

Version history (PyOTE description – out of date - is at end)

Version 5.6.5 22 December 2024

- Fixed failure to allow 2 point identification of integration block to use when a LiMovie file is being processed that also requires manual timestamp entry.

Version 5.6.4 22 December 2024

- Fixed failure to Accept block integration when a LiMovie file is being processed that also requires manual timestamp entry.

Version 5.6.3 21 December 2024

- Added a pop-up message when the Start over button is clicked when block integration is present, warning the user that re-reading the csv file is required if the block integration needs to be changed.

Version 5.6.1 20 December 2024

This version changes the way in which PyOTE estimates the extent to which an apparent 'event' might have been caused by fortuitous noise in the observation rather than an actual occultation.

It removes the Gumbel Extreme Value calculations used to predict drop values at the 3, 4, and 5 sigma positions. Besides being hard to grasp intuitively how to use these values, there are cases where, while the calculations are technically correct, they result in values that are physically impossible yet were being used in deciding whether to report an event – or not.

The new approach is described this way in the 'help' screen that appears any time the NIE (Noise Induced Event) calculation is invoked:

Noise Induced Event (NIE) distributions and noiseSigmaDistance metric (sigma distance between position of noise-only peak and actual observed drop)

Two plots are shown, the result of 50,000 simulations using the noise parameters extracted from the actual observation lightcurve, including the effects of correlated noise.

The blue plot is a histogram of the maximum drops found in a search of 50,000 noise-only simulated observations matching the duration of the found/candidate event (number of readings in the found event), the length of the observation (total number of readings), and with no constraint as to where the drop occurs in the simulated observation lightcurve.

The green plot shows the distribution of drops around the found drop, using a gaussian distribution with the peak positioned at the found drop value and a sigma value computed from the quadrature sum of the baseline mean sigma and the event mean sigma.

The amplitude of the green plot is chosen so that the area under the blue and green plots are equal. This choice makes the green plot reflect 50,000 simulations too. Now the green plot shows the distribution of drops that would be found if there was an 'event' while the blue distribution shows the distribution of drops when only noise was present in the observation lightcurve.

The further to the right of the peak of the noise-only distribution the observed drop occurs, the less likely it is that the found event is due to noise.

A metric is provided for reporting purposes calculated by dividing the drop distance between the noise-only distribution peak and the observed drop by the sigma parameter of the green plot: this metric is called **noiseSigmaDistance**.

If **noiseSigmaDistance** is positive and greater than 2.0, one can have confidence that the found event is highly unlikely to be due to noise in the observation. There will be a clear and obvious separation between the two distributions.

Note: there are many other factors that go into an assessment as to whether an observation has recorded an 'event', such as size of drop compared to expectations, duration of the event compared to expectations, other observer's results for the same occultation, etc. The **noiseSigmaDistance** metric is just one factor and should never be used alone to decide whether or not an occultation event has been recorded. not an occultation event has been recorded.

Version 5.6.0 11 November 2024

- Fixes a problem in the diffraction model calculation involving parameters supplied to the ellipse calculation subroutine.

Version 5.5.9 31 October 2024

- Fixes a problem in the disk-on-disk model calculation involving parameters supplied to the ellipse calculation subroutine.

Version 5.5.8 17 October 2024

- Fixes a problem in the NIE calculation that was causing the program to hang in an infinite loop for some data sets.
- Fixes some edge case errors that were introduced when I incorporated suggested changes from Graeme resulting in version 5.4.9 My regression testing at that time was not complete enough to find all the unintended consequences from those changes - my bad – now remedied. Thanks again to Graeme.

Version 5.5.2 22 July 2024

- This change is invisible. We report it here for the record. It was necessary during the NIE calculations, which involve numerical integration and curve fitting, to normalize two lightcurve parameters (observed_drop and noise sigma) to avoid numerical instability that resulted in infinite loops when intensity values got very large. The normalization is performed just before the NIE calculations are performed and immediately reversed after the NIE numbers have been calculated, thus making the changes invisible.

Version 5.5.1 3 February 2024

- fixes the display of NIE probability in the plot that accompanies a single-point-drop event. The logged information was correct, so no change there, but now the legend on the NIE histogram plot matches that value reported in the log.

Version 5.5.0 30 December 2023

- Incorporates ViZieR changes from Phil Stuart. Thanks Phil.

Version 5.4.9 30 December 2023

- Incorporates changes from Graeme Schmidt to fix r-only and d-only solutions. Thanks Graeme.

Version 5.4.8 19 November 2023

- Minor change to format of fit metric report involving NIE

Version 5.4.7 16 November 2023

- Fixed an error in the reporting of NIE probability

Version 5.4.6 22 October 2023

- Fixed Noise Induced Event calculation slow processing problem (for some lightcurves, it was so slow that it appeared to 'hang')
- Fixed the probability printed when a single point event is evaluated for validity.

Version 5.4.5 8 October 2023

- Corrects a math error in the estimation of the sigma level of a drop that is more than 5 sigma (it was off by 5 – too low)

Version 5.4.4 7 October 2023

- Edits around wording regarding NIE green error bars

Version 5.4.3 5 October 2023

- Fixes flash tag processing (some earlier version broke it)
- Fixes ADV/AAV usage (also broken by an earlier version change)
- Much work was done in the **Other models** tab to make it more user friendly and intuitive.

Version 5.4.2 1 September 2023

- The False-Positive test is now referred to as the Noise Induced Event (NIE) test. This more accurately describes what the test is by its title and avoids the confusing contradiction in 'false' and 'positive' being used in close conjunction for the old title.

Version 5.4.1 28 August 2023

- Added the computation of a '3 sigma line' to the False Positive test. That line (shown in green) is added at the drop value that has 99.7% of all noise-induced drop to its left.

The max_drop from the 50,000 trials against a noise-only baseline remains and is always shown to the right of the green three_sigma_line as a black line.

The three_sigma_line is very repeatable from run to run, unlike the max_drop line which has large amount of jitter.

Because the three_sigma_line is stable and repeatable, it is a better choice for comparing against the observed drop – if the observed drop is to the right of the three_sigma_line, PyOTE now reports that it passes the false-positive test in the metrics printout. If the observed drop lies between the green line and the black line, it will have a non-zero false-positive probability value (but less than 0.003).

Reporting an event with a brightness drop that puts it in between the green line and the black line should be done in consultation with a regional coordinator.

Version 5.4.0 24 August 2023

- Adds a red circle around any reading in either a 'signal' or an 'appsum' light curve that included 1 or more defect pixels in its sampling mask. For such red circles to appear, PyMovie version 3.8.6 or greater and the use of a dark video recording in conjunction with the observation recording is required.

Version 5.3.8 7 August 2023

- Adds csv source file information to fit metrics.xlsx file

Version 5.3.6 29 July 2023

- Colors a FP FAILED metric log entry red

Version 5.3.5 22 July 2023

- Fixes Mac filepath issues

Version 5.3.4 18 July 2023

- Adds formatting to the .xlsx fit metrics file and provides a default name of the enclosing folder

Version 5.3.3 17 July 2023

- Fixes a bug in the fit-metric.txt file duration calculation

Version 5.3.2 15 July 2023

- Fixes a bug in the calculation of the duration fit metric when an already integrated csv file is in use.
- Added specialized folders inside an observation folder to reduce clutter. There are new folders titled: FP-plots, Detectability, Normalization, ApertureGroups, OCR, and Finders inside an observation folder

Version 5.3.1 10 July 2023

- Squeezed the lightcurves tab to allow a little more space for the plot for those that need to run PyOTE on a small laptop screen

To deal with pixel-challenged laptop screens ---

1: the plot can be made to span the entire width of the GUI by dragging splitter all the way to the left

2: or you can change your pixel resolution to 150% instead of the default 200%

Version 5.3.0 8 July 2023

- Changed the initial test that can result in 'no event found' to accept more events – this test was extremely conservative and was a hold-over from a time when there was no False-positive test. Now we can rely on the FP test to weed out events that might be due to noise and so we can loosen up the initial test.
- Fixed some issues around removing and adding data sets on the light curve tab. Now the data grid at the lower left tracks the deletes and adds properly and the entire list of light curves can be cleared.

Version 5.2.8 30 June 2023

- Added detection of duplicate column names (only possible in a csv file from PyMovie).

Duplicate names cause the column data to be concatenated and this cannot be tolerated, so when this error is found, the data import is aborted.

The user will need to edit the csv file to resolve the name duplication.

