Detecting stellar flares using conditional volatility

Giovanni Motta
Data Science Institute, Columbia University

joint with **Johnny Wang** (*Texas A&M*) and **Vinay Kashyap** (*cfa*)

Center for Astrophysics, Harvard & Smithsonian

5-8-2024

The 2-minute cadence flux light curves of TIC13955147 & TIC269797536

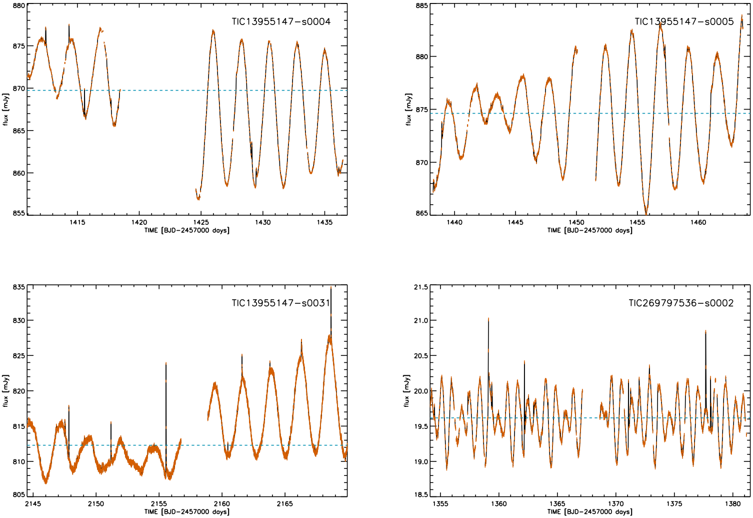

Flares

- Flares are impulsive events that release large bursts of energy over short timescales.
- The aftereffects of solar flares include spectacular aurorae, interference in satellite communications, expansion of the thermosphere, and in particularly strong cases could cause widespread electrical grid disruptions and blackouts.
- Flares occur when magnetic energy stored in twisted flux tubes is released into the hot plasma of the stellar corona. They occur predominantly where strong magnetic fields are present, around sunspots, in areas called "Active Regions".
- However, flare onset is poorly understood. One of the major results of our work is that the residual light curve can be fit as GARCH(p, q).

Our contribution

- Flares are believed to occur due to a process called self-organized criticality, which also governs the onset of avalanches and earthquakes, and the process is known to be similar on stars as on the Sun (see e.g., Aschwanden et al. 2014, 2018).
- Studying stars is an excellent way to improve our understanding of the Sun. Effectively, they allows us to look at how the Sun would behave if it had a different mass, or age, or composition, or rotational rate, or magnetic dynamo strength, or even if it had a close-in giant hot Jupiter.
- Studying ensembles of stars will let us pin down parameters and processes that are important and control the nature of the corona.
- Here we develop a method to detect flare occurrences in TESS light curves in order to gain insight into flare distributions, cadence, and onset.

Our contribution

- We have studied several white light flares from the Transiting Exoplanet Survey Satellite (TESS), and modeled them via a GARCH process.
- The GARCH process has been applied extensively in econometrics to understand stock price fluctuations, and our preliminary findings show that stellar flares can also be described by the same mathematics.
- We study the so-called volatility index, and find that it is a good guide to flare occurrence. This is a remarkable new result, and promises to open new avenues of exploration, to understand the precise nature of flare onset.
- Our goal: to define the modeling process, and to analyze the large dataset of light curves obtained from TESS to characterize the different stars. This will allow us to explore differences with spectral type, with stellar age and activity level, and even whether tidal and magnetic interactions with nearby exoplanets changes flaring behavior
- We develop a new and powerful method to analyze time series to detect flares in TESS light curves.

Our novel semi-parametric approach

Our approach combines a non-parametric method that removes the harmonic trend with a parametric method that detects the flares

1. **Non-parametric**
   First, we remove the trend using a time-varying harmonic fit combined with a NP-test for the local variance, so to capture changes in the deterministic amplitude of the light curve.
2. **Parametric**
   Then we enlighten the analogy between the stochastic part of the light curves and GARCH processes, and detect the flares by exploiting the parametric structure.

- We demonstrate that flares can be detected as significantly large deviations from the baseline.
- We apply the method on exemplar light curves from flaring stars **TIC13955147** (s4, s5 and s31) and **TIC269797536** (s2).

A snapshot of our novel method

TIC13955147 s-31

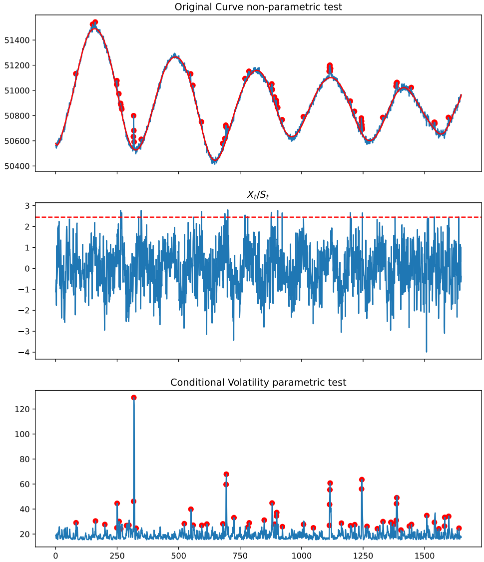

Outline

Flares

Our Statistical Model

- Deterministic Trend: Time-varying Harmonic Oscillations
  - *Modulated* luminosity of variable stars
  - Amplitude modulation
  - Frequency modulation
- Stochastic Errors: Conditional Heteroskedastic Volatility

Our novel semi-parametric approach

Data: TESS Light Curves

Our Statistical Model

Our novel model is the sum of a non-random trend and a stochastic error:

$$Y_t = \mu(t) + X_t, \qquad t = 1, \ldots, n \tag{1}$$

- The function $\mu(t)$ is a deterministic, periodic time-trend that allows for time-varation in its coefficients, whereas
- the errors X_t follow a GARCH model

$$X_t = \sigma_t \varepsilon_t, \qquad \varepsilon_t \overset{iid}{\sim} \mathcal{N}(0,1), \tag{2}$$

where the stochastic conditional variance $\sigma_t^2 = \mathrm{Var}(X_t | \mathcal{I}_t)$ depends linearly on p past squared shocks:

$$\sigma_t^2 = \alpha_0 + \sum_{i=1}^{p} \alpha_i X_{t-i}^2 + \sum_{j=1}^{q} \beta_j \sigma_{t-j}^2 \tag{3}$$

Deterministic Trend: Time-varying Harmonics

Data:

- random observed vector $\boldsymbol{Y} = (Y_1, Y_2, \ldots, Y_N)'$
- time vector $\boldsymbol{t} = (t_1, t_2, \ldots, t_N)'$

Modulation model

$$Y_i = \sum_{k=1}^{K} \{g_{1k}(t_i)\cos(w_k t_i) + g_{2k}(t_i)\sin(w_k t_i)\} + X_i, \qquad i = 1, \ldots, N,$$

- smooth time-varying amplitudes
- K number of harmonics
- $\boldsymbol{X} = (X_1, \ldots, X_N)'$ follow a GARCH(p, q) process

We estimate the parameters using penalized B-splines (Motta et al., 2002)

Outline

Modulation

Continuous wave modulation can be divided into two sets:

- amplitude modulation (AM)
- angle modulation
  - frequency modulation (PM)
  - phase modulation (FM)

Amplitude modulation

- Carrier wave $c(t) = U_c \sin(2\pi f_c t + \phi_c)$
where U_c, f_c, and ϕ_c are the amplitude, frequency, and phase of the carrier wave, respectively.

