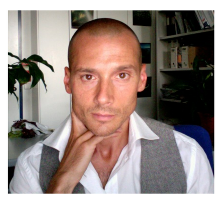

On Detecting Transient Phenomena

A Unified Approach

G. Belanger

European Space Astronomy Centre, ESA, Spain

Abstract

Transient phenomena are interesting and potentially highly revealing of details that could otherwise go unnoticed. It is therefore important to maximize the sensitivity of the method used to identify such events. We present a general procedure based on the use of the likelihood function for identifying transients that is particularly suited for real-time applications, because it requires no grouping or pre-processing of the data. The method makes use of all the information that is available in the data throughout the statistical decision making process, and is suitable for a wide range of applications. We consider those most common in astrophysics which involve working with images, time series, energy spectra, and power spectra, and demonstrate the use of the method in the case of a weak X-ray flare in a time series and a short-lived QPO in a power spectrum. (Work published in Belanger 2013, ApJ, 773, 66)

Detection and Identification

The detection of any kind of transient involves identifying something that was not there before. All transient phenomena share in common, independently of their particular time scale, that they appear and fade away.

The process of detection and identification of a transient feature in a set of measurements is a statistical procedure that involves comparing numbers and making decisions based on the probability ratios, and in other words, on *likelihoods, which measure statistical evidence.* We always want to maximise sensitivity to transients and minimise the frequency of false detections. Therefore, we must use all the information that is available, and interpret the data as statistical evidence.

A transient is identified in relation to background conditions. The background process can be constant or it can be variable. Our basic working assumption here is that we are dealing with cases where the background is constant on the time scales that are relevant to the problem of identifying transients, whether it is nil or not.

A Unified Approach

The method is straight forward and is based on the use of the likelihood function. All details pertaining to the inherent statistical properties of the observed random variable (e.g., normal, Poisson, exponential) are automatically taken into account and integrated in every aspect of the procedure, which makes use of each measurement and does not require any kind of grouping or approximations. The approach is unified because it treats in the same manner the detection of transients in any domain (time, space, frequency, etc).

The very first measurement gives the first estimate of the reference value: the value we expect to measure under usual conditions when there are no transients. With this we draw the curve that expresses the likelihood of all possible reference values given the evidence from that measurement informed by the probability distribution of the measured variable (Fig 1, panel a). Most common are the normal seen in continuous processes, and the Poisson when counting events.

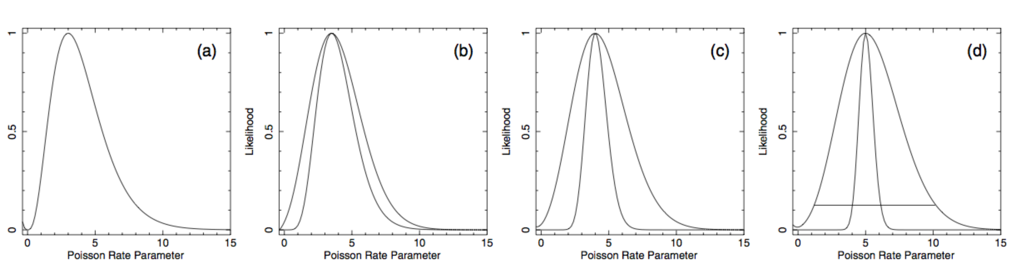

Figure 1. Graphical representation of the single-measurement and joint likelihood functions after 1 (panel (a)), 2 (panel (b)), 7 (panel (c)) and 19 (panel (d)) measurements of a Poisson variable with a true rate parameter of ν = 5. The 1/8 likelihood interval is shown in panel (d).

The second measurement gives a second estimate of the reference value. It is evaluated for its potential of being a transient by computing its likelihood ratio with respect to the previous ML. This is the only mathematically correct way to evaluate the likelihood of measuring that second value in light of the first. If it is not statistically different from the first beyond the established threshold, they are combined to better estimate the reference. The joint likelihood function is computed from the two measurements, and immediately begins to grow narrower (Fig 1, panel b).

With the third and each subsequent measurement the procedure is: 1) compute the likelihood of the newly measured value based on the single-measurement function defined by the current ML reference; 2) if the likelihood is less than the defined threshold, issue a transient event warning. Do not update the reference; 3) if the likelihood is within the likelihood interval (Fig 1, panel d), recalculate the joint likelihood function including the new measurement and update the reference.

The joint likelihood is the likelihood function of the reference value *given the entire set of measurements*, and with each additional measurement, it gets narrower and more finely peaked on the ML reference value; the single-measurement likelihood is the function that shows how likely it is to measure any given value *each time a measurement is made*. The better the reference value, the more reliable the location of the single-measurement likelihood function. However, its shape depends only on the probability density of the random variable and on the reference value. The single or multiple thresholds used to identify the occurrence of a transient event must be defined and optimised according to the application.