- Added a warning message that will only appear if a user manually calculates baseline statistics and then does a **find event**. Sometimes that is exactly what is wanted, but those statistics become fixed and will be used for all subsequent **find event** clicks.

That would be very wrong if the user is stepping through multiple light-curves in search of the "best" as they would all use the same baseline numbers (B and sigmaB).

- Added **false-positive metric** to fit-metrics and rearranged output order.

Version 5.2.7 24 June 2023

- Minor change: corrected titles in fit-metric file from **D frame** to **D reading num** and **R frame** to **R reading num**

Version 5.2.6 17 June 2023

- Changed order of buttons on Noise analysis/Detectability tab to better reflect work flow
- If a minEvent of 1 is entered for the min/max search, it is automatically changed to 2 (with an advisory message to this effect) to ease work flow.
- Added buttons to **Lightcurves** tab to make it more convenient to calculate and recalculate baseline statistics for multiple light curves.

- Added a **Disable err bar plots** check box to the **Square wave model** tab to suppress error bar and false positive plots to ease the task of comparing light curve extractions from multiple apertures – no need to close two plots that are only needed for the final analysis.
- Added a block of the most useful metrics for comparing light curve extractions from multiple apertures at the end of the text report window and also in the **.log.txt** report.
- The metrics for each ‘solution’ are added to a text file (fit_metrics.txt) – this text file can be converted to a .xlsx file which can be used to support the choice of ‘solution’ submitted in a report.
- **Changed snr to dnr (drop to noise ratio) in the places where that is the actual calculation being performed.**

Version 5.2.4 11 May 2023

- Fixes an edge case – when only two points were in the event bottom, sigmaA was being set to 0.0 (causing an assertion error) rather than setting it to the value previously determined for sigmaB.

Version 5.2.3 2 May 2023

- Restores the calculation of timeDelta to 6 digit precision (it had been inadvertently limited to 4 digits by a previous change)
- Fixes a bug in the calculation of flash edge times that was discovered when very high SNR recording was attempted

Version 5.2.2 13 April 2023

- Adds diffraction modelling of an elliptical asteroid shape

Version 5.2.1 27 March 2023

- Fixed problem with Vizier loading from North American .xlsx caused by Occult sometimes putting a string in the asteroid field and sometimes an integer.
- SER files with colorCode other than mono are now accepted.

Version 5.2.0 17 March 2023

- Fixes a bug in code introduced in version 5.1.0 where dropped frames are detected – that code did not handle an observation that passed through midnight (23:59:59)

Version 5.1.9 10 March 2023

- Fixes issues involving Tangra light curve names
- Clears previous light curves names when reading a new csv file.

Version 5.1.8 9 March 2023

- Fixed GUI problem that just keeps re-occurring behind the scenes – frustrating.

Version 5.1.7 9 March 2023

- Modified the csv file reader to read the AOTA compatible Tangra csv files: binned with background subtraction; un-binned with background subtraction; un-binned without background subtraction.
- At first 'read', a list of all available light curves for display is popped up and all the 'signal' light curves are populated in the light curves list on the Lightcurves tab.

This was done so that the user is always made aware of the existence of multiple light curves read from the csv file.

Version 5.1.6 11 February 2023

- Fixes Vizier file preparation code to detect when lat or long is given with a + sign (common when input comes from reading a .xlsx file) and avoid adding a redundant + (double) character.
- Reverts to using pyqtgraph 0.12.4

Version 5.1.0 7 February 2023

- Changed the way data sets (lightcurves) read from a PyMovie csv file are selected for viewing. Whereas previously it was possible to specify a 'prefix' to be used when reading PyMovie csv files (such as 'signal' or 'appsum'), it was not possible to mix-and-match data sets with different prefixes. With this version, that now becomes possible.

With this feature, it now becomes possible to demonstrate for yourself that there is no reason to process a target lightcurve in 'aperture sum' mode (no background subtraction) versus the regular 'signal' mode (where background subtraction is performed). Simply display the 'signal' and 'appsum' versions of a target lightcurve and use the y position spinBox on one of the curves until they overlap. This should show that during the event, the points almost exactly overlap, so event timing will be the same.

- Added a button to the Vizier tab that will fill most of the entry boxes in the Vizier form from a North American .xlsx file for that event (derived from OW/Occult 4), if one is available.

Version 5.0.8 1 February 2023

- Changed the Vizier form filling requirements to allow for all star catalog entries to remain unfilled.

Version 5.0.7 22 January 2023

- Fixed insidious and dangerous bug in Vizier tab that caused site coordinates to be saved and restored incorrectly.

Version 5.0.6 19 January 2023

- Fixed some data entry/edit issues in the new Other models tab.

- Fixed a bug in the edge-on-disk model that was causing edges to be misplaced when D and R limb angles were unequal and the D angle was not 90 degrees. This was noticeable in the model-example provided for testing/training.

Version 5.0.5 17 January 2023

- Fixes squarewave solution display to (once again) honor the checkboxes that control camera response, raw light, and edge plotting.
- Adds a tab for exporting VizieR archive lightcurves

Version 5.0.3 9 January 2023

- model-examples folder is now being placed in the Documents folder rather than the Desktop folder.
- Launch from **python3** or **python** or **py** now detected and dealt with when checking PyPI repository for latest version.

Version 5.0.2 8 January 2023

- Added timing problem report – counts of dropped readings; duplicated readings; cadence error. These appear in the log box (lower right corner) and are written to the log file.
- Made D and R edge time calculations work properly even when there are dropped readings in the observation

Version 5.0.0 6 January 2023

- Added ability to specify a block size to be used during automatic block integration. When a block size is entered, the program will determine the best offset value to use for the integration.
- Added a comprehensive set of lightcurve models together with tools to help automate the fitting of the new models to observation data. Models included:
 - Diffraction (including central spot for low rho events) using physical optics
 - Disk on disk (geometrical optics) for large stars
 - Edge on disk (geometrical optics) for large stars and limb angles on asteroid

Version 4.9.2 14 August 2022

- Added various tests involving the penumbral fit process to gently lead the user back to the path of rightness should that user try to do a 'fit' without loading a light-curve and/or not marking D and R regions.

Version 4.9.1 10 August 2022

- In the diffraction tab, added a selection box so that asteroid distance can be entered either in parallax (arcseconds) or AU. Also reviewed and added to the context help on this tab to be more explanatory.

Version 4.9.0 9 August 2022

- Added a button to the diffraction tab for use during penumbral fits to do a re-fit.

Improved (hopefully) the context help regarding the selection of D and R regions for a penumbral fit (cautioning the user to only include points from the penumbral transition).

Version 4.8.9 8 August 2022

- Timestamp error lines now distinguish between 'cadence jitter' and 'dropped frames' by using an additional color.

Cadence jitter happens when high frame rates create a situation where frame recording cannot keep up and the traffic control implemented by the computer changes the arrival/timestamp time of frames so that some appear a little earlier and others a little late.

Cadence jitter is defined as a timing discrepancy of between 20% and 80% of a frame time and is represented by a yellow line.

Timing discrepancies greater than 80% of a frame time continue to be shown as a red line and are assumed to come from dropped frames. In any case, they deserve attention.

Version 4.8.8 7 August 2022

- Added message at normal program close to the effect that QBasicTimer messages that appear after the program has closed are harmless artifacts of the order in which QUI elements are closed.

Version 4.8.7 29 July 2022

- changed numpy version requirement to match that of PyMovie so that ultimately the two programs can share a virtual environment, thus saving much space.

Version 4.8.6 23 July 2022

- added instructions telling the user how to install a new version when one is found. There are instructions for pip based installations and pipenv installations

Version 4.8.4

- there are no changes. This is a new version just to test the change introduced in 4.8.3 about version detection and handling thereof.

Version 4.8.3

- this version removes the offer to update to a newer version when one is found. This has been done primarily to allow an easy transition to a distribution method that does not require the installation of Anaconda3, but it also gives the user more control. When a newer version is available, a message to that effect will be shown but otherwise the user's workflow is not interrupted by a 'nag' about 'do you want to download the newer version?'.

Version 4.8.2

- changed the format of the false-positive probability report because very small but non-zero values were showing (confusingly) as probability 0.0000 That no longer can

happen.

Version 4.8.1

- fixes automatic block integration failure under Win11.
- makes pyote install for new version update use pipenv if present

Version 4.8.0

- adds PyQt5 to list of required packages in support of using pipenv to install PyOTE without the need for an Anaconda3 installation.
- removes the writing of the startup batch file as a different method will be used in conjunction with pipenv.

