NP-FOCuS: Online Nonparametric Changepoint Detection via Functional Pruning LR tests

Gaetano Romano, Idris Eckley, Paul Fearnhead
14-16 Dec 2022

The Setting
Offline Changepoint Detection

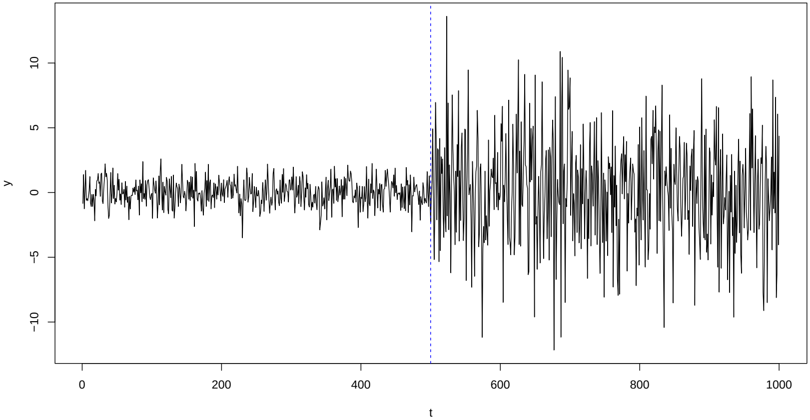

Online changepoint detection

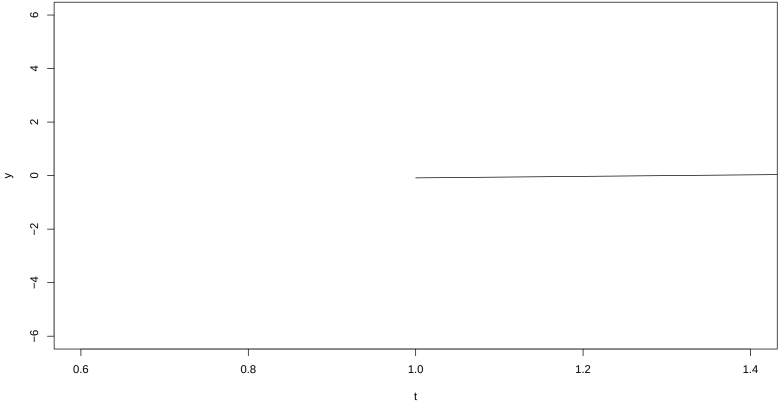

One change...
Online changepoint detection

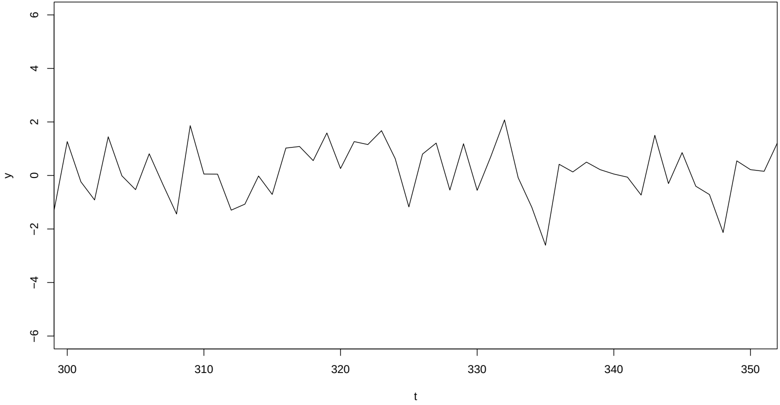

The Setting

Online changepoint detection

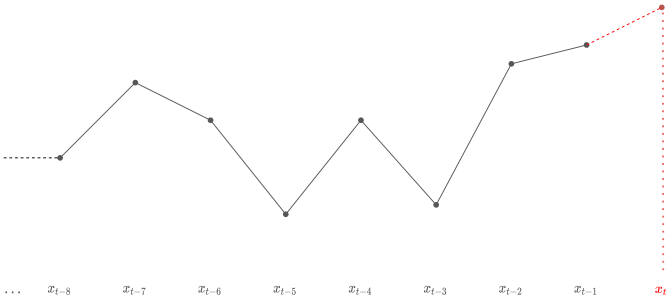

At each iteration, we have to decide whether we are observing a change or not.