- Modulation signal $U_m(t)$: waveform (the message) to be transmitted

- Modulated signal

$$\begin{aligned} U_{AM}(t) &= [U_c + U_m(t)] \sin(2\pi f_c t + \phi_c) \\ &= \left[1 + \frac{U_m(t)}{U_c}\right] c(t) \qquad (4) \end{aligned}$$

Amplitude modulation: Example 1

- Carrier wave $c(t) = U_c \sin(2\pi f_c t + \phi_c)$

- Modulation signal $U_m(t) = U_m^A \sin(2\pi f_m t + \phi_m^A)$

- Modulated signal

$$U_{AM}(t) = [U_c + U_m^A \sin(2\pi f_m t + \phi_m^A)] \sin(2\pi f_c t + \phi_c),$$

or equivalently, using

$$\sin(a)\sin(b) = \tfrac{1}{2}[\cos(a-b) - \cos(a+b)] \quad \text{and} \quad \sin(a) = \cos(a - \tfrac{\pi}{2})$$

$$\begin{aligned} U_{AM}(t) = {} & U_c \sin(2\pi f_c t + \phi_c) \\ & + \frac{U_m^A}{2}\left[\sin(2\pi(f_m - f_c) + (\phi_m^A - \phi_c + \frac{\pi}{2}))\right] \\ & - \frac{U_m^A}{2}\left[\sin(2\pi(f_m + f_c) + (\phi_m^A + \phi_c + \frac{\pi}{2}))\right] \end{aligned}$$

Amplitude modulation: Example 2

Suppose $U_{c1} = U_{c2} = U_c$, then the carrier wave

$$c(t) = U_{c1}\sin(2\pi f_{c1}t + \phi_{c1}) + U_{c2}\sin(2\pi f_{c2}t + \phi_{c2}). \tag{5}$$

is given by

$$c(t) = U_c\sin(2\pi f_{c1}t + \phi_{c1}) + U_c\sin(2\pi f_{c2}t + \phi_{c2}). \tag{6}$$

Using $\sin(a) + \sin(b) = 2\cos(\frac{a-b}{2})\sin(\frac{a+b}{2})$,

$$c(t) = \underbrace{\left[2U_c\cos\left(2\pi\frac{(f_{c1}-f_{c2})}{2}t + \frac{\phi_{c1}-\phi_{c2}}{2}\right)\right]}_{[U_c(t)=time-varying\ \ amplitude]}\underbrace{\sin\left(2\pi\frac{(f_{c1}+f_{c2})}{2}t + \frac{\phi_{c1}+\phi_{c2}}{2}\right)}_{average\ \ wave}$$

If $U_m(t) = U_m^A\sin(2\pi f_m t + \phi_m^A)$, the modulated signal $U_{AM}(t)$ is as follows

$$U_{AM}(t) = [U_c(t) + U_m^A\sin(2\pi f_m t + \phi_m^A)]\sin\left(2\pi\frac{(f_{c1}+f_{c2})}{2}t + \frac{\phi_{c1}+\phi_{c2}}{2}\right),$$

or in form of equation (4) we have

$$U_{AM}(t) = \left[1 + \frac{U_m^A\sin(2\pi f_m t+\phi_m^A)}{U_c(t)}\right]\left[U_c\sum_{i=1}^{2}\sin(2\pi f_{ci}t + \phi_{ci})\right],$$

where $U_c(t)$ is the amplitude of the non-modulated curve $c(t)$.

Frequency modulation: Example 3

A simple case: $\int_0^t U_m(\tau)d\tau$ takes the form

$$\int_0^t U_m(\tau)d\tau = \sin(2\pi f_m t + \phi_c^F). \tag{7}$$

Substituting (7) in

$$U_{FM}(t) = U_c \sin(2\pi f_c t + \int_0^t U_m(\tau)d\tau + \phi_c), \tag{8}$$

the modulated signal is

$$U_{FM}(t) = U_c \sin(2\pi f_c t + \sin(2\pi f_m t + \phi_c^F) + \phi_c).$$

Examples 1 & 2 (AM) and 3 (FM)

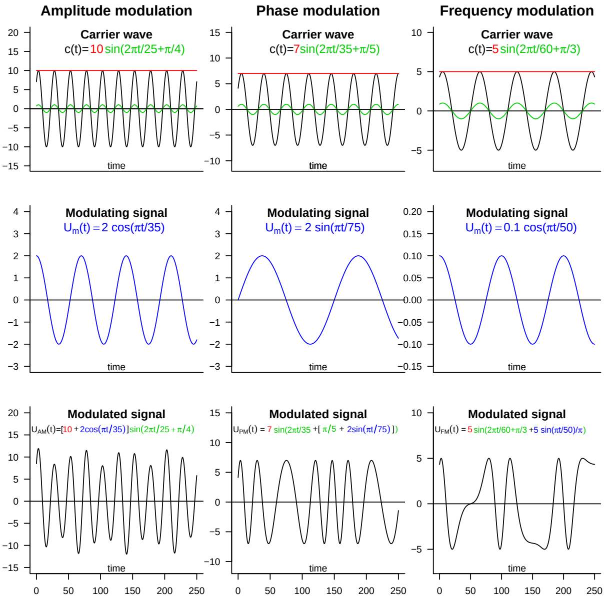

Outline

Stochastic Errors: Heteroskedastic Volatility

ARCH(p) model has a multiplicative form

$$X_t = \sigma_t \varepsilon_t, \qquad \varepsilon_t \sim N(0,1)\,, \tag{9}$$

where

$$\sigma_t^2 = \alpha_0 + \sum_{i=1}^{p} \alpha_i X_{t-i}^2, \quad t \geq p\,, \tag{10}$$

with $\alpha_0 > 0, \quad \alpha_1, \ldots, \alpha_p \geq 0.$

Properties

- $\mathbb{E}(X_t | X_s, s < t) = \mathbb{E}(X_t) = 0$
- $\mathbb{E}(X_t X_s) = 0$, $s \neq t$
- $\mathbb{E}(X_t^2 | X_s, s < t) = \sigma_t^2 = \alpha_0 + \sum_{i=1}^{p} \alpha_i X_{t-i}^2$
- $\mathbb{E}(X_t^2) = \alpha_0 / (1 - \sum_{i=1}^{p} \alpha_i)$
- $X_t^2 \sim \mathrm{AR}(p)$ and $\rho_{X^2}(h) \geq 0$ for all h (generating persistence of volatility)
- X_t has heavier tails than ε_t, in the sense that its kurtosis

$$\frac{\mathbb{E}X_t^4}{(\mathbb{E}X_t^2)^2}$$

  is $\geq$ than that of ε_t

Stochastic Errors: Heteroskedastic Volatility

A GARCH(p, q) model has a multiplicative form

$$X_t = \sigma_t \varepsilon_t, \qquad \varepsilon_t \sim N(0,1)\,, \tag{11}$$

where

$$\sigma_t^2 = \alpha_0 + \sum_{i=1}^{p} \alpha_i X_{t-i}^2 + \sum_{j=1}^{q} \beta_j \sigma_{t-j}^2, \quad t \geq \max(p,q)\,, \tag{12}$$

with $\alpha_0 > 0, \quad \alpha_1, \ldots, \alpha_p \geq 0, \quad \beta_1, \ldots, \beta_q \geq 0.$