Version 4.7.9

- disables the compiling to C code (which was done to increase execution speed) of routines used for finding an event using min/max size. There is compelling (but indirect) evidence that the C compiler on Win11 (and maybe only for the i5 processor) is generating incorrect code. This issue showed up recently with 2 users who had installed Win11 on i5 processor-based computer. On those machines, finding an event using min/max would always hang at 99% completion. Allowing the original Python code to be used instead of compiling to C-code resolved that issue. As finding an event using min/max size minimizes the number of candidate solutions naturally, the loss of speed is not so important – it's fast enough.

version 4.6.4

- made the 'camera response' checkbox (in the Settings/misc tab) sticky
- corrected the log messaging about Donly, Ronly, DandR, and min/max directions to the solver consistent

version 4.6.3

- in detectability: logged new 'drop message' to the detectability.log

version 4.6.2

- in detectability: when event duration is or becomes less than the camera exposure time, the observed drop is reduced proportionately

version 4.6.1

- in detectability: changed calculation of number of event points from $\text{ceiling}(\text{dur}/\text{rdgTime})$ to $\text{round}(\text{dur}/\text{rdgTime})$
- in detectability: allow a single point event.
- In detectability: stop stepping down duration when calculation of number of event points reaches zero

version 4.6.0

- clarified the titles on detectability plots

version 4.5.7

- suppressed a 'min event limit too high' message when min event is set to lowest possible value of 2

version 4.5.6

- removed the 'busy code' which changed the mouse cursor to "I'm busy". This seems to have the side effect (on some computers – but none of mine) of suppressing the cursor change when hovering over a splitter. As this is confusing, the busy cursor feature has been removed.

version 4.5.5

- fixes a bug where, if a full D and R light-curve was split (trimmed) into a D-only and an R-only light-curve, the R-only result did not apply subframe timing adjustments

version 4.5.4

- cleans up the the first bin of the False-Positive histogram

version 4.5.3

- adds an option to validate a single point event to the Analysis tab. It uses the false-positive test procedure to determine if the selected point had a non-zero chance to have been caused by baseline noise alone.

version 4.5.2

- fixes a bug where transition points (only present in high SNR light-curves) were not properly ignored during calculation of event noise. This manifested itself most clearly in too large error bars for magDrop

version 4.5.0

- fixes a bug introduced with the NE3 code changes that inhibited sub-frame timing adjustments.

version 4.4.9

- fixed yet another manisfestation of the bug introduced in version 4.2.1

version 4.4.8

- fixed another manifestation of the bug introduced in version 4.2.1

version 4.4.7

- fixes a bug introduced in version 4.2.1 (February 8, 2022) that caused D values to be calculated as occurring 1 reading early

version 4.4.5

- added a checkbox to enable selection of whether or not the camera response curve is shown on the lightcurve plot

version 4.4.4

- detectability and normalization activities are written to dedicated log files:

*.detectability.PYOTE.log
*.normalization.PYOTE.log

version 4.3.8

- suppressed writing of detectability results and normalization activities to the log file. The 'writings' still appear in the normal place, they just don't get written to the log file.

(Later I may direct such stuff to a separate log file.)

version 4.3.7

- removed the Pearson R metric – it was confusing and ineffective compared to the simple standard deviation metric
- made it possible to select the points of the normalized target lightcurve to be included in the std metric

version 4.3.6

- added a simple metric as an aid to selecting an appropriate number of readings to use to smooth the reference curve used for normalization. It is printed in green whenever normalization is active. It is simply the standard deviation of the normalized target lightcurve. Although it is quite simple, it is a value that is minimized the more level and the less bumpy the normalized target curve is.

When normalization is active, at each change, a pair of metrics are printed in the log panel.

The red metric is the Pearson R value of the correlation between the target lightcurve and the reference lightcurve; it should be maximized (using the X offset spinbox of the reference curve) at the point where the two curves are properly time-aligned. It is also affected by the smoothing interval, but this connection should be ignored.

The green metric is the standard deviation of the target lightcurve – it is lowered when the target lightcurve has a minimum slope and the fewest and shallowest bumps. Use the spinbox for the number of readings in the smoothing interval to minimize this number.

version 4.3.5

- made the 'step-by' buttons radio buttons so that it is clear what the current step size is.
- removed the automatic switch to the Lightcurve panel

version 4.3.3

- added a Pearson R calculation whenever normalization is being performed. This may be an aid in aligning time shifts between the target lightcurve and the reference lightcurve.

The idea is that the user will observe this value as left and right (X offset) shifts are performed and choose the position that gives the highest positive value for this number

(it will always be between 1.0 and -1.0).

version 4.3.2

- adds step size selection for the Y offset spinboxes
- activates the 'show' checkbox on reference curves so that the reference curve display can be turned off – but it will still be active behind the scenes, just not visible.

version 4.3.1

- fixes an error message that occurs during the initial installation of a version that has a change to the number of tabs in the GUI

version 4.3.0

- minor cosmetic cleanup to deal with Windows GUI differences (I hope it works – this is always a frustrating and mysterious area)

version 4.2.9

- adds the lightcurve used to the report generated during the detectability analysis
- automatically clears the effect of previous normalization when the reference lightcurve is changed (or deselected)

version 4.2.8

- fixes 'trim problem'
- adds 'step by' buttons that change the smoothing interval spinbox step size.

version 4.2.6

- tidies up the lightcurve panel a bit and changes color samples for target and reference curves.

version 4.2.5

- completes implementation of the time shift (left/right shift) of the reference curve for normalization.
- adds dot color samples next to the lightcurve titles.

version 4.2.4

- makes tab/panel order sticky
- analysis tab always opens first, but automatic switch to Lightcurves tab occurs upon file read.
- 'target' lightcurve cannot no longer be moved up or down. If there is an overlapping lightcurve, move the non-target lightcurve to clear the overlap.

version 4.2.3

- fixes block integration that was broken by the new lightcurve selection code.

- partially implements time shift of the reference curve for normalization – the visual shift is implemented, but the supporting math is not in place so the normalization is currently unaffected by the reference lightcurve shift; normalization is done with the unshifted reference curve – this is just a necessary step in the code development..
- adds logging/reporting of the curve (if any) used for normalization and the smoothing interval used.

version 4.2.2

- adds a new tab/panel for controlling the display of lightcurves. From the 'help button' on that tab:

This panel allows up to 10 lightcurves to be displayed at the same time. If the csv file has more than 10 lightcurves, the first 10 are displayed.

The **target** lightcurve is always drawn with bright blue dots.

If a lightcurve is selected as a **reference** to be used for normalization, it is always drawn with bright green dots.

Unless a lightcurve is designated as a **target** (curve to be analyzed for an event) or is designated as a **reference** lightcurve, its dot color depends on the row it is in - every row has a unique color other than blue or green.

Lightcurves can be displaced up or down using the **Y offset** spinner to control the displacement. This affects the display position only; the underlying values are not affected. This facility was added to allow the separation of lightcurves that would otherwise overlap in a confusing manner.

There can only be one lightcurve selected as **target**.

There can be either 0 or 1 lightcurve selected as a **reference** for normalization.

Normalization is applied whenever the **normalization reference curve smoothing interval** spinbox is changed from 0. Whenever this number is changed, a new normalization will result. If this number is returned to zero, all normalization is removed and the original values restored.

The **X offset** spinbox is not yet implemented. Its (future) purpose is to allow the **reference** curve used for normalization to be 'time shifted' for those cases where a drifting cloud affects lightcurves at slightly different times.