The Setting

Comparison to the offline setting

Compared to the offline setting we have additional requirements:

- The method needs to be sequential;
- The methodology needs to process observations at a faster rate than the sampling rate.

Delayed Evaluation

What happens if our method is too slow?

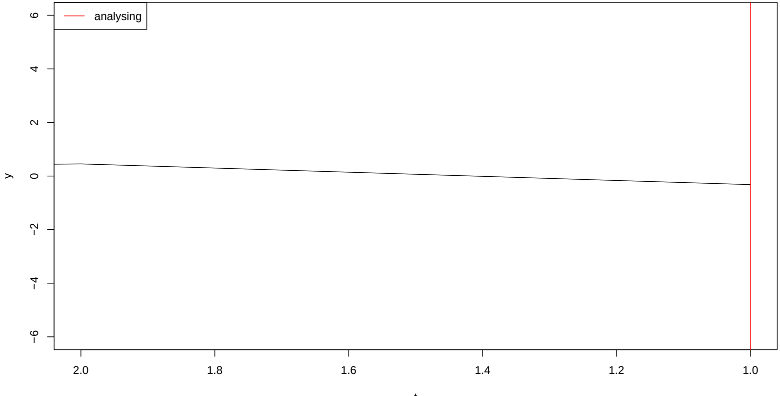

Different Scenarios

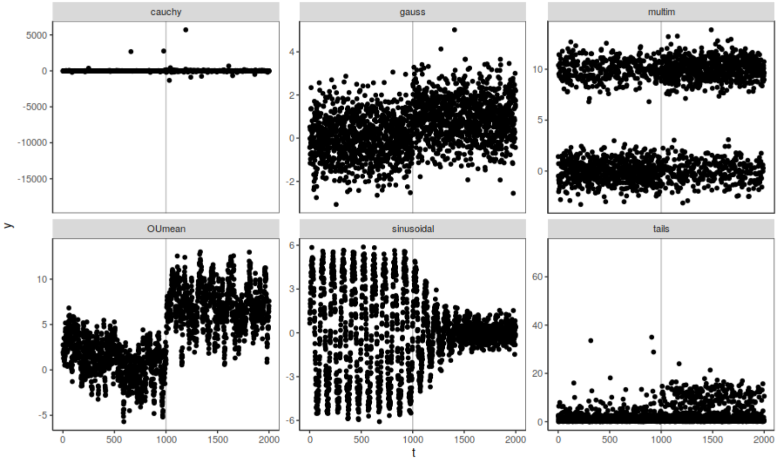

Non-Parametric Approach I

Empirical Cumulative Distribution Function (eCDF).

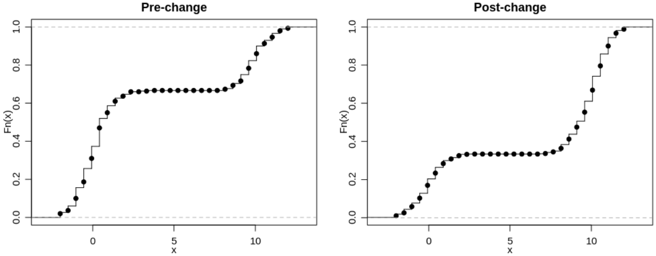

Figure: Pre and post change eCDF of the multimodal example

Non-Parametric Approach II

Let $\theta = F_{1:n}(p)$ be the unknown CDF for $y_1, \ldots, y_n$, and let $\hat{F}_{1:n}(p)$ be its relative eCDF:

$$\hat{\theta} = \hat{F}_{1:n}(p) = \frac{1}{n} \sum_{t=1}^{n} \left[\mathbb{I}(y_t < p)\right],$$

for a fixed value of $p \in \mathbb{R}$, a quantile of our full eCDF.

Non-Parametric Approach III