Properties

- $\mathbb{E}(X_t | X_s, s < t) = \mathbb{E}(X_t) = 0$
- $\mathbb{E}(X_t X_s) = 0$, $s \neq t$
- $\mathbb{E}(X_t^2 | X_s, s < t) = \sigma_t^2 = \alpha_0 + \sum_{i=1}^{p} \alpha_i X_{t-i}^2 + \sum_{j=1}^{q} \beta_j \sigma_{t-j}^2$
- $\mathbb{E}(X_t^2) = \alpha_0 / (1 - \sum_{i=1}^{p} \alpha_i - \sum_{j=1}^{q} \beta_j)$
- $X_t^2 \sim \text{ARMA}(p, q)$ and $\rho_{X^2}(h) \geq 0$ for all h

Stochastic Errors: Heteroskedastic Volatility

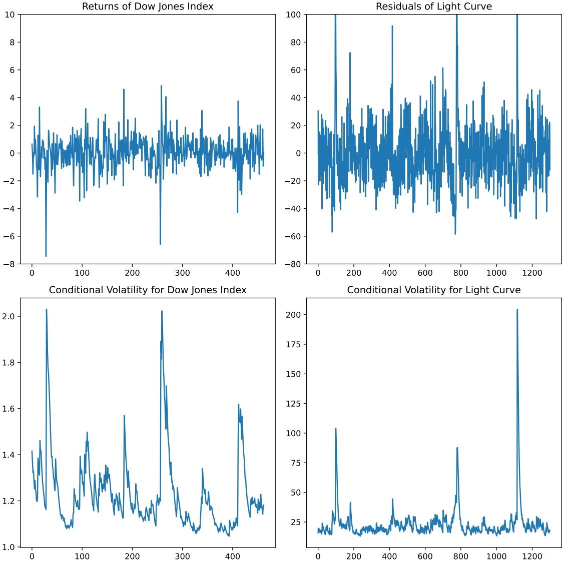

Heteroskedastic Volatility

Stylized facts of financial resturs

- **heavy tailed**

$$\Pr(|X| > x) \sim C x^{-\alpha}, \quad 0 < \alpha < 4$$

- **uncorrelated**
  $\rho_X(h)$ near 0 for all lags $h > 0$

- $|X_t|$ and X_t^2 have **slowly decaying autocorrelations**

$$\rho_{|X|}(h) \text{ and } \rho_{X^2}(h) \text{ converge to 0 slowly as } h \text{ increases}$$

- **stochastic** and **persistent** volatility
  - large (small) fluctuations in the data tend to be followed by fluctuations of comparable magnitude
  - this is reflected by GARCH models through the correlation in the sequence $\{\sigma_t^2\}$ of conditional variances

Outline

A semi-parametric approach

NON-PARAMETRIC de-trend

- We detect flares in TESS light curves via hypothesis testing. We propose a uniform threshold, computed using the residuals of light curves $X_t = Y_t - \mu(t)$ by examining each time-varying volatility within its local neighborhood.

- Since the conditional variance of a GARCH is time-varying, given the residual value $X_t = Y_t - \mu(t)$, we can test

$$H_0: \ X_t|\mathcal{I}_t \sim WN(0, \sigma_0^2) \qquad \text{vs} \qquad X_t|\mathcal{I}_t \sim WN(0, \sigma_t^2).$$

  Under H_0, $\text{Var}(X_t|\mathcal{I}_t) = \sigma_0^2 \ \forall t$. Under H_1, $\text{Var}(X_t|\mathcal{I}_t) = \sigma_t^2 > \sigma_0^2$ for some t.

- To this end, we define our time-varying test statistic as X_t^2/S_t^2, where S_t is a **robust**, **time-varying** estimator of the conditional volatility.

- H_0: $X_t = \sigma_0 \varepsilon_t \sim \mathcal{N}(0, \sigma_0^2)$ for all t, whereas $H_1 : X_t|\mathcal{I}_t \sim \mathcal{N}(0, \sigma_t^2 > \sigma_0^2)$. Under the null, X_t^2/S_t^2 is distributed according to a $F_{1,n-1}$ distribution.

- We reject the null if $X_t^2/S_t^2 > F_{1,n-1}(1-\alpha)$.

A semi-parametric approach

NON-PARAMETRIC de-trend

- As $n \to \infty$, $F_{1,n-1} \xrightarrow{d} \chi^2_1$.

- $\big(X_t/S_t\big)^2 > \chi^2_{1(\alpha)}$ is equivalent to $|X_t|/S_t > Q^{\text{HN}}_{(\alpha)}$, where $\alpha = \Pr\big(\mathcal{H} > Q^{\text{HN}}_{(\alpha)}\big)$, $\mathcal{H}$ being a Standard Half-Normal random variable.

- Our ejection region: $\{\,|Y_t - \mu(t)| > S_t\, Q^{\text{HN}}_{(\alpha)}\}$.

- We reject the H_0 if $\Pr\big(\mathcal{H} > Q^{\text{HN}}_{(\alpha)}\big) < \alpha$, where the tuning parameter α (the significance level of the test) is often set by the astronomy *ex-ante*.

- We estimate α rather than assigning an arbitrary valued to this parameter: we define an upper bound $\alpha_{\max}$ and then select, among those values of the above probability that are smaller than $\alpha_{\max}$, the largest one.

- The presence of flares makes $\hat{\mu}$ biased. For this reason, we (detect and) delete the flares from Y and then re-fit μ to the 'flares-filtered' data.

A semi-parametric approach: robust estimation of σ_t

Hampel (1974)

$$MAD_n = 1.4826 \times \underset{1 \leq i \leq n}{\text{med}} |X_i - \underset{1 \leq j \leq n}{\text{med}} X_j|$$

- The MAD has the best possible breakdown point (50%, twice as much as the interquartile range)
- A breakdown point of 50% is the highest possible, because the estimate remains bounded when fewer than 50% of the data points are replaced by arbitrary numbers.
- The MAD also has some drawbacks.

1) Its efficiency at Gaussian distributions is very low; whereas the location median's asymptotic efficiency is still 64%, the MAD is only 37% efficient.
2) MAD takes a symmetric view on dispersion, because one first estimates a central value (the median) and then attaches equal importance to positive and negative deviations from it.

A semi-parametric approach: robust estimation of σ_t

Rousseeuw & Croux (1993)

$$S_n = 1.1926 \times \underset{1\leq i\leq n}{\text{med}} \{ \underset{1\leq j\leq n}{\text{med}} |X_i - X_j| \}$$

- More efficient and not slanted towards symmetric distributions.
- Unlike the MAD, it does not need any location estimate of the data. Instead of measuring how far away the observations are from a central value, S_n looks at a typical distance between observations, which is still valid at asymmetric distributions.

A semi-parametric approach: robust estimation of σ_t

Rousseeuw & Croux (1993)

Simulate

$$\sigma_t^2 = \alpha_0 + \alpha_1 X_{t-1}^2 + \beta_1 \sigma_{t-1}^2 \, ,$$

with $\alpha_0 = 0.1$, $\alpha_1 = 0.2$, and $\beta_1 = 0.5$.