- the normalization itself has been simplified so that no user input is needed other than the number of readings to be used in the smoothing procedure. The smoothing procedure is a double application of a first order Savitzky-Golay filter (a straight line) to the points included with extrapolation.
- added (back again) a spinbox on the Setting/Misc. Tab that allows the dot size used in lightcurve plots to be changed.

version 4.2.1

- fixed final report showing NE3 stuff even when NE3 not in use was checked.

version 4.2.0

- improves annotation on **Detectability** plots by including the magDrop information.
- eliminates the display of negative drops in the False-Positive histogram

version 4.1.9

- fixes (hopefully) Win 11 issue involving the directory separator character \ Win 10 accepts / as a separator – Win 11 apparently does not (unless there is some setting that will persuade Win 11 to accept either separator).

version 4.1.8

- adds option to write sample light-curve from detectability analysis to a csv file that be imported to PyOTE.

version 4.1.7

- adds the requirement to specify an observation duration when doing a detectability test
- during a detectability test, if a detectable event was found, a sample light-curve showing such an event is plotted, otherwise the normal False-Positive plot will be shown.

version 4.1.6

- reenables Cholesky failure messages (disabled for testing)
- Adds test to detectability routine so that the user cannot give an event duration that requires more points than are available in the observation.

version 4.1.5

- changed error bar calculation reporting so that the the value reported in the containment interval report matches that reported in the Excel final report when there is no asymmetry present.

version 4.1.4

- This version adds tools needed to make best use of the Night Eagle 3 camera, the successor to the Night Eagle Astro camera (which is no longer in production). They are all grouped on a new tab page titled **Night Eagle 3**.
- The Night Eagle 3 is a rolling shutter CMOS camera. As a result, with this camera, the timing of an occultation depends on which row the occulted star is at when the occultation occurs. PyOTE will automatically calculate the needed time correction from the y (row) position that you enter in the spin/entry box at the bottom of this tab.
- For the Night Eagle 3, times extracted by PyOTE require only a VTI correction (if you use an IOTA VTI, there is no correction at all needed).
- Note: if the recording was **not** made with a tracking telescope (so the target star is moving across the image), you will need to watch the video and record the row of the

occultation for use in this program. Normally, with a tracking telescope, you will be able to use the y position of the 'target' aperture directly. PyMovie includes this information in the comment lines of the csv file.

- The Night Eagle 3 has a very effective noise reduction system called DNR (Dynamic Noise Reduction). There are 4 levels of noise reduction and you will need to indicate on this tab page what DNR setting your recording was made with.
- The noise reduction is not without a small cost however - the edges of an occultation light-curve will no longer be a step change. Instead, the edges will follow an exponential curve, approaching maximum and minimum intensities asymptotically. PyOTE has the ability to fit such an exponential curve to the D and R transitions, so the time resolution will be restored during the least squares fit and you should feel free to use whatever level of noise reduction you may need. It is recommended that you run your NE3 at a gamma of 0.75 (1.0 gamma is not available) AND that you use PyMovie to linearize the recording (i.e., invert the 0.75 gamma curve of the camera).
- The D and R exponential curves each have their own time constant, measured in frames, that control the frame rate of exponential curve growth. The default values provided are usually enough to provide the starting point for a good fit. This starting point is used by PyOTE during a least-squares driven search for better time constant values. If the starting point given is too far from 'correct', the least-squares search may settle in a local minimum that produces a fit that is visually bad. In this case, change the starting value to something closer to 'correct' and try again.
- When the exponential curve fit algorithm is in use, the light-curve plot PyOTE displays will be changed. Gone is the blue 'camera response' curve, replaced by brown dotted 'theoretical' exponential edge curves that you can use to judge for yourself the goodness of a 'fit'.
- You will probably want to use the most aggressive noise reduction in most cases, but make sure that your expected event duration is long enough that the D transition exponential curve has time to settle to the bottom event plus some time to allow a good determination of the event bottom intensity. At DNR:HIGH, your expected event time should be greater than 30 frames (1 second) to use this setting.
- I suggest the following rules-of-thumb: can use DNR:HIGH for events 2 seconds and longer; can use DNR:MIDDLE for events 1 second and longer; can use DNR:LOW for events 0.5 seconds and longer.

version 4.1.2

- fixes bug when D only search has been selected

version 4.1.1

- fixes bug where signal columns beyond four were being read as duplicates of column 5.

version 4.1.0

- adds error bars to reported magDrop values.

This was an interesting exercise in the propagation of errors because the magDrop calculation involves $\ln(A/B)$ where we know the error bars for A and B but have to propagate those errors through the ratio A/B to get its error bars and then through the $\ln()$ function to get the final error bars. The equations I use were verified through simulations.

version 4.0.9

- fixes issue where PyOTE could not read files that it had written.

version 4.0.8

- fixes bug (introduced by a misspelling in version 4.0.7) that kept csv files with more than four apertures from running.

version 4.0.7

- adds the ability to select the PyMovie data column type to be analyzed. There are two main column types that are likely to be used: signal, which has background subtracted and appsum, which has no background subtracted – it's the raw mask pixels summed.

Some others are available: avgbkd (shows the average pixel value in the aperture); stdbkg (shows the pixel noise in the background pixels in the aperture); nmaskpx (which shows the number of pixels in the sampling mask – there are times when this curve can be used to extract event times).

Note: this only applies to csv files generated by PyMovie.

version 4.0.6

- minor GUI changes

version 4.0.5

- if you click outside the light-curve, the closest point will be toggled (so you don't have to zoom in to select the leftmost or rightmost point in a light-curve).
- When baseline statistics are calculated, they are reported as **baseline mean** and **baseline snr** (instead of just snr, which is the term we use when a drop is present).

version 4.0.4

- a few GUI changes
- removed the Use as secondary check box. Now if secondary curve is selected (spinner is non-zero), it will be used automatically.

version 4.0.3

- implements a major GUI change: uses tabbed folders to group buttons and other controls.
- Removes radio buttons that indicated a D and R solution wanted, or a D-only solution, or an R-only solution. The type of solution is now inferred from D and R region setting.

When min/max event size is used as the search criteria, then a D and R solution is assumed and found. This type of search criteria does not make sense for a D-only or an R-only solution. Use a D region or R region selection for these cases

- The 'selected curve to analyze' number and the 'curve to normalize' number are now allowed to be the same – no funny jumping around.

It is even possible to normalize a curve against itself. There may be a use case where this is useful, so I didn't lock it out.

version 4.0.2

- removed a debug print statement that was inadvertently left in. It is responsible for the “mouse event” messages polluting the log file.

version 4.0.1

- increased width of detectibility plots to (hopefully) deal with Windows tendency to truncate the plots. If this doesn't work, then this will be a permanent feature (not a bug) for Win10 installations.

version 4.0.0

- detectibility graphics plots are now being put in their own folder. The root folder is that of the directory where the light-curve came from. The plots will be found in `lightCurveDirectory\DetectibilityPlots\`

In addition, if you supply a duration step size, which is the way to ask the detectibility calculator to step from the given duration down until an event of that duration fails the False-positive test, only the final plot (where False-positive became non-zero) will be plotted.

All of the above is just to cut down on clutter in the light-curve directory and to only plot meaningful graphics generated during the minimum duration search.

version 3.9.9

- whew! 3.9.8 did solve the sizing issue. So it is safe to get back in the water so...

This version just improves the right-click context help for buttons in the new Detectibility calculator.

version 3.9.8

- one more attempt to solve the gui sizing issue.
- 'sticky' settings are back but are now stored in `pyote.ini` instead of the obscure `simple-ote.ini`

version 3.9.7

- Something inexplicable has happened regarding gui size. Suddenly, the gui size is too large for many people – but I have set the gui size smaller than it has ever been. So, best guess is that an overly large gui size got distributed once (there was a strange

event where an extremely large gui size appeared) and is being 'remembered' in the simple-ote.ini file that preserves the 'sticky' values like size and position of the gui and a few other things.

This version refuses to read any existing .ini files. Hopefully you will be able to run the program, but nothing will be 'sticky'.

If this 'cures' the problem, I can restore the use of the .ini file and we can get back the sticky stuff.

This version also removes the 'scrunch' of the lefthand button panel.

version 3.9.6

- another attempt fix the startup gui size issue on smaller legacy monitors. This version has a design size of 1900x1000. It may be necessary to delete the simple-ote.ini file that 'remembers' size and position settings. It is located in whatever directory PyOTE starts up in.

version 3.9.5

- adds automation to the 'detectability' calculator. If a duration step size is entered, a series of 'detectability' calculations will be made at decreasing durations until the false-positive probability becomes non-zero. A unique .png (incorporates duration and magDrop in the file name) will be written to the directory of the light-curve and the final plot where detection failed will be left on-screen.

If the duration step size is left unfilled or is zero, a single 'detectability' calculation will be made as before.

version 3.9.4

- attempts to fix over-size initial gui.

version 3.9.3

- makes various improvements to the 'detectability' calculator to make it easier to use.

version 3.9.2

- adds a 'detectability' calculator for use in pre-planning event observation.
- 'scrunched' the left panel buttons together to make more room for the user to move the horizontal splitter up (which gives more room to view the lower right panel where the report gets printed).

