solar System Dynamics. Small Body Identification API. https://ssd-api.jpl.nasa.gov/doc/sb_ident.html

Ochsenbein, F., Bauer, P., and Marcout, J. 2000, A&AS, 143, 23. The VizieR database of astronomical catalogues.

Wenger, M., Ochsenbein, F., Egret, D., et al. 2000, A&AS, 143, 9. The SIMBAD astronomical database.

Wright, E. L., Eisenhardt, P. R. M., Mainzer, A. K., et al. 2010, AJ, 140, 1868. The Wide field Infrared Survey Explorer (WISE): Mission Description and Initial On orbit Performance.

Appendix A. Pipeline Parameters

Table A1. Parameters of the change detection and motion measurement.

Parameter	Value
Epoch window at each end of the data	30 days
Catalog limit, AllWISE W1	15.5 mag
Aperture radius	2 pixels
Background annulus	4 to 7 pixels
Margin from the image edge	8 pixels
Quality flags	all bits except SOURCE (21)
Spectral channels	102, 0.75 to 5.0 microns
Calibration reference stars	SNR above 10, at least 20 per channel
Systematic floor reference stars	SNR above 50
Isolation from any source	15 arcsec
Isolation from a source 1 mag brighter	45 arcsec
Significance threshold	5 sigma in at least 3 channels
Maximum flagged fraction	0.5
Detection and non detection, appeared or vanished	SNR 5 and 2
Brightness change, span and coherence	0.30 and 0.80
SIMBAD match radius	6 arcsec
Motion stacks	up to 8 exposures per pass and detector
Motion detection threshold in stacks	5 times the robust noise
Motion search radius and threshold	4 pixels; 3 arcsec
Gaia match radius and margin	6 arcsec; 150 arcsec
JPL query window and uncertainty limit	12 hours; 120 arcsec

Appendix B. Blind Test Protocol

1. **Selection.** High proper motion stars were retrieved from SIMBAD by name; static stars were drawn at random, with a fixed seed, from the isolated NEP_off sources with W1 between 9 and 13, and bright static stars from those with W1 brighter than 6.5.
2. **Ground truth.** A case is moving when the catalogued proper motion times the interval between the first and last pass is at least 3 arcsec. The truth was written to a separate file before any classification.
3. **Images.** For each case, four exposures per detector in the first and last pass were stacked and rendered with the same intensity scale for both passes, north up, with the difference in a symmetric gray scale saturating at 5 times its robust noise.
4. **Blinding.** Images were copied under random names to a folder containing nothing else. The evaluator received a fixed instruction and the list of file names, and was not allowed to read any other file.
5. **Scoring.** Answers were mapped back to cases only after all of them were recorded. Unclear answers count as not moving.
6. **Availability.** The 47 cases, the truth file, the answers and the scoring script are kept with the project for inspection.