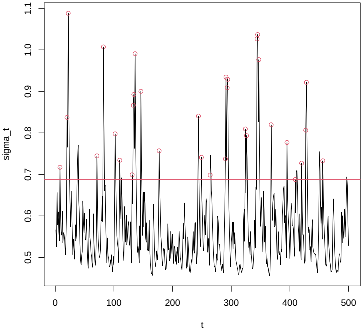

A semi-parametric approach

PARAMETRIC flares-detection

$$\sigma_t^2 = \alpha_0 + \sum_{i=1}^{p} \alpha_i \sigma_{t-i}^2 \varepsilon_{t-i}^2 + \sum_{j=1}^{q} \beta_j \sigma_{t-j}^2, \quad t \geq \max\{p, q\}, \tag{13}$$

- Parametric test statistics based on the past conditional volatility.
- Generalized χ^2 distribution to model the sequence of conditional variances:

$$\tilde{\chi}(\boldsymbol{w}, \boldsymbol{k}, \boldsymbol{\lambda}, C_t) = \sum_i w_i \varepsilon_{t-i,k_i,\lambda_i}^2 + C_t,$$

  a weighted sum of non-central Chi-square random variables $\varepsilon_{t-i,k_i,\lambda_i}^2$, with degrees of freedom k_i and non-centralities λ_i.
- Test statistics:

$$\sigma_t^2 | \sigma_{\mathcal{I}_t}^2 = \tilde{\chi}_t(\boldsymbol{w}_t, \boldsymbol{k}, \boldsymbol{\lambda}, C_t),$$

  with $w_{t,i} = \alpha_i \sigma_{t-i}^2$, degrees of freedom $k_i = 1$, and non-centralized parameters $\lambda_i = 0 \ \forall i$, and $C_t = \alpha_0 + \sum_{j=1}^{q} \beta_j \sigma_{t-j}^2$. We decide that t is an extreme point if

$$P_t = \Pr(\tilde{\chi}_t(\boldsymbol{w}_t, \boldsymbol{k}, \boldsymbol{\lambda}, C_t) > \sigma_t^2) \leq \alpha \,.$$

Our algorithm

Non-parametric detrending

Define $X^{(0)} = Y^{(0)} - \mu^{(0)}$ and $\alpha^{(0)} = \alpha_{\max}$

for i in $0 : (I-1)$ **do**

for all t in 1:n **do** $\pi_t^{(i)} = \Pr\big(\mathcal{H} > \frac{|X_t^{(i)}|}{S_t^{(i)}}\big)$

if $\pi_t^{(i)} < \alpha^{(i)}$ **then**

delete $\{t, Y_t\}$

else

$Y_t^{(i+1)} = Y_t^{(i)}$

$X^{(i+1)} = Y^{(i+1)} - \mu^{(i+1)}$, and $\alpha^{(i+1)} = \max_{t:\, \pi_t^{(i)} < \alpha^{(i)}} \pi_t^{(i)}$

Parametric Flares-detection

Fit a GARCH(p, q) to $X^{(I)} = Y^{(0)} - \mu^{(I)}$

for all t in 1:n **do** $P_t = \Pr[\tilde{\chi}_t(\boldsymbol{w}_t, \boldsymbol{k}, \boldsymbol{\lambda}, C_t) > \sigma_t^2]$

if $\mathbb{1}_{\{P_t < \alpha\}} \times \mathbb{1}_{\{X_t^{(I)} > 0\}} = 1$ **then**

Define $F_t = Y_t$

Our algorithm

Non-parametric detrending

Define $X^{(0)} = Y^{(0)} - \mu^{(0)}$ and $\alpha^{(0)} = \alpha_{\max}$

for i in $0 : (I-1)$ **do**

for all t in 1:n **do** $\pi_t^{(i)} = \Pr\big(\mathcal{H} > \frac{|X_t^{(i)}|}{S_t^{(i)}}\big)$

if $\pi_t^{(i)} < \alpha^{(i)}$ **then**

delete $\{t, Y_t\}$

else

$Y_t^{(i+1)} = Y_t^{(i)}$

$X^{(i+1)} = Y^{(i+1)} - \mu^{(i+1)}$, and $\alpha^{(i+1)} = \max_{t:\, \pi_t^{(i)} < \alpha^{(i)}} \pi_t^{(i)}$

Parametric Flares-detection

Fit a GARCH(p, q) to $X^{(I)} = Y^{(0)} - \mu^{(I)}$

for all t in 1:n **do** $P_t = \Pr[\tilde{\chi}_t(\boldsymbol{w}_t, \boldsymbol{k}, \boldsymbol{\lambda}, C_t) > \sigma_t^2]$

if $\mathbb{1}_{\{P_t < \alpha\}} \times \mathbb{1}_{\{X_t^{(I)} > 0\}} = 1$ **then**

Define $F_t = Y_t$

The value of α is guaranteed to decrease over i because the inequality appearing in both the condition $\pi_t^{(i)} < \alpha^{(i)}$ and the update $\alpha^{(i+1)}$ is strict.

Outline

Flares

Our Statistical Model

- Deterministic Trend: Time-varying Harmonic Oscillations
  - *Modulated* luminosity of variable stars
  - Amplitude modulation
  - Frequency modulation
- Stochastic Errors: Conditional Heteroskedastic Volatility

Our novel semi-parametric approach

Data: TESS Light Curves

Data: TESS Light Curves

- We demonstrate the application of our method to TESS (Transiting Exoplanet Survey Satellite)[1] datasets which have been observed multiple times.
- We choose two stars on which flares have been unambiguously observed (see Feinstein et al. 2022).
- Both are eruptive variables, at similar distances, and one is solar-like, and the other is a low-mass star.
- We apply our method to detect flares in the `PDCSAP` (Pre-Search Data Conditioning Simple Aperture Photometry) electron count rate light curves, which have backgrounds subtracted and systematics corrected for cotrending basis vectors[2].
- We show the light curves, converted from electron count rate to mJy[3] for TIC 13955147 and for TIC 269797536.

---

[1] https://tess.mit.edu

[2] https://spacetelescope.github.io/mast_notebooks/notebooks/TESS/beginner_tour_lc_tp/beginner_tour_lc_tp.html

[3] https://tess.mit.edu/public/tesstransients/pages/readme.html

List of stars and datasets selected

We downloaded the processed light curve data from the TESS Science Processing Pipeline (TESS-SPOC)[4] from the MAST Portal (Mikulski Archive for Space Telescopes)[5] for the sectors listed in the Table below.