I think that this is kind of ugly, and it can be easily returned to the more spacious format if users would prefer the old look over more view space for the report.

version 3.9.1

- adds an snr calculation to the **Calc baseline stats** button.

version 3.9.0

- restored the right-click help to the **Mark baseline region** | **Clear baseline region** | **Calc baseline mean/sigma/corr coeffs** buttons that got lost during the scrollable gui experiment.

version 3.8.9

- Changed the label on the **Write csv file** button to **Save current light-curve** to more intuitively suggest what is hiding behind this button.

version 3.8.8

- fixed a bug that sometimes kept events that contained a single point at the bottom from producing a report.

version 3.8.7

- changed the way curve to analyze and normalization curve are selected to make it less confusing. Now, if there is a conflict (same curve selected for analysis and normalization) the other one is set to 0 (a new value) to 'get it out of the way'. It may still be a little bit confusing when you need to crossover, but if you set the higher numbered curve first, it will be easy.

version 3.8.6

- checks for only a single light-curve in the csv file which was causing a 'list index out of range' error message

version 3.8.5

- changed the way QGridLayout was referenced.

version 3.8.4

- changed the way QTableWidgetItem and QApplication were referenced to accommodate the latest Anaconda version, which has reorganized where those modules are stored. A similar change was made in version 3.7.6 and I do not know how this has gone uncorrected for so long – I hope that I have not missed something.

version 3.8.3

- adds display of light-curve name (if available in the csv file) to the curve-to-be-analyzed and the normalization curve next to their selection spinners.

version 3.8.2

- 3.7.5 introduced a scrollable GUI for small screens. Unfortunately, that caused resizing problem for all users. The version reverts back to the original GUI design.
- Added a print to the command window of the version number for diagnostic purposes. It will help clarify the situation when an installation has somehow managed to have multiple locations for PyOTE storage as to exactly what version is running.

version 3.8.1

- adds a checkbox to use reading number to annotate the x-axis, even when timestamps are available to use for this annotation. The setting is 'sticky', so your last preference will be remembered and used in the next run.

version 3.8.0

- adds a tutorial button to supply a quick-start for new or infrequent users.

version 3.7.9

- adds a tool that allows a user to manually specify which baseline points to use for the calculation of B (baseline mean) baseline noise (σ_B) and noise correlation coefficients. This was added primarily to support GPS-flash-tagged light-curves where the flash itself adversely affects the baseline values. With this new tool, the regions of the baseline that are outside the 'flash-zone' can be specified. Look for a button labelled **Mark baseline region** located near the bottom-left of the GUI panel.

This tool may also be useful for wind-gust situations, headlights from passing cars, bumping the telescope, etc. In the past, such accidents might have required extensive 'clipping' of the light-curve to avoid the affected areas, with a consequent loss of information. Hopefully this tool will alleviate those situations.

version 3.7.8

- added visible left, top, and bottom axes to main plot (with ticks to make easier to associate timestamp with point).

version 3.7.7

- if timestamps are available in the csv file, they are used as x axis labels in the main plot.

version 3.7.6

- changed the way QMainWindow and QApplication were referenced to accommodate the latest Anaconda version, which has reorganized where those modules are stored.

This was causing a fatal startup error.

version 3.7.5

- changes the routine that looks for the latest version of PyOTE to one provided by Kia Getrost. His version contacts the PyPI repository via a json query and is the officially supported way to get version info. My version was based on a 'hack' that depended on a special feature of pip (the loader that get programs from the PyPI repository) that was marked as 'unsupported'. That worked for many years until the pip programmers removed the 'special feature' as was their right (and the implicit promise/warning, I guess). This caused several users to always get a message that they were not running the most recent version of PyOTE, even though they were. Again: thanks to Kia Getrost for researching the problem and even supplying correct code (worked first time!) for me to use.
- The GUI has been changed so that scroll bars will appear when the standard GUI size exceeds your screen dimensions. THIS IS EXPERIMENTAL BECAUSE I DO NOT HAVE SUCH A SMALL SCREEN TO TEST WITH. If this does not work for you, please let me know immediately!

version 3.7.4

- fixed issue with 'save main plot' that caused the graphic to sometimes be written to a previous folder rather than the current folder containing the csv file (FYI: it was a 'feature' of the routine that saved this information that, in the interest of efficiency, it would sometimes not write a new folder name to the 'sticky file' until later – the fix was to force it to update the 'sticky file' on every change)

version 3.7.3

- adds code line needed by any Mac users that are running Big Sur

version 3.7.2

- fixes bug that caused all analysis attempts to 'stall'

version 3.7.1

- fixed a bug in the r limb angle entry

version 3.7.0

- adds right-click help to ast speed label giving the equation to use if asteroid speed is not available but asteroid diameter and maximum duration are specified
- adds suggestion to right-click help on penumbral fit checkbox for how to find the penumbral curve csv file that is provided for practice purposes.

version 3.6.8

- adds modeling of off centerline observations to lightcurve calibration curve generation

version 3.6.7

- clarified the location of the Enable Manual Timestamp Entry checkbox in the pop-up message appears when there are no timestamps in the csv file.

version 3.6.6

- modified the lightcurve demo to show more clearly the camera exposure function and the star intensity function that are convolved with the diffraction lightcurve to produce the lightcurve as seen by the camera.

The right-click help connected to the Demo button has been expanded to include a discussion of the convolution operation and hopefully provide some context to aid in understanding the star and camera function plots.

version 3.6.5

- added advisory message when False Positive probability plot appears in hopes of stemming in-discriminant use of this number as a 'decider' between a 'miss' and a 'positive' for an observation.

version 3.6.4

- version 3.6.2 would open .xlsx Report file on Mac, but not Windows. This version attempts to fix that (Windows needs a different command to open a file).

version 3.6.2

- In Excel report:

... at end of filling, I call the OS to open the file. Requires correct association of .xlsx file type with Excel or LibreOffice, or whatever you use to examine spreadsheets

... now writing numbers into numeric cells rather than text. This allows the cell formatting to control rounding, etc

... if I can't write to the selected file, I ask whether you might have it open somewhere else already

version 3.6.1

- reduced number of digits in the error bar numbers to 2 written to the Asteroid Occultation Report Form so that the resulting number fits within the allotted cell size. (I'm told that Occult only uses 2 digits anyway.) I also updated the context help associated with the ... fill Excel report button

version 3.6.0

- provides the ability to fill entries in the Excel spreadsheet Asteroid Occultation Report Form from PyOTE results, thus eliminating transcription errors.

A button to allow the user to activate this 'fill' has been added just to the right of



the ... write report button

NOTE: the normal .xls report form that OccultWatcher creates and prefills (when requested) during your report to OccultWatcher ...

... MUST BE CONVERTED to .xlsx for use by PyOTE!

... For Windows users, Excel will read an .xls file and save it as .xlsx

... For Mac/Linux users, LibreOffice will read an .xls file and save it as .xlsx

Sorry about this extra step, but it was necessary. The downstream tools used by area coordinators work equally well with the .xls and .xlsx forms of the spreadsheet, so there is no problem sending in the .xlsx version.

What gets filled in for you is:

...D/R uncorrected times

...D/R error bars

...SNR

...OTA used

...nominal magDrop (entered in Comments cell of the spreadsheet).

NOTE: you must still open the spreadsheet and enter the exposure setting used for your camera; it was not possible to do this from PyOTE.

After filling in the exposure setting in the spreadsheet, double check the form but it is likely that it is ready for submission.

version 3.5.9

- fixed another bug in penumbral curve fit and removed diagnostic printouts.

version 3.5.8

- fixed a number of bugs in the penumbral curve fit code

version 3.5.7

- automatically turns off the display of the normalizing lightcurve when a normalization is performed. It was a source of confusion to leave the normalizing lightcurve visible because it was sometimes the case that the normalization appeared not to have occurred (when the normalization effect was subtle/minor).

version 3.5.6

- fixed a bug that kept a user from selecting a new file to read if PyOTE had been started from PyMovie. Previously, it had only reopened the same file instead of giving the user a file select dialog to choose from.

version 3.5.5

- fixed a bug that kept occultations from being extracted from lightcurves 5 and up. The lightcurves above 4 could be viewed --- they just couldn't be processed through the event finder because they were a different type (coding error)

version 3.5.4

- made the 'get newest version' code identical to that in PyMovie in hopes that that will resolve the issue that some people experience with pip (or python) installing downloaded pyote in a directory where it subsequently cannot be discovered. The change is minute, so I'm not optimistic, but it's worth a shot.

version 3.5.3

- removed the blank lines between header lines extracted from the csv file and placed in the log file --- this makes it easier to look at the newly added aperture settings (so just a tiny cosmetic change).

version 3.5.2

- Changed the 'smooth reference star procedure' to no longer display the points at the left and right ends; such points are actually extrapolated points with all the hazards that extrapolation can engender. Smoothing functions that use sliding windows always have a problem at either the left edge, the right edge, or both

(when a symmetrical smoothing algorithm is employed). They run out of points and have to extrapolate/fake a number of points equal to the window size. Such extrapolated points can exhibit extreme behavior, zooming up or down unexpectedly and unrealistically.