X-ray Transient in Time Series

We are observing a hypothetical bursting X-ray source embedded in a region from which the average event rate is 1 s^{-1}. The observation lasts 1 hour, and in it occurs a weak burst lasting about 30 s, and from which 33 events are detected. Even though a hint of its presence is seen in the 30 s binned time series, (but not really with 60 s bins), this event could easily have gone unnoticed if the detection relied on the examination of similar time series. With the procedure described above, the burst is detected at a log-likelihood of -48.41 (likelihood of 10^{-21}).

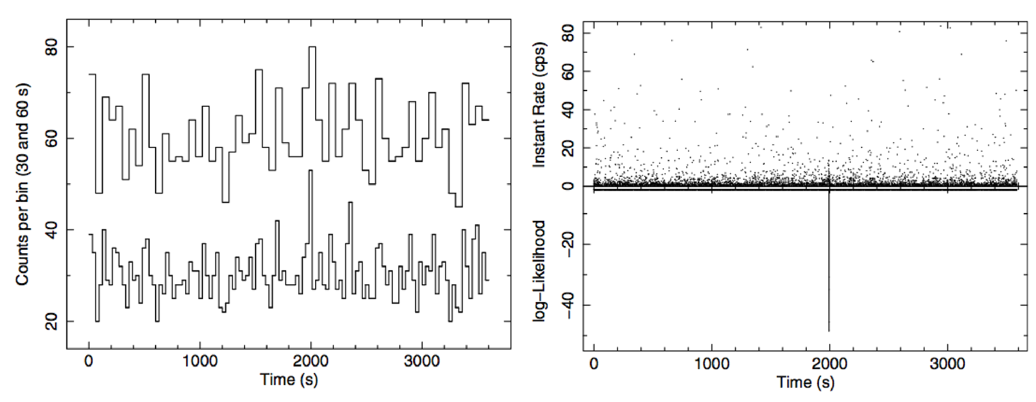

Figure 2. Time series of the observation shown in counts per bin for 30 and 60 s bins (left panel, bottom and top, respectively). Instantaneous count rate calculated as the inverse of the time between events shown as a function of each event's arrival time above the transient detection likelihood also evaluated in real time. The maximum value of the instantaneous rate is 2950 s^{-1}, but the scale was truncated to 86 s^{-1} to match the scale of the 60 s time series and better show the scatter. Values of the log-likelihood that do not meet the trigger criterion are shown at the warning threshold level of −2.1 (likelihood of 0.14). The sole detection is that of the transient event, and it dips down to −48.41 (likelihood of 10^{-21}).

This is the combined likelihood of detecting at least eight consecutive events, each of which met the warning threshold of log-likelihood -2.1, when expecting the average detection rate. In contrast, looking at the peak that stands out in the 30 s binned time series, we would compare a total intensity of 53 events against the expectation of 30, and find a likelihood of 5.8×10^{-4} , which might be enough to make us raise an eyebrow, but not much more.

The strategy was established using simulations: observations that did not include a burst gave rise to false transients less than 10% of the time, but that the when a burst was included it was detected 60% of the time. For the purpose of detecting strong but extremely short-lived transients, it would better to use a very low, single-point threshold. Each application has its own optimal settings that can be precisely tuned using simulations.

Transient QPO in X-ray Binary

We are now observing in X-rays of a bright (500 s^{-1}) accreting system whose emission comes from two components: the accretion disk and the hot inner flow. The emission processes give rise to red noise with different spectral indexes. The accretion disk is much larger in extent and has a sharp inner radius. It dominates at lower frequencies with a power-law index of -1, and has a high-frequency cutoff beyond which it does not contribute to the power spectrum. The turbulent inner flow is much smaller in extent because it is bounded by the inner edge of the disk. Its emission is more variable and dominates the high-frequency part of the spectrum with a power-law index of -3.

We are interested in monitoring the range of frequencies between 0.1 and 10 Hz for the appearance of a weak, short-lived, transient QPO that we expect to appear at or very near the break in the power spectrum at 1 Hz that marks the boundary between the disk and the turbulent inner flow. We make a periodogram every 10 seconds with the events accumulated during this time interval, and monitor the power at one or any number of frequencies. For a short-lived QPO, we cannot rely on the transient persisting in more than one "measurement", and therefore we must establish a single detection threshold using simulations. In this case, log-likelihood of -10.1 (likelihood of 4×10^{-5}), ensures a low level of false detections (5%).