TIC	Other Names†	Spectral Type†	Distance† [pc]	Sectors
13955147	HD 32372 2MASS J05005186-4101065 Gaia DR3 4813691219557127808 1RXS J050051.7-410100	G5 V	78	4, 5, 31, 32
269797536	2MASS J04363294-7851021 Gaia DR3 4622912654918835200	M4 V	70	2, 5, 8, 11, 12, 13, 27, 28, 29, 32, 35, 38, 39

†: From the SIMBAD database (Wenger et al. 2000)

[4]https://archive.stsci.edu/hlsp/tess-spoc

[5]https://archive.stsci.edu/missions-and-data/tess

TIC13955147-s31

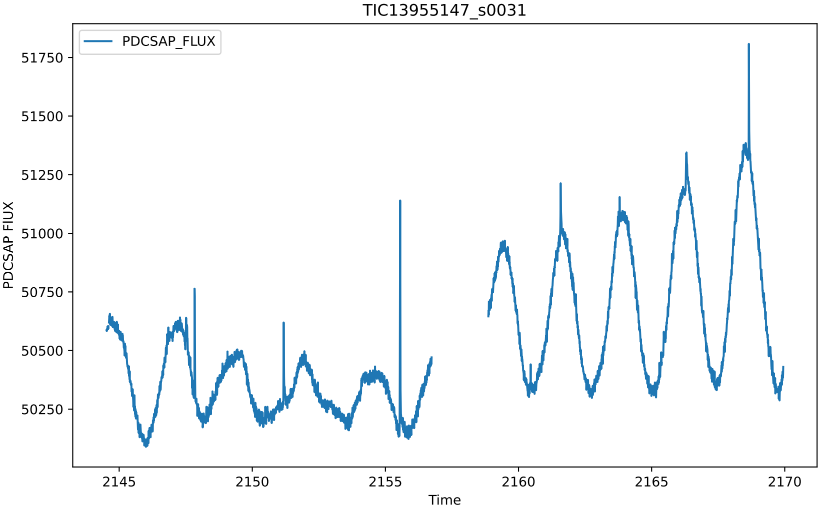

TIC13955147-s31-I

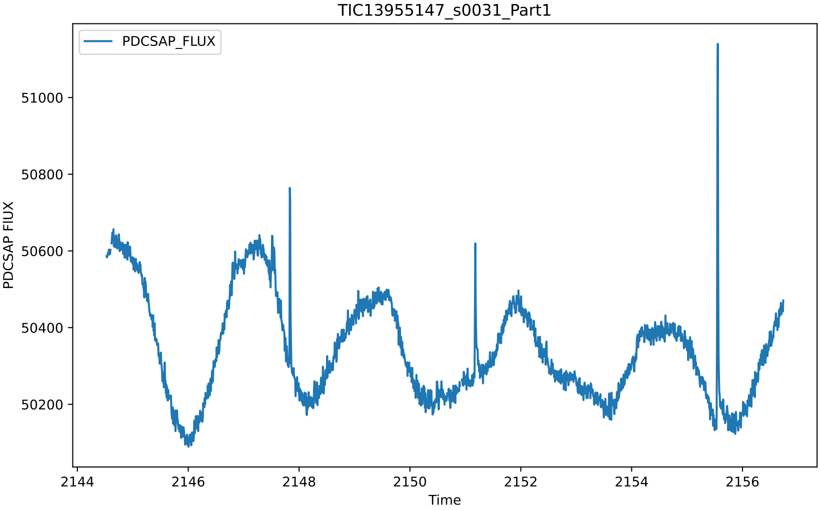

TIC13955147-s31-I – Harmonic Fit

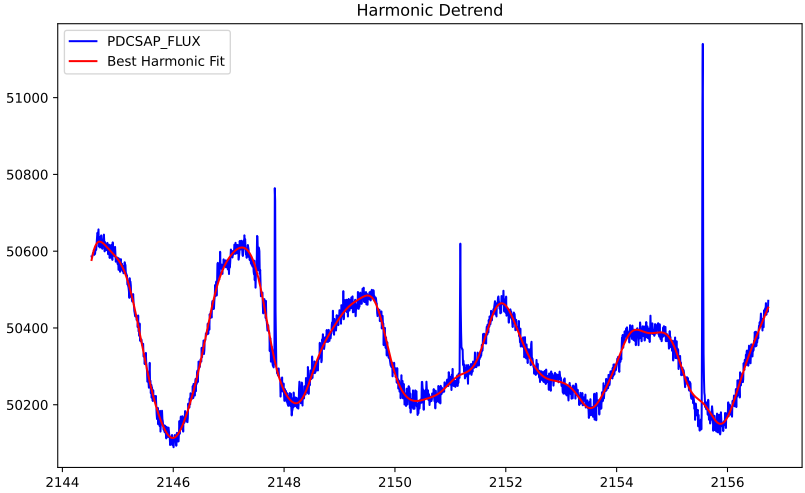

TIC13955147-s31-I – Residuals

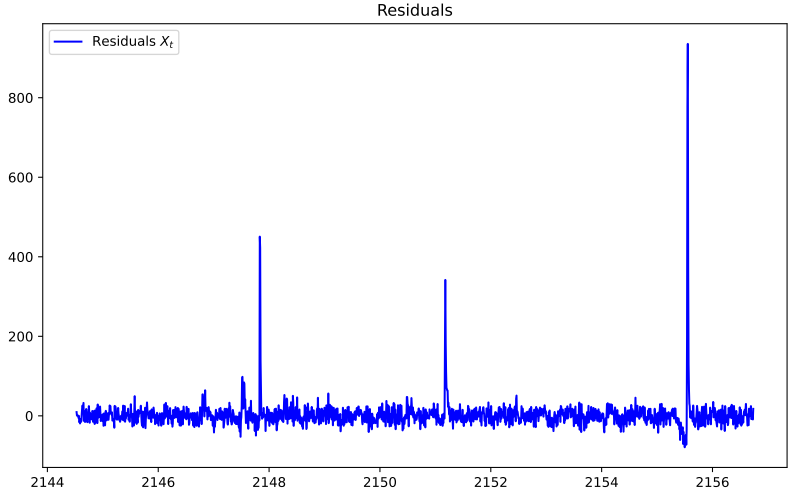

TIC13955147-s31-I – Non-parametric flares detection

$\pi_t = \Pr(\mathcal{H} > \frac{|X_t|}{S_t}) < \alpha$

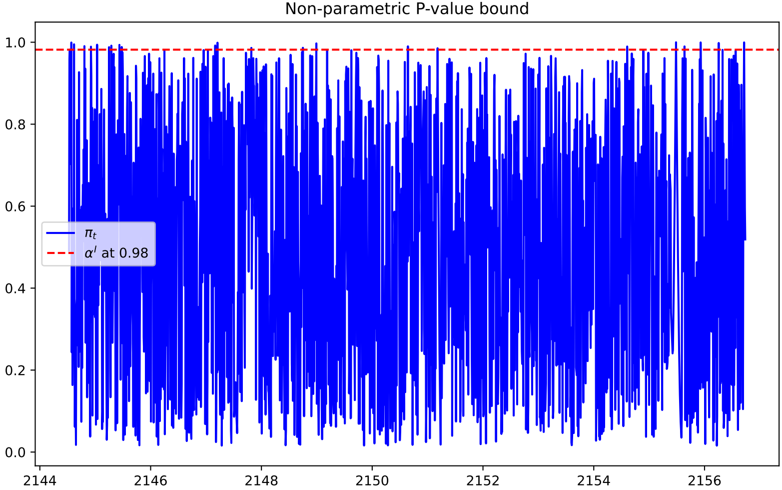