Previously PyOTE treated this as a cosmetic problem and relied on the observer to be aware of the end point effects and ignore them. But that puts a burden on the user to be well informed about what's going on. As one of the goals of PyOTE is to enable infrequent/inexperienced users to get dependable results without requiring in-depth understanding of the internal workings of the program, we have decided to make the end-point smoothing issue very apparent by doing an automatic 'trim' of the points affected by extrapolation.

version 3.5.1

- Added additional references to the North American Excel Spreadsheet report in the new section of the final report that bangs on about the start-of-exposure timing convention.

version 3.5.0

- When PyMovie files are read, the aperture names are now being used in the data table (lower left panel) as column headings and used during the 'write csv' process. This makes the format of the PyOTE csv file match the PyMovie format so that AOTA can read both PyMovie csv files AND PyOTE csv files.
- NEW: when 'trims' have been placed, a 'write csv' process will honor those values and produce a 'trimmed' output file.
- NEW: when a light curve has been 'normalized', the changed values are written to the data table where, once again, a 'write csv' process will capture the results.
- Added additional reminders that the start-of-exposure timestamp/timing convention is employed.

version 3.4.9

- Added some explanatory language to the "Excel report" section regarding the proper interpretation of the 'false positive probability' number.

version 3.4.8

- This version deals more realistically with high magDrop lightcurves by defining a 'limiting magDrop' as:

$$\text{limMagDrop} = 2.5 * \log_{10}(B / \text{std}(A))$$

$\text{std}(A)$ is the noise level of A.

B = average baseline intensity

A = average event intensity

Normally, we report/calculate $\text{magDrop} = 2.5 * \log_{10}(B / A)$, but this calculation becomes

increasingly unreliable as the value for A gets very small. And when A is noisy, it is even possible to statistically have A become negative for large mag drop lightcurves. This happens more and more as A approaches and then becomes smaller than $\text{std}(A)$. For example, if A happened to be equal to $\text{std}(A)$, the normal distribution of A values tells us that 84% of possible A values are > 0 and so can be used in the regular magDrop equation. The other 16% of the time we can only report that the calculation could not be performed.

The above observation suggests that reliable estimates of magDrop require that A be greater than $\text{std}(A)$ --- that is the ad hoc reason that we have defined limMagDrop as we have.

This value is substituted for a calculated magDrop whenever A is less than $\text{std}(A)$, i.e., whenever A is at or below $\text{std}(A)$.

limMagDrop values are reported with a leading $>$ symbol to signify that that value is a limMagDrop value. They are easy to spot in the report.

version 3.4.7

- Automatically loads the correct version of Adv2

version 3.4.6

- Adds AAV Version 2 file as a type that can be read (important when Do OCR check is enabled)

version 3.4.5

- Fixes block integration which was failing when more than 4 lightcurves were being processed.
- Made use diff and Do OCR checkboxes sticky.

version 3.4.4

- PyMovie files can have lightcurves extracted from more than 4 apertures (with user supplied names). This version allows all lightcurves from PyMovie files to be read and made available for processing. Prior to this change, only the first 4 lightcurves were read.

Note: when you change the lightcurve to be analyzed (with the spinner), the log panel will show the aperture name for that lightcurve. That happens when the reference lightcurve is changed as well.

These changes are to PyMovie file treatment ONLY.

version 3.4.3

- Fixed bug that required an entry in $\text{dist}(\text{AU})$ and $\text{speed}(\text{km/sec})$ edit boxes for a solution to be found (the empty entries were causing an uncaught exception).

The intention is that PyOTE should work as it always did if a user ignores the new lightcurve parameter panel and makes no entries. This 'fix' was required to make that happen.

version 3.4.2

- Cosmetic change again: added a spinner to control line widths in plots so that a user can adjust for the resolution of the screen in use. I design on a 5120x2880 screen and needed lines to be 3 pixel wide to suit my taste. But some users have screens with 1280 horizontal resolution and those same 3 pixels became unsightly fat lines --- now there's a choice.

version 3.4.1

- Some cosmetic changes: thinner vertical lines for edge position and error bars; checkboxes to control whether the underlying lightcurve is plotted, error bars are plotted, or edges are plotted --- a cleaner plot is the major result and you have better control over the 'look'
- Added a checkbox to disable the automatic display of D and R frames from the video for OCR quality control checks. When there are no concerns about OCR reliability (true for me nearly all the time), it can be annoying to have to close the frames all the time.
- The BIG change is the addition of a penumbral curve fit procedure. It's a bit fiddly, so I included a test lightcurve with the download. I can't give you a specific location for the file because it depends on details of your particular installation. Find where it is by searching for example-penumbral.csv. When you find it, copy or move it to some other folder because if you process it where it resides, other files will get added in your installation directory --- we really don't want extraneous non-program files floating around in your installation directory.

To learn how to use the new procedure (which is a bit 'fiddly'), right-click on the penumbral fit checkbox --- be patient; play around.

version 3.4.0

- Adds the ability to specify a diffraction lightcurve for use in timing the event. A new panel with edit boxes for asteroid/occulting body distance (in AU --- astronomical units) and asteroid/shadow speed (in km per second) has been added. These values are needed in order to calculate a diffraction light curve.

In addition to modeling diffraction effects, one can add the effect of a finite star disk to produce a penumbral curve. NOTE: PyOTE does not yet have the ability to correctly analyze a penumbral curve where it takes more than 1 or 2 readings for the transition. That project is under way and will be in the next version.

version 3.3.9

- Automatically installs cv2 if not already present. This package is needed for the new frame view feature.

version 3.3.8

- If the video referenced in the csv file can be found, there is now an automatic display of the D and R frames relevant to calculating correct D and R times so

that the user can verify that timestamp OCR extracted the correct timestamp values.

version 3.3.7

- A new button (View frame) with an associated spinner for entry of a frame number has been added:

Use this button to view a frame from the video that was used by PyMovie or Limovie to prepare the .csv file that is currently being analyzed. .avi and .ser files are viewable in this manner as well as .fits files inside a FITS folder.

If this button is disabled, it is because the .csv file did not come from PyMovie or Limovie or simply cannot be found/opened.

This feature can/should be used as a final quality control check for a video that contains timestamps extracted using OCR. It is possible for OCR to fail in manner that is **not detected** by PyOTE because the program only verifies that there is a consistent step (delta time) between frames. If a high order digit in the timestamp has been consistently misread, substituting a 8 for a 9 in the minutes field for example, the steps can be consistent while the reported time of the event will be seriously in error.

ALL time reporting is derived from the timestamp(s) associated with D and/or R (the integer values, not the sub-frame values). If those timestamps are correct, the reported times will be correct even when there may be a few missing or duplicated frames. So best practice is to enter the D frame value in the spin box and visually confirm that the timestamp that you can see is the same as that extracted by the OCR procedure. Repeat for R.

Another use for this feature is to handle the case where there is a visual timestamp, but either OCR was not activated during the .csv preparation, or the timestamp overlay came from an unsupported VTI type. The workflow would be to let PyOTE find the D and R frame values, but before pressing **... write report**, do a **Manual timestamp entry** for the D and R frame entries found by viewing the relevant frames and entering the correct times in the Manual timestamp dialog.

It should be noted that the manual timestamp entry can be performed even when timestamps were already present in the file --- your manual entries will cause all the timestamps to be recalculated.

version 3.3.5:

- Changed usage of `max([a, b, c])` to `max(a, b, c)` to see if this allows the Numba JIT compiler to work for one user that found version 3.3.4 failed to load/compile.