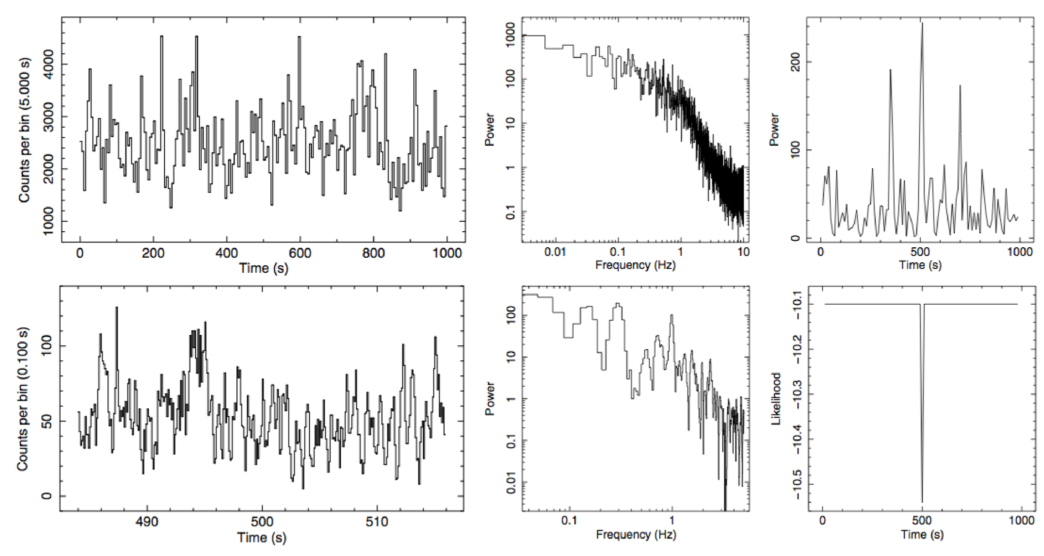

Figure 3. The top row shows the time series of the entire observation (binned to a resolution of 5 s for clarity of presentation); the periodogram made from the Kalman filtered, 0.05 s time series of the arrival times; and the power at 1 Hz estimated at 10 s intervals as a function of time. The bottom row shows a zoom on the time series during the transient QPO from its start at 485 s until its end at 515 s after the beginning of the observation; the periodogram of the Kalman-filtered 0.05 s resolution time series; and the log-likelihood as a function of time where only detections beyond the established threshold are shown. The QPO is characterised by 30 cycles of an almost periodic signal centred on 1 Hz with a standard deviation of 1/20 about that frequency and a pulsed fraction of 27%.

Concluding Remarks

Treating and interpreting data as statistical evidence in seeking to further our understanding of physical processes and of the behavior of complex systems such as those we observe in astrophysics, using all the information carried by these data, is most directly done through the use of the likelihood function, appropriately chosen according to the measured random variable.

The procedure presented is well suited to handle the first two classes of transients with constant background. It is obvious that identification efficiency depends intimately on the strength of the transient signal. The approach is perfectly well suited for analysing archival data, but especially powerful for real-time applications.

If the process is variable but predictable, then this is an extension of the procedure using a model, but in which it itself evolves as a function of time. The formalism is otherwise identical. If the process is variable and unpredictable, it implies that the measurements in each pixel or channel are not distributed according to a fixed probability distribution. Therefore, each pixel or channel is treated individually, but because we have no a priori expression for the likelihood function, the intensity and how it is distributed is characterised by the running mean and variance.

For highly variable processes, where deviations in shape from known probability distributions are large, looking at the distribution of measured values is not really useful, because the changing intensity in each cannot be described by random variables with stationary probability distributions. However, a variable process can be highly non-stationary in the time domain but stationary in frequency space. This is analogous to a variable point source whose intensity varies markedly in successive images but whose location in the sky remains the same, and whose shape as it appears in each image is as always given by the PSF.

Combining the information carried by the data in the time and frequency domains, and treating it simultaneously in the fashion described in this paper is a most powerful means for detecting transient features in highly variable processes. A detailed investigation of how to best characterise such processes, be it for identifying transients or for some other purpose, will be addressed elsewhere.

References

Cowan, G. 1997, Statistical data analysis (Oxford)

Fisher, R. A. 1912, Messenger Math., 41, 155

Fisher, R. A. 1922, RSPTA, 222, 309

Kalman, R. E. 1960, J. Basic Eng., 82, 35

Konig, M., & Timmer, J. 1997, A&AS, 124, 589

Royall, R. M. 1997, Statistical Evidence, A Likelihood Paradigm (New York: Chapman & Hall/CRC)