TIC13955147-s31-I – Non-parametric flares detection

$\pi_t = \Pr(\mathcal{H} > \frac{|X_t|}{S_t}) < \alpha$

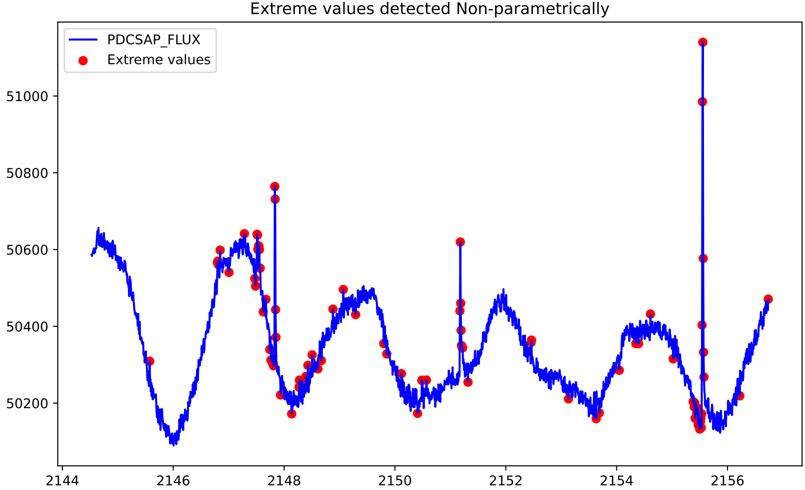

TIC13955147-s31-I – GARCH fitted to the residuals

$X = Y^{(0)} - \mu^{(I)}$

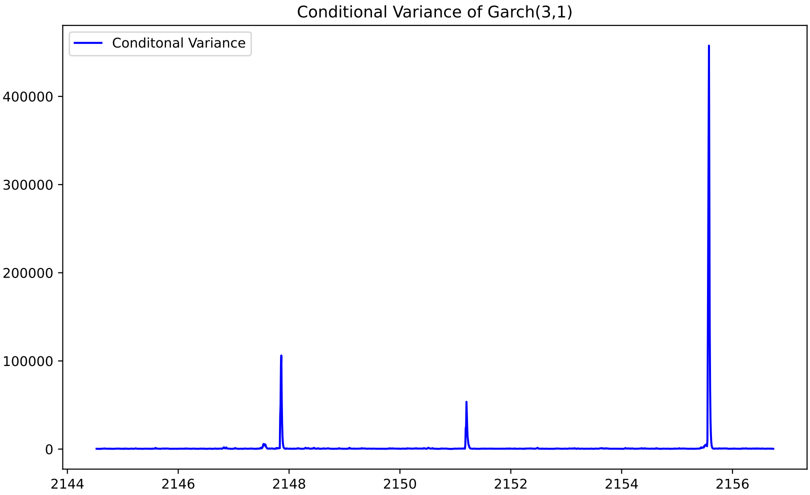

TIC13955147-s31-I – Non-parametric flares detection

$X = Y^{(0)} - \mu^{(I)}$

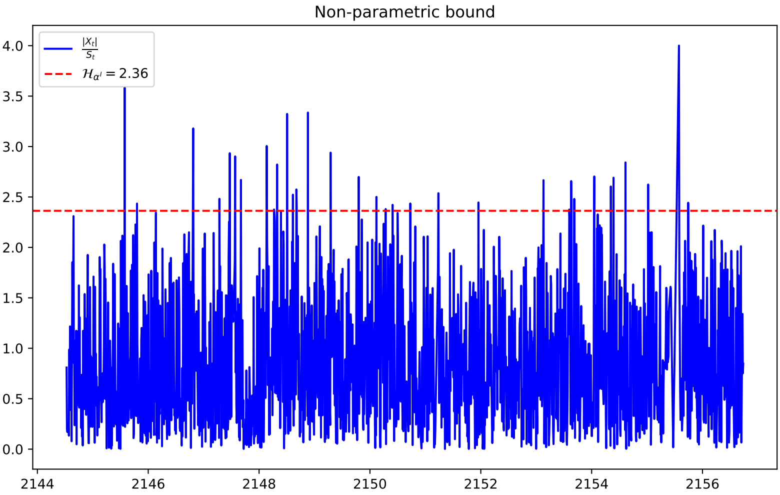

TIC13955147-s31-I – Parametric flares detection

$X = Y^{(0)} - \mu^{(I)}$

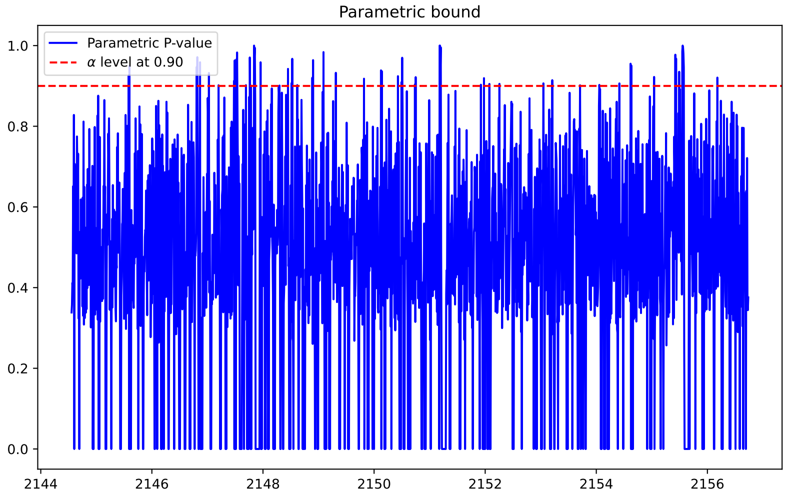

TIC13955147-s31-I – parametric flares detection

$P_t = \Pr(\tilde{\chi}_t(\boldsymbol{w}_t, \boldsymbol{k}, \boldsymbol{\lambda}, C_t) > \sigma_t^2) \leq \alpha$

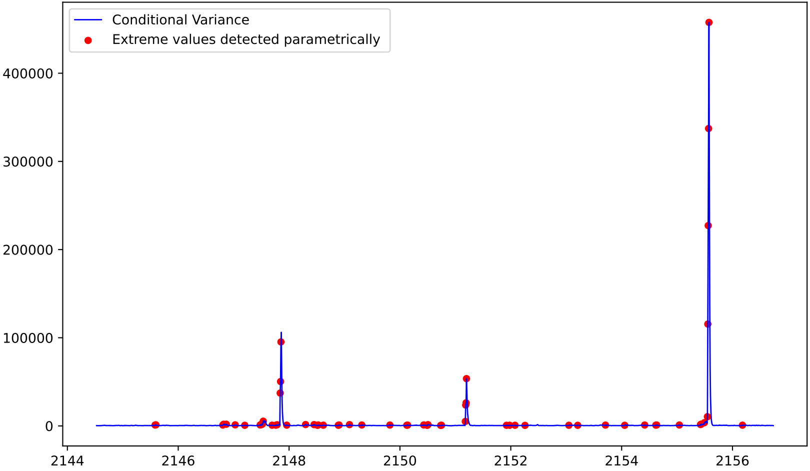

TIC13955147-s31-I – parametric flares detection

Flares are positive extreme values: $\mathbb{1}_{\{P_t \leq \alpha\}} \times \mathbb{1}_{\{X_t^{(I)} > 0\}} = 1$