This should have no effect on users that already have version 3.3.4 working.

version 3.3.4:

- To shorten the time to find 'solutions', I used the Numba JIT (just-in-time) platform independent compiler that produces machine code from Python byte-code. You may notice a very slight increase in the time to start-up PyOTE because I do those compile operations while PyOTE is being loaded.

version 3.3.2:

- Adds a false-positive probability calculation and printout in the final report. This number is the fraction of 'false drops' found during the 50,000 tests that are greater than or equal to the drop value extracted from the actual observation.

version 3.3.1:

- Adds a 'false positive' detection to the final report. A new plot has been added to the error bar plot. It shows the distribution of drop sizes (B-A) for an event of the size (duration) extracted from the actual observation, but with only correlated noise in the sample (the number of points in this sample is equal to the number of points used in the lightcurve extraction). 50,000 attempts are made to find the deepest event that appears (falsely) when there is only noise being analyzed. If the drop from the actual observation is greater than the maximum size of a 'false drop', we have some assurance that the event extracted from the actual observation did not happen 'by chance'.

version 3.2.9

- Changed main plot so that the scroll wheel only zooms the x axis.
- Changed lightcurve plot so that it conforms properly to 'start-of-exposure'.

version 3.2.8

- Changed font size in help files --- it was fine for Mac but too big for Win10

version 3.2.7

- Removed the 'hover-for-help' and replaced it with a 'right-click-on-item' to get help. This scheme was introduced in PyMovie and I found it easier to use than the 'hover' scheme. In practice, the 'hover' popped up when it was not needed, so most users eventually disabled it. As a result, it became so tedious to look at help --- enable hover; hover; read; disable hover --- that the help system was used less and less. The right-click-for-help is always available and easily invoked --- hopefully this will encourage more frequent reference to it.

version 3.2.6

- This is a 'cosmetic' release --- there should be NO detectable differences from version 3.2.5 in terms of functionality.
- All python files were brought into compliance with PEP 8 coding standards. Only I care about that.
- More significantly, I removed the dependency on C code by using Numba as a code accelerator instead of Cython. As a result, I no longer need to compile separate code versions for Mac, Windows, and Linux. That makes my life easier, but you should experience no operational changes.
- All this 'cosmetic' work is in preparation for working on PyOTE issues again.

version 3.2.5

- Added special test for Tangra files to detect the empty fields (which MUST be fixed) that Tangra outputs whenever it has trouble extracted a value from an aperture. It prints a message and stops all further processing, forcing the user

to attend to and deal with the missing values.

version 3.2.4

- Modified the test for newer version to accommodate the different strings returned by pip 18.1 and pip 19.0+
- Added ability to invoke PyOTE from PyMovie with an externally supplied csv file that is automatically opened.

version 3.2.3

- fixed a long overlooked bug in the loading of the data table (at lower left corner of GUI): when there are four lightcurves, LC4 was set in the table from LC3 (i.e., $LC3 == LC4$ whenever there was an actual LC4. It was correct in the lightcurves themselves, so no observation analyses have been affected by this bug. It was cosmetic only.
- Added support for the PyMovie csv format

version 3.2.1

- this version makes PyOTE more robust to a common 'cockpit error' that users have been making with Tangra files. Specifically, if a Tangra csv file is opened in a spreadsheet program, then saved from that spreadsheet program, the original csv file gets modified and overwritten by the addition of empty fields at every row sufficient to match the number of columns in the longest header/comment row --- the spreadsheet program did this to satisfy its internal requirement that every row have an equal number of columns. The result is superfluous commas at the end of data lines (that Tangra did NOT put there). Until this version, that 'butchered' file could not be read. This version adds code to parse data lines only up to the first non-empty column. Hopefully this will not have ramifications in the future (like a format change that has empty fields followed by non-empty fields --- not a likely expectation).

version 3.2.0

- Changed GUI to better align text on min max edit boxes to avoid confusion.

version 3.1.9

- Fixed a bug in the test for a min/max solution search being constrained by a too large min value.

version 3.1.8

- version 3.1.7 was released without an updated version history. Here is what was changed in 3.1.7:
- Added the ability to write the data table that is displayed in the lower left corner of the GUI out as a csv file. Now, if timestamps and block integration operations are performed on the input file, those results can be preserved in a csv file.

A 'file save' dialog is provided should you wish to change the default name and location of the resulting file. The default name is that of the input file with the

text .PYOTE inserted to the left of the .csv extension. The default location is the directory of the input file. It is recommended that you accept these defaults unless you have compelling reasons to do otherwise.

version 3.1.6:

- Values entered in the Manual Timestamp Entry dialog box are now 'sticky', thus making corrections easy to do without requiring re-entry of all data.

Also trapped is the case where a user has entered a custom frame time but failed to click the radio button indicating that it is to be used.

version 3.1.5:

- Added additional tests of candidate solutions against a straight line so that there should always be agreement between a solution found by a min/max event size search and a marked D and R region search of the same area.

Previously it was possible for the min/max search, which searches the entire light curve, to be tripped up by what we call a 'competitor'. A 'competitor' is an 'event' with good statistics. However, that 'competitor' may have a small magDrop and so later be rejected when we compare with a straight line solution. That 'competitor' would thus mask an event with slightly worse statistics but a larger magDrop. The change was to test every candidate against a straight line during the search. This does make the search time longer, but not too much longer.

version 3.1.4:

- Fixed error in new dropped reading detection logic when light curve was processed in field mode.
- Cleaned up some language in tooltips.

version 3.1.3:

- Expanded manual timestamp preset time deltas to include NTSC and PAL field times. Also added ability to evaluate numeric expressions entered in the 'Custom time' box: now you can type $1.001/60.0$ in that box if you wish.
- Eliminated the 'entry num' column in the data matrix at the lower left of the GUI. The 'entry num' is unused and a possible source of confusion with the frame or field number for the unwary.
- Added **all** the light curves read from the input file to the data matrix display. Previously, only the first light curve values were displayed. This is done in anticipation of adding a 'write csv' button to memorialize the result of a manual timestamp entry.

version 3.1.2:

- Added a test for possible dropped frames identical to that done in R-OTE when manual timestamp is utilized. The test is to calculate the expected number of

frames based on standard NTSC/PAL frame times and compare that number with the count of frames enclosed by the early and late timestamps. If there is a mismatch of more than 0.12 frames, a warning is popped up and a log entry made. It is possible to use a 'custom' frame time if your camera differs from either of those standards.

version 3.1.1:

- A convenient way to search for a 'solution' is to set a min and max event size rather than mark D and R regions. This is particularly useful in low snr situations where the D and R edges may be quite diffuse. However, if one sets the min event too large or the max event too small, the resulting 'solution' will be artificially constrained and thus be wrong. This situation is now detected and a log entry as well as a pop-up alert will tell the user to change the limits and try again.
- Three magDrop values are now calculated for each confidence level: the largest magDrop calculated using $B + \text{err}(B)$ along with $A - \text{err}(A)$; the nominal magDrop calculated using B and A; the minimum magDrop calculated using $B - \text{err}(b)$ along with $A + \text{err}(A)$
- The labels on the Find Event button and the Calc Err Bar button were changed to more clearly suggest that after finding an 'event', one should then press the 'report' button to the right in order to complete the process.

version 3.1.0:

- Added a Mac version of a pyote startup file. It is automatically placed on the Desktop the first time pyote is run. Double-clicking on that Desktop file icon will start pyote thereafter.

version 3.0.8:

- Added a Windows batch file to the distribution that, when executed, will startup pyote. The file is called PYOTE.bat and is automatically copied to C:\Anaconda3 (if it is not already there) when pyote is first run. Now, to create a clickable desktop icon for starting up pyote, a user need only go to the C:\Anaconda3 directory, locate the PYOTE.bat file, create a shortcut to it, and drag the shortcut to the desktop. Remember, that file does not appear until the first run of pyote.

The 'skipped' version numbers were caused by the need for repeated testing of this new feature, each test requiring a new version, even though functionality did not change,

version 3.0.1:

- Restored the vertical splitter between the command/plot area and the table/report area. Somewhere along the line this capability was accidentally removed, and the lack of the splitter was not noticed. Now it's back.

version 3.0.0:

- No code changes. This version is the same as 2.1.6 except that it is built on python 3.7. The previous versions used python 3.6. This allows new users to install the latest Anaconda3 version (which installs python 3.7) without fiddling with archived Anaconda3 versions.

version 2.1.6:

- We now disable the Accept integration button on the first left click in the light curve. As such a click removes the color bars that result from the automatic block integration analysis, it seems intuitive to disable the Accept integration button at that time as well.

version 2.1.5:

- Disable the Accept integration button when user overrides an automatic block analysis with a manual block selection followed by a click on the Block integrate button.

version 2.1.4:

- Corrected a bug that kept manual selection of block integration from being performed after a refusal to accept the automatic block analysis results.

version 2.1.3:

- A minor change to how color bars are plotted when the automatic block integration feature is employed. The edges now appear between data points so the bands are easier to see, particularly for 2 point block sizes.

version 2.1.2:

- To ease usage of the automatic block integration feature, accepting the automatically determined block integration parameters no longer uses a modal query box, which interfered with the ability to explore/expand the light curve plot. Now there is separate button which gets enabled after an automatic block integration completes.

version 2.1.1

- Added progress bar tracking of block integration analysis because it can take an extended amount of time to complete the analysis when the light curve has many points.

version 2.1.0

- Added automatic determination of 'correct' block size and offset for block integration when user clicks **Block integrate** button without selecting the two points normally required to specify integration block beginning and end. The user can choose to accept or reject **pyote**'s opinion of the correct parameters to use when the automation determination is invoked.

version 2.0.9

- Made the selection of Tooltip display 'sticky'
- Duration calculation when D and R span midnight now handled correctly

version 2.0.8

- toolTips changed to invoke and display in a custom dialog box that can be moved and resized to better accommodate legacy displays
- Calc flash timing calculation fixed to properly deal with the non-integer frame numbers that can result from field processed csv files
- Flash timing has been verified to work with integrated light curves
- Made block integration 'sticky' in that a 'Start over' no longer undoes a previous block integration. As a result, once block integration has been performed after a file read, it cannot be done again; a reread of the original file is now required.

version 2.0.7

This version provides several features to ease the processing of light curves that are timed with LED flashes from iPhones (John Grismore's AstroFlashTimer) or Android phones (Eric Couto's Occult Flash) rather than VTI timestamped files

- Adds a button to calculate the edge position of an LED timing flash.
- Adds a checkbox to enable/disable the tooltip messages that appear when a control is hovered over. Tooltip display defaults to 'enabled' because tooltips are an important aid for guiding users initially. Later, when such help is no longer needed, the user can turn them off (they are annoying when you don't need them).
- Adds the ability to select which light curve is to be analyzed. Previous versions would only analyze the first light curve for D and R events. This flexibility is useful in general, but was particularly needed to support LED flash timing.
- Adds a checkbox to force manual entry of timestamp info. This is useful when OCR on a VTI timed light curve has catastrophic errors. It is always employed when using LED flash timing.
- During the error bar calculation, it is possible for the Cholesky decomposition needed for treating correlated noise to fail. Previous versions treated this as a fatal error and would not produce a final report. This version instead treats the noise as uncorrelated and continues processing to produce a final report.

version 2.0.6

- Added additional instruction in the popup that appears when no timestamps are found in the csv file. This will give casual users additional guidance and clarification for the manual timestamp entry process.

version 2.0.5

- files generated by pyote now contain PYOTE in the filename.
- Timestamps can be corrupted to the point that a timeDelta of 0.0 can result. This version traps that event and reports it clearly --- 2.0.4 failed silently with a divide by zero exception

version 2.0.4

- improves the handling of errors during the reading of Tangra files by showing the offending line in the report panel. Tangra, if it has a tracking problem (i.e., loses it) will emit an empty field for that measurement, leaving it up to the user

to decide how to fill in the missing value. Prior pyote versions simply reported 'format error' without providing a printout of the offending line. This version fixes that.

version 2.0.3

- detects and handles situations in which fewer than 14 baseline points are available for calculation of correlated noise coefficients. When fewer than 14 points are available, the correlation coefficients are set to: [1, 0, ...] (i.e., coefficients are set to 'no correlated noise')

version 2.0.2

- Note: this version has many significant changes. If you lose confidence in this version, remember that you can always go back to version 1.47 by typing ---

```
pip install pyote==1.47
```

in an Anaconda console. (Be sure to use double == signs in the command.)

- improved handling of D and R region selection so that one cannot enter an invalid configuration --- automatic corrections/changes are applied.
- incorporates a new 'solver' that no longer requires an initial estimation of baseline noise. This 'solver' is also much faster. With this 'solver', the two-pass modification added in version 1.46 is no longer needed.
- removes unneeded 'analyze noise' buttons and rearranged other buttons to be in-line rather than one above the other to allow the vertical splitter between the plot area and report area more room to change (a help to those using screens with relatively low pixel densities).

version 1.47

- adds bold red highlighting to message:

! There is something wrong with timestamps at D and/or R or frames have been dropped !
so that it is harder to miss.

version 1.46

- adds automatic recalculation of baseline and event noise parameters utilizing all available data points during a second solution pass; this removes the variability in calculated error bars due to user selection of a necessarily less complete set of data points for noise analysis during the first solution pass.
- adds bold blue text in the 'Excel' portion of the final report to indicate whether or not the light curve was block integrated, trimmed, or normalized. Failing to block integrate a light curve that needed it is a common error. Highlighting the presence or absence of block integration in the most looked at portion of the final report will hopefully help reduce the number of such errors.

Version 1.45

- the initial fully functional release of pyote.

Introduction to *pyote*
Bob Anderson (bob.anderson.ok@gmail.com)

pyote is an occultation timing extraction utility program written primarily in python and distributed through PyPI (the python package repository).

This program is specifically designed for those who will use such a program infrequently; it has been designed to the best of my ability to produce consistent results in the hands of both infrequent and frequent users --- the same results should be obtained no matter who processed the data.

One important feature of the program intended to give confidence to the occasional user is the production of a log file that documents all processing steps/decisions made in sufficient detail that anyones result can be reviewed by more experienced users easily --- it is sufficient to simply send such a reviewer just two things: the light curve and the log file.

1. *pyote* is designed for ease-of-use in the analysis of occultation light curves that can be modeled reasonably well with a model based on geometrical optics. Such light curves are common with star/asteroid occultations when the star is effectively a point source and the asteroid transit speed is such that diffraction effects are masked by the natural integration effect of the camera operation coupled with the frame rate of the video recorder.
2. Correlated noise caused by atmospheric scintillation is frequently present in occultation observations recorded at normal video rates of 25 or 30 frames per second. *pyote* utilizes statistically rigorous calculations to properly characterize the increased uncertainty in D/R time estimates due to such correlated noise. Additional noise correlation is often present due to the relatively slow electronics present in low-cost frame grabbers and in the camera electronics responsible for generating the composite video output.
3. Physically realistic models (but based on geometrical optics) are fit to the light curves with all decisions about details (complexity) of the model used made using the Akaike Information Criterion (AIC). In particular, an AIC calculation is always used to justify or reject sub-frame timing.
4. Maximum Likelihood Estimation is used throughout to determine 'best fit' of model light curves to the actual data.

The gui for *pyote* is designed to lead the user through the necessary steps by enabling the buttons in sequence as each task is performed. So, initially, only two principal buttons are enabled: the 'info' button that brought up this document and the 'Read light curve' button. After reading this document, open a light curve, and follow the enabled buttons.

All of the major buttons have hover text associated. To learn (or refresh) how to use the program to analyze a light curve, spending a little time 'hovering' on the buttons will pay dividends.

pyote will never change the input light curve, so experimentation is encouraged. There is a 'Start Over' button at the bottom that I encourage you to use freely.

Every step you make in the analysis is recorded in a log file. This is done because experience has shown that some light curves are touchy to analyze and it is useful to ask someone more experienced in running the program to look over your work. With the original light curve and a copy of the log file, your work can be exactly duplicated by someone else. And that log file is never deleted once it is opened for a particular light curve; it is simply appended to, so a record of each 'experiment' is thus always available.

Every time *pyote* is started, it connects to PyPI (assuming you have an internet connection) and checks to see if a more recent version of *pyote* has been added to the repository. If your



version is completely up-to-date, you will see this



in the log file panel in the lower right-hand corner of the gui. Otherwise, this will appear: Normally, you will want to click 'yes'. That will cause your current version of *pyote* to install (but not run) the newest version. Of course, to execute that new version, you will need to do a close and reopen.

As convenient as this is, there is always a small risk that a new version will actually 'break' something and that the 'cure' may take some time to be posted. But it is always possible to return to a specific previous version of *pyote*. The procedure to do this is explained below.

Open an Anaconda Prompt window if you are running Windows.

For a Mac installation, open a command window and type ***source activate***.

Then, type the following line in that command window:

1. `pip install pyote==1.42`

This command will uninstall the current (flawed) version of *pyote* and installs a specific version, in this case, version 1.42. Note the double == followed by the specific version number to be installed. (You can always determine a version of *pyote* that was working for you by opening a recent log file --- the *pyote* version that produced that log file is recorded there.)

