The Consecutive Signed-Slope (CSS) method

Last modified by F. Masci, April 27, 2021

Method:

- Register and interpolate a sequence of image exposures (epochal images) onto a common pixel grid to construct an image cube.
- Consider the flux time series (lightcurve) for a single pixel in this grid from the sequence of *time-ordered* images.
- Optionally apply a median filter to this pixel time series to suppress noise between consecutive time-points. The window size depends on the timescale of variability sought and the observing-cadence of your imaging. The temporal window size can be fixed given knowledge of the specific survey cadence and science case. The goal (see below) is to make this window choice "science agnostic".
- Compute the pairwise first-differenced time series for this pixel ($D_t = f_{t+1} - f_t$ for all t = 0, 1, 2,...).
- Find the **longest** run (count) of **consecutive** pairwise slopes that are either positive OR negative ($D_t >$ or $D_t < 0$ respectively).
- Generate an image of these longest runs from "collapsing" all pixel-based time series: a positive count implies a time period where the pixel flux is growing in flux; a negative count implies a time period where the flux is decaying.
- Apply a 2D spatial sum-filter to the *absolute* values of the *longest runs* image to accentuate spatially-correlated neighboring pixels with either large negative or large positive run lengths; these would be candidates for variable or transient sources, i.e., which may have exhibited a systematic trend in flux anywhere in the time-ordered image stack.
- Using the background distribution in pixel values in the 2D spatial sum-filtered image from the previous step (which includes pixels falling on sources with significantly non-zero flux and where the majority are likely to be non-variable relative to the original detector noise), compute a threshold above which a pixel or cluster thereof is likely to indicate variability or transient behavior. In other words, this threshold is derived empirically from the same data by histogramming the pixels and picking a quantile from a *Pvalue* of interest (e.g., < 0.1% or < 1% chance occurrence). See below for examples. Computing thresholds empirically will implicitly include any systematics tied to your specific field of interest following all processing to mitigate them. I.e., the *null distribution* used to define your detection process must include all possible sources of error. This means the associated *Pvalues* are strictly for your specific field and must not be interpreted in an absolute sense for applicability to the entire survey.
- Using the threshold derived in the previous step, execute a peak detector on the 2D spatial sum-filtered run length image. One may impose additional constraints to make use of spatial correlations in *run length* and hence maximize reliability, e.g., by requiring at least *N* pixels to be above the threshold within a local 3x3 pixel region. This avoids single pixel outliers that are unlikely to be associated with real source signal unless the input image pixels grossly under-sample the PSF of the optical system with which the data were acquired.

Expectation from pure Gaussian noise:

- Here I pose the question: given a pure-noise time series of length N, what is the probability of obtaining a consecutive run length (in either positive or negative pairwise slopes) >= some value NC? In other words, how likely is it to obtain a run length >= NC by chance due to noise alone?
- These predictions will guide us in determining the significance of a specific run-length in a flux time series **under the *null hypothesis of pure Gaussian noise.***
- It's important to note that fluctuations in a real flux time-series' on a presumed-"static" source background are not guaranteed to be Gaussian. Systematics will be present: e.g., drifts in pointing and hence flux-interpolation accuracy when interpolating all image epochs onto a common grid; flux calibration (photometric throughput) biases across epochs; complex varying backgrounds and inaccurate subtraction thereof; PSF variations (spatial and temporal); source blending bias; intra-pixel responsivity (relative source pixel-position or phase), etc. Hence, the real "null hypothesis" must include all these systematics, which presumably, will be present in all other sources that you are comparing against in your field. **Therefore in practice, measurements of your local background population of "static" sources will define your true null hypothesis.**
- I performed a Monte-Carlo simulation to explore the pure Gaussian noise null hypothesis. This consisted of simulating 5000 time series trials for each length N and counting the number consecutive longest run lengths above each NC value.
- Results are shown in the attached plot. In summary, it's unlikely (Pr < 5%) to get a run of >= 6 consecutive same-sign slopes in a pure-noise time series of length N <= 100. For N <= 400, it's even more unlikely (Pr <~ 1%) to get a run of >= 7 consecutive slopes with the same sign.
- As mentioned, this exercise is purely academic. The author of this document does not recommend using this *Gaussian* prediction to assess the significance of a *run length* in real data unless it can be shown that measurement errors are *dominated* by a *random Gaussian process*. By "random", we mean where the errors associated with the individual image epochs, following interpolation are *independent*. An example is when one is dominated by Poisson noise from the original photon collection process that's also well into the high-count (Gaussian) limit. Systematics from any post-processing (e.g., astrometric calibration, registration, photometric throughput matching across epochs) are assumed negligible.

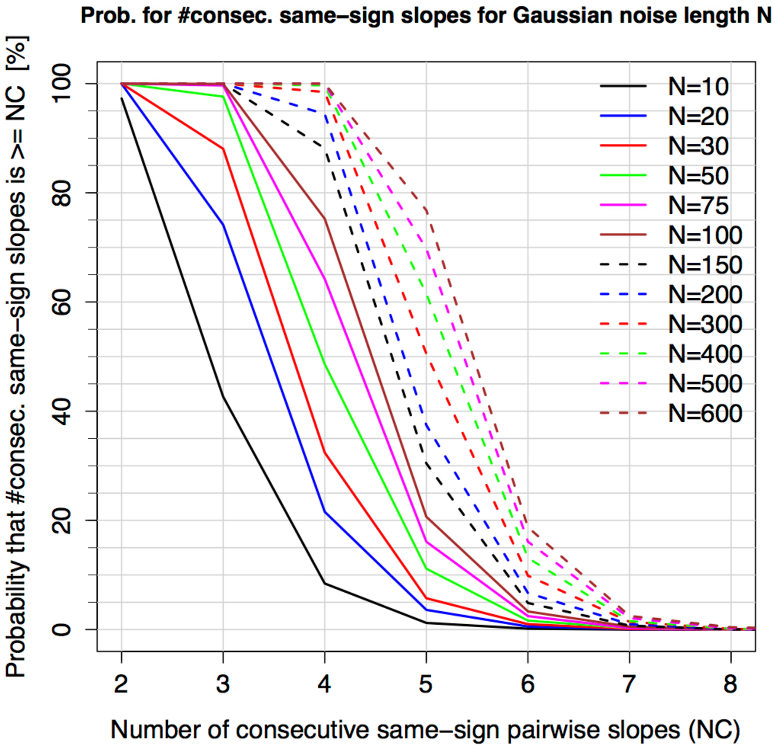

Putting it into practice:

The above method was implemented in a Python script (contact F. Masci for code), then executed on two cases from the WISE survey.

CASE 1:

The first case was on a stack consisting of 608 epochal images in the W1 band on the "fireworks galaxy", NGC6946. This galaxy had a Type IIP supernova (peaking around early June 2017). Shown below are the epochal image with the supernova at its peak (circled green); the longest same-slope-sign run-length image; the 2D spatially sum-filtered image; lightcurve. The supernova is clearly discerned above the background.

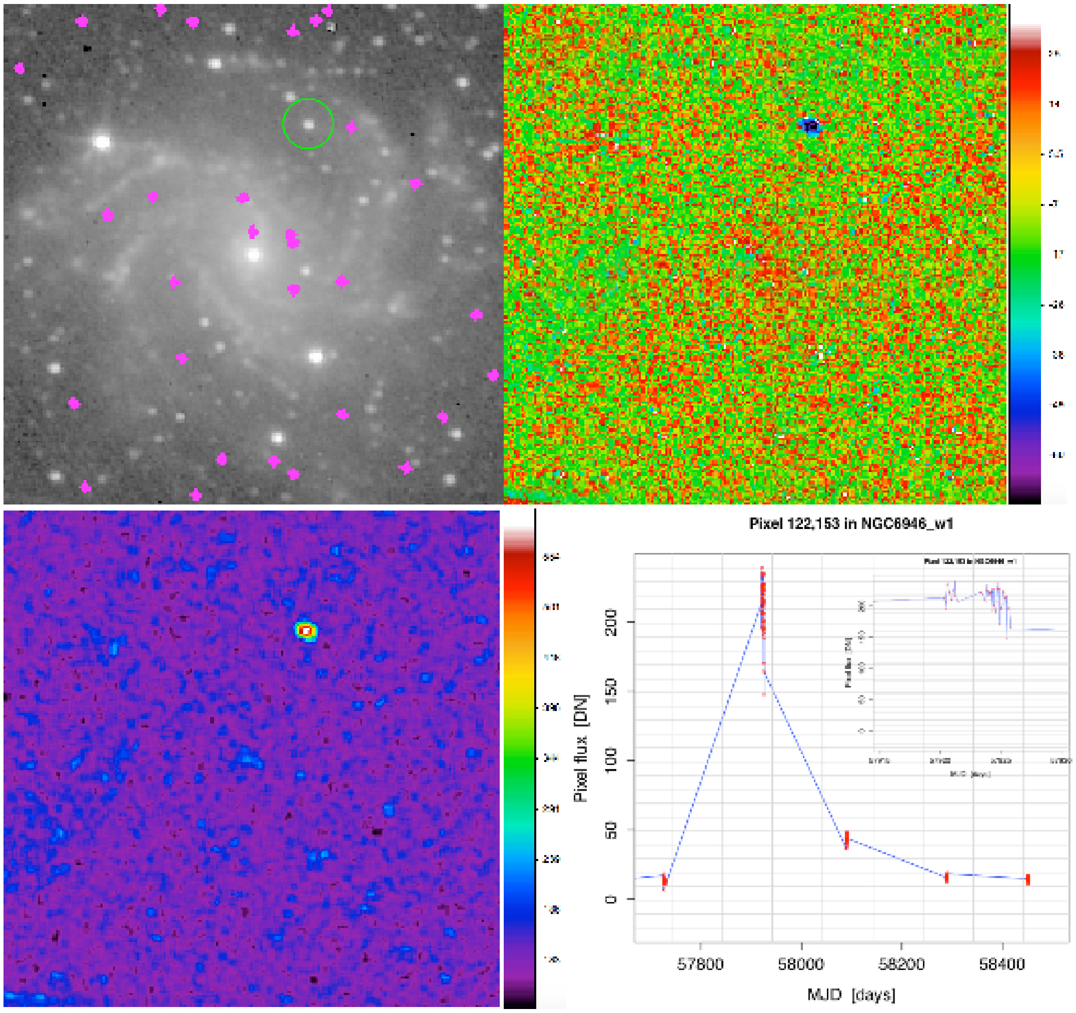

Associated with the NGC6946 test field above are four additional files:

- A [2D histogram](#) of *median-collapsed* pixel values in the stack versus pixel value in the spatial sum-filtered *run length* image. The vertical magenta line corresponds to the 0.5% upper quantile on the X-axis values alone (same as the magenta line in plot 2. below):

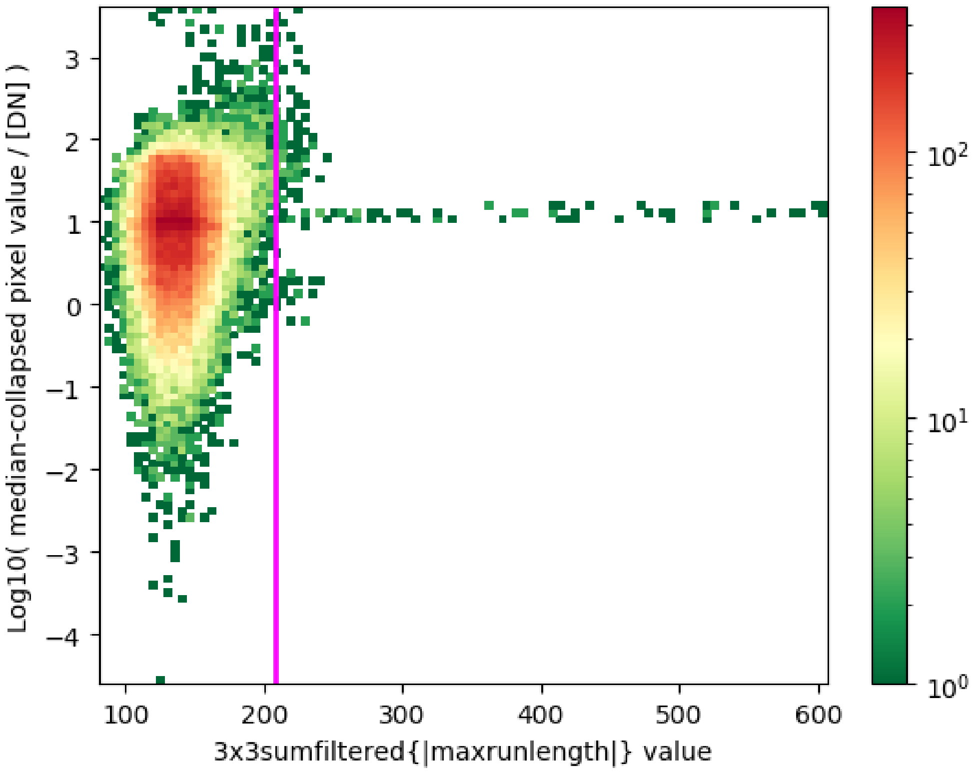

- A 1D histogram of pixel values in the spatial sum-filtered *run length* image with a few quantiles (possible thresholds to use) corresponding to *Pvalues* = 0.5, 1, 3, 5%:

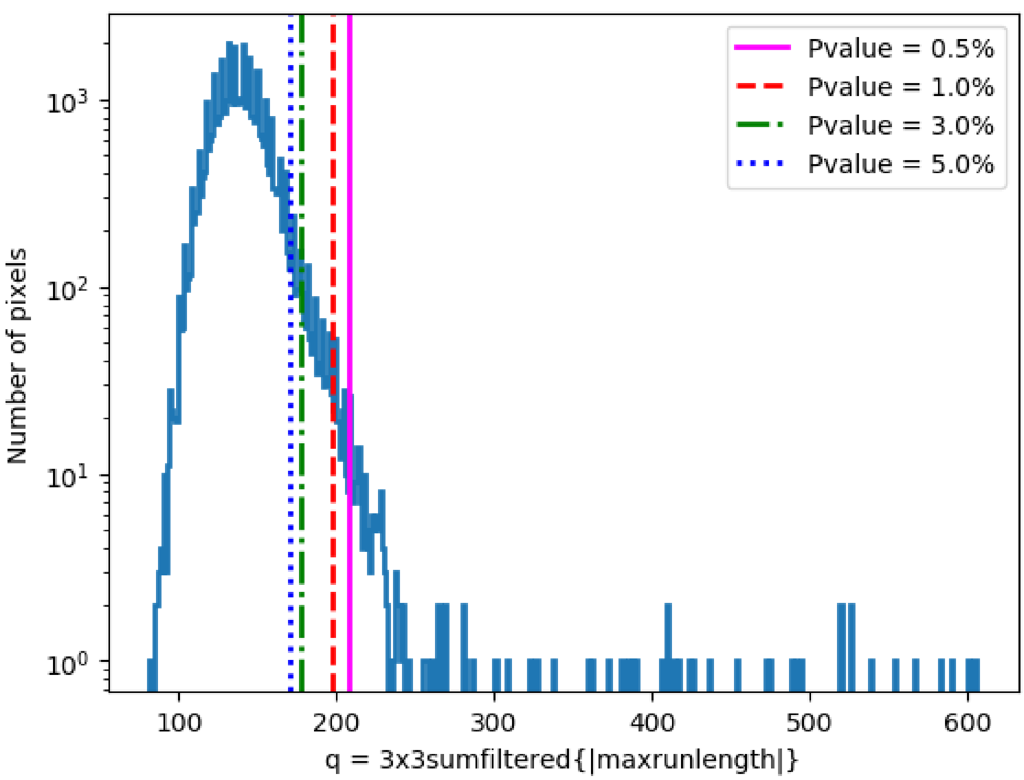

- A text file listing diagnostic information for peaks (local maxima on clustered pixels) detected above the threshold with *q_value* = 208, corresponding to *Pvalues* <= 0.5%. This includes a criterion where 6 or more pixels within a 3x3 region centered on each peak are also above the threshold. The rows in this file are in order of decreasing peak-pixel significance. A measure of the corresponding *N*sigma* deviation from the mode of the histogram is also given, where *sigma* is a robust estimate using the high-tail of the 1D histogram in 2.
- A DS9 region file (based J2000 RA, Dec coordinates) corresponding to the above text file to assist users when overlaying detections on any archived image data.

The following image shows the locations (green crosses) of the significant peaks (local pixel maxima) from the output text file overlaid on the spatial sum-filtered *run length* image.

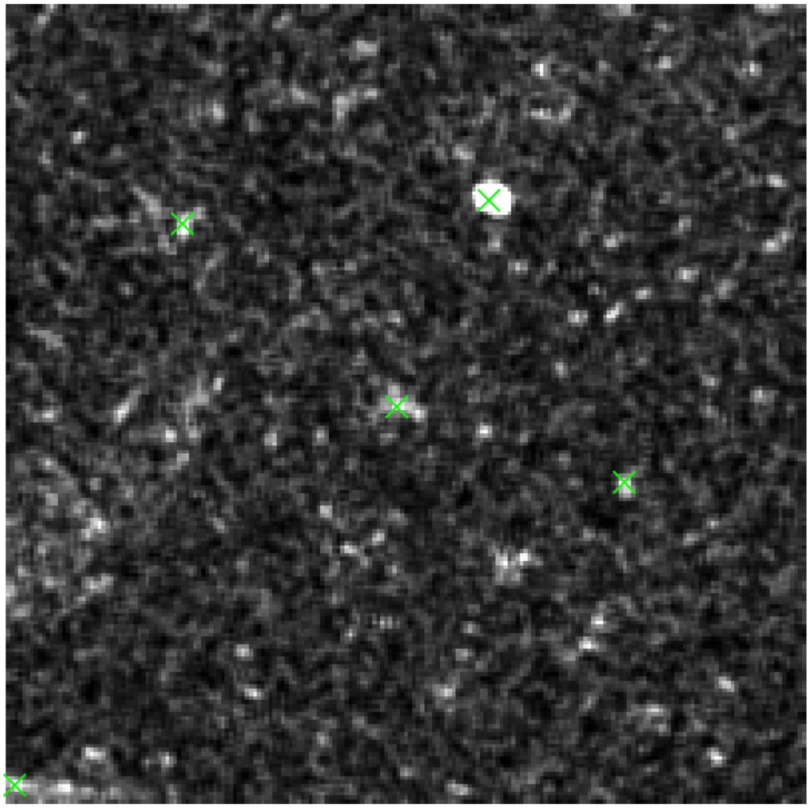

CASE 2:

- The second case was on a stack consisting of 174 epochal images in the W1 band on the RR Lyra variable, AV Men. This variable has a period of ~ 0.55 days. Shown below are the median of all epochal images with the variable circled green; the longest same-slope-sign run-length image; the 2D spatially sum-filtered image; lightcurve. Although not as significant as the supernova above, this variable still shows up as the strongest spatially-correlated peak in the images.

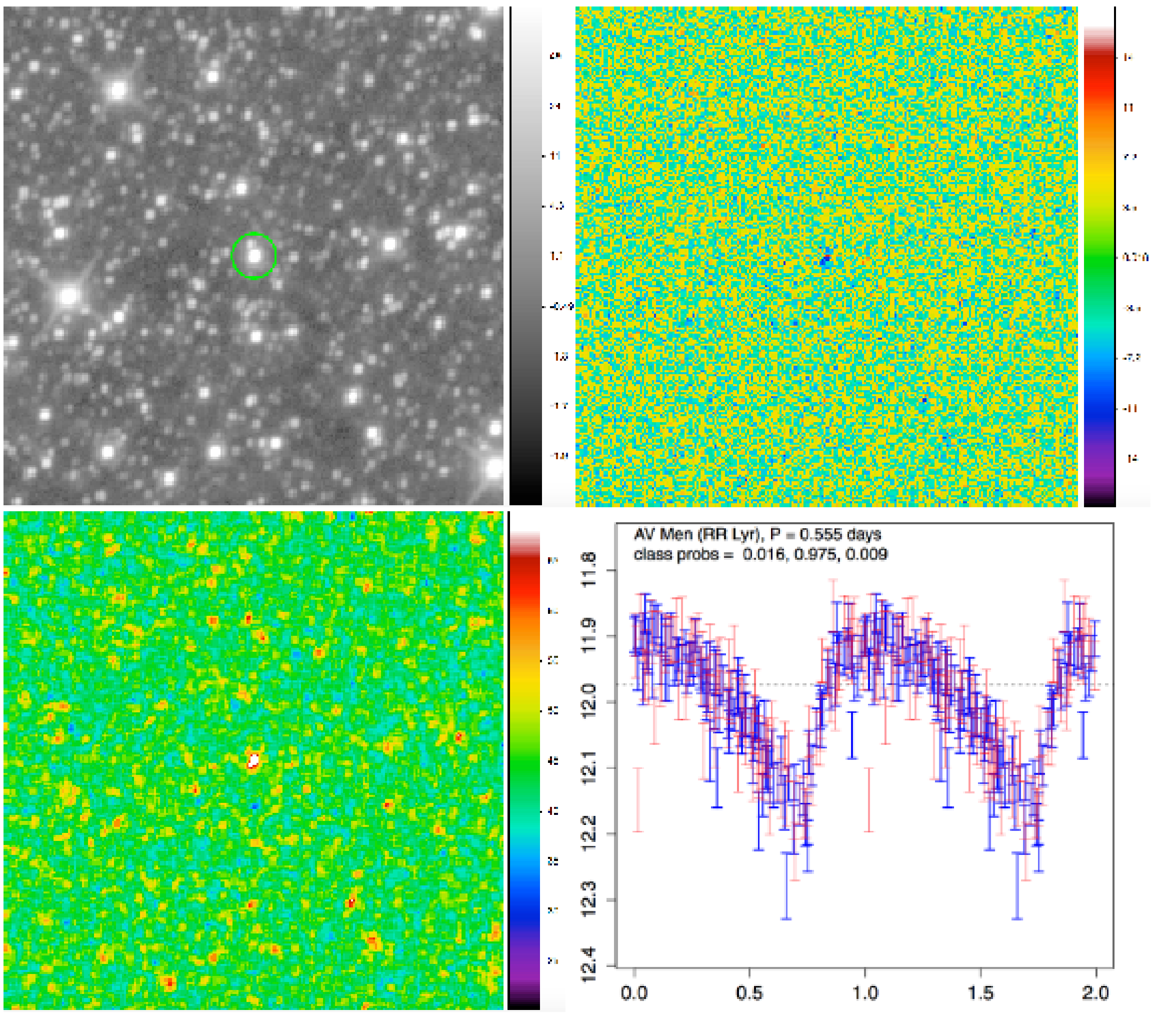

Further exploration and considerations:

- How science-agnostic is this method? What are the limitations on the types of variability (fast vs slow) this method is sensitive to for a given survey cadence?
- Can we optimize and parameterize the method according to the flavor of variability sought (e.g., reoccurring/periodic variables, fast/slow transients)?