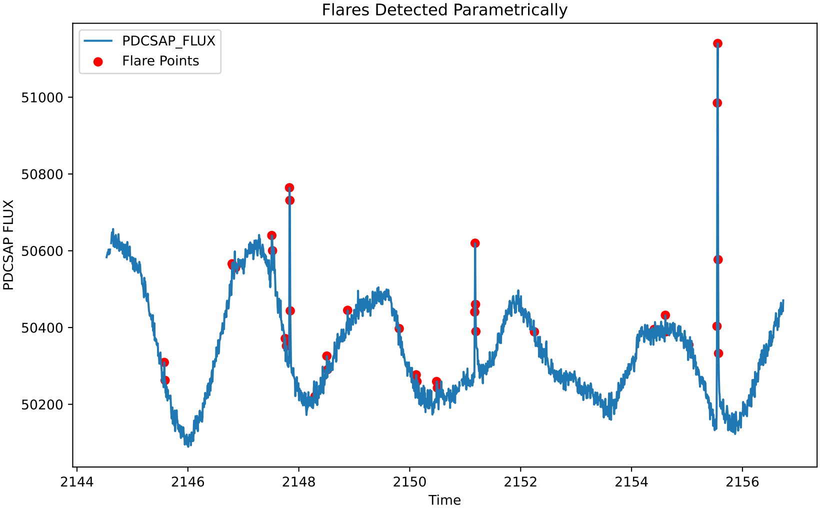

TIC13955147-s4

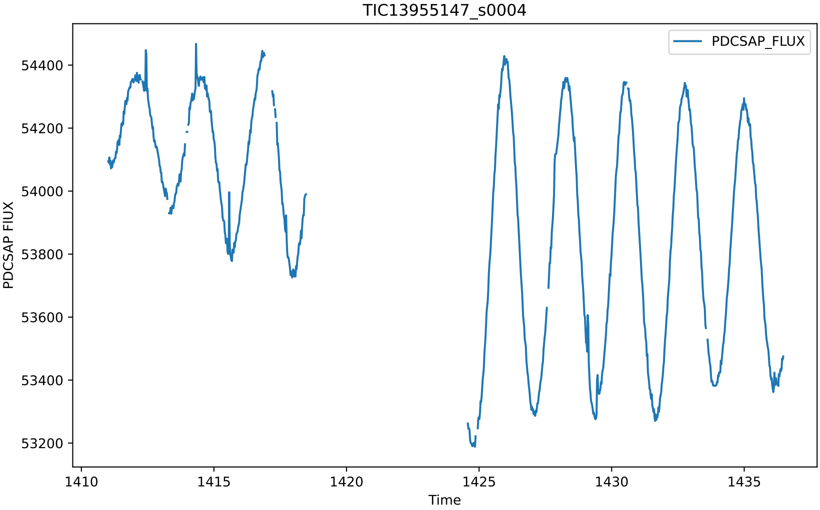

TIC13955147-s4-I

non-parametric detrend

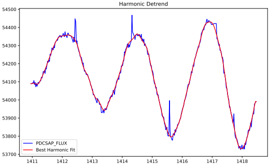

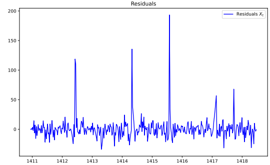

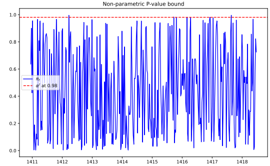

parametric flares detection

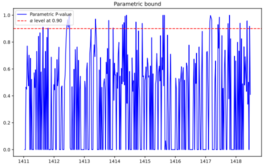

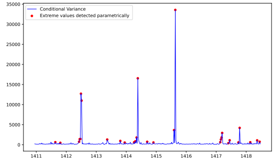

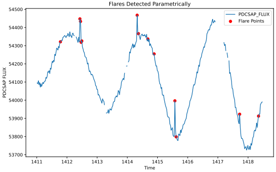

TIC13955147-s4-II

non-parametric detrend

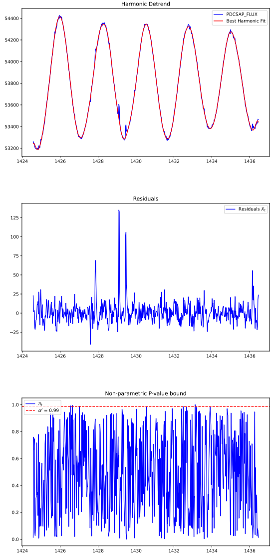

parametric flares detection

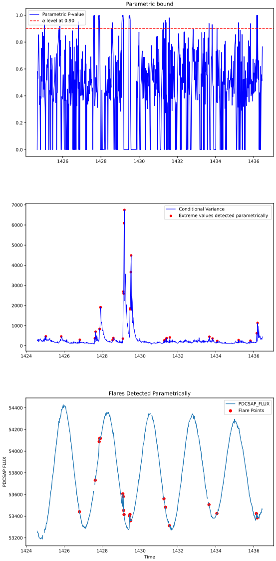

TIC13955147-s5

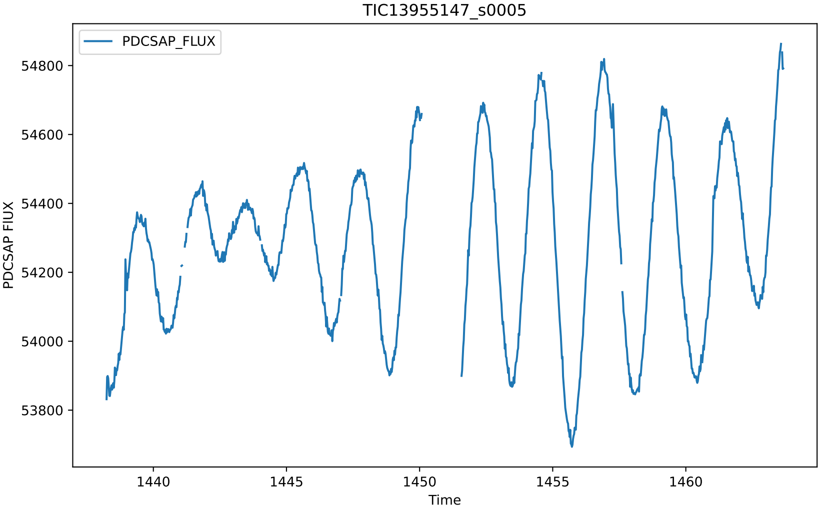

TIC13955147-s5-I

non-parametric detrend

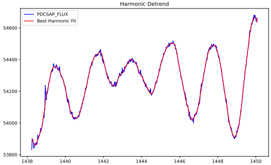

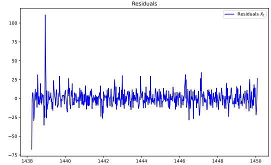

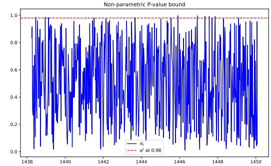

parametric flares detection

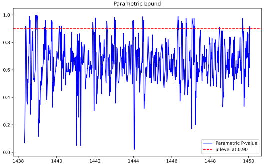

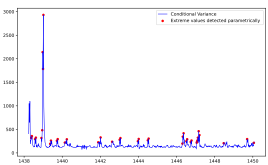

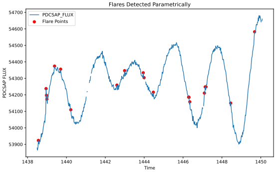

TIC13955147-s5-II

non-parametric detrend

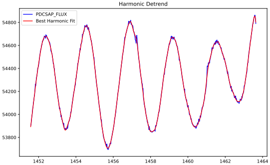

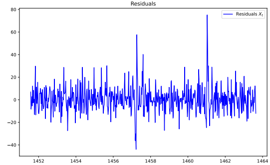

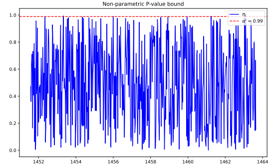

parametric flares detection

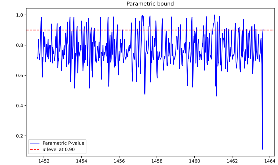

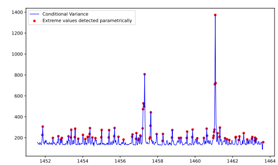

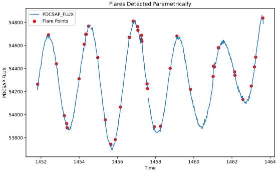

TIC269797536-s2

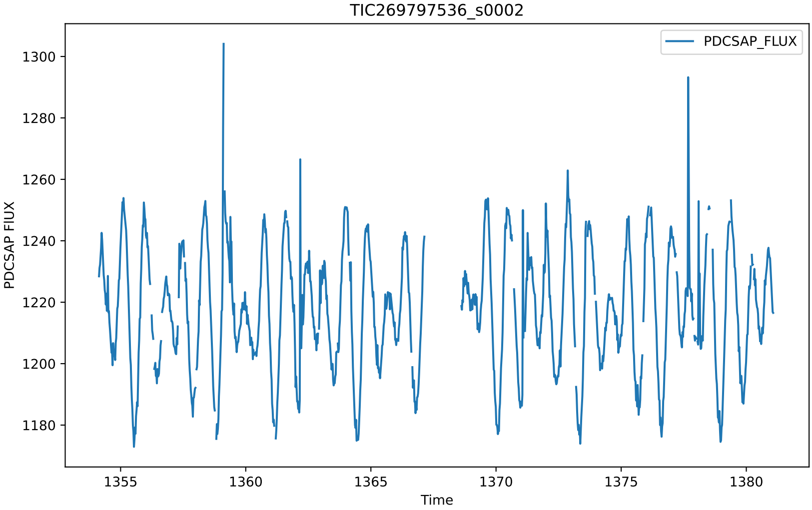

TIC269797536-s2-I

non-parametric detrend

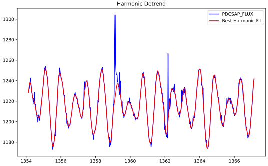

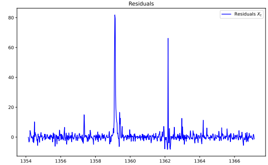

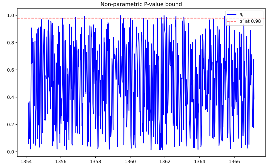

parametric flares detection

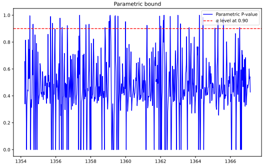

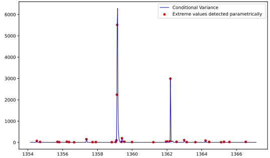

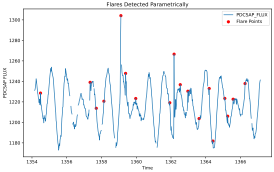

TIC269797536-s2-II

non-parametric detrend

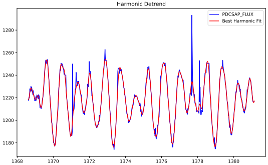

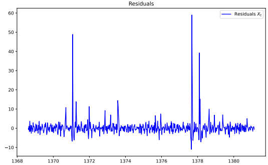

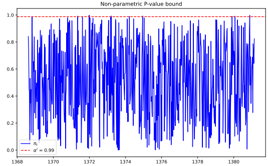

parametric flares detection

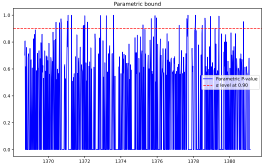

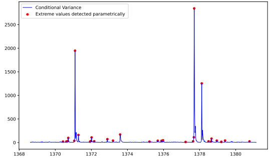

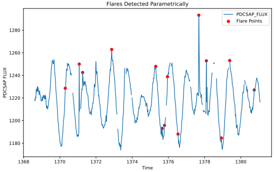

That's it

Thanks ☺

ARFIMA

- The Autoregressive Fractionally Integrated Moving Average (ARFIMA) model extends the classical ARIMA model by incorporating a fractional differencing parameter, enabling the model to capture both short-term autocorrelation and long-range dependence within time series data.
- Particularly valuable for analyzing data sets where phenomena persist over time, such as in economics and finance.

$$\left(1 - \sum_{i=1}^{p} \phi_i L^i\right) (1 - L)^d Y_t = \left(1 + \sum_{j=1}^{q} \theta_j L^j\right) \varepsilon_t,$$

where:

L is the lag operator such that $LY_t = Y_{t-1}$,

p is the order of the autoregressive (AR) part,

d is the order of fractional differencing, allowing the model to handle processes with long memory,

q is the order of the moving average (MA) part,

ϕ_i are the coefficients of the AR part,

θ_j are the coefficients of the MA part,

ε_t is the error term, assumed to be white noise.

ARFIMA

- The unconditional variance of the model is influenced by the parameters of both the AR and MA parts as well as the differencing parameter d, which affects the long-term dependence properties of the series.
- The conditional mean of the ARFIMA model is given by the AR part:

$$E\left(Y_t \mid Y_{t-1}, Y_{t-2}, \ldots\right) = \sum_{i=1}^{p} \phi_i Y_{t-i}$$

  while the conditional variance remains constant.
- Primary motivation ARFIMA: their ability to model series with highly complex stochastic structures that exhibit both short-term and long-term dependencies.
- For instance, financial time series often display volatility clustering (short-term) and long memory (long-term dependencies in volatility), both of which can be effectively captured by the ARFIMA model.

ARFIMA-GARCH fit with 1 iteration

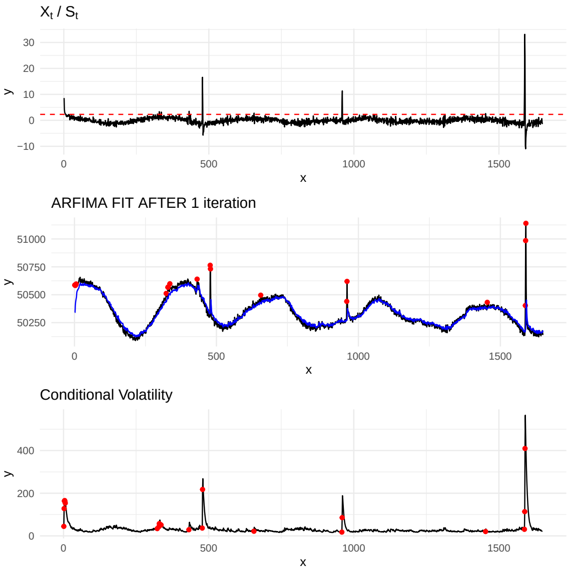

ARFIMA-GARCH fit after 5 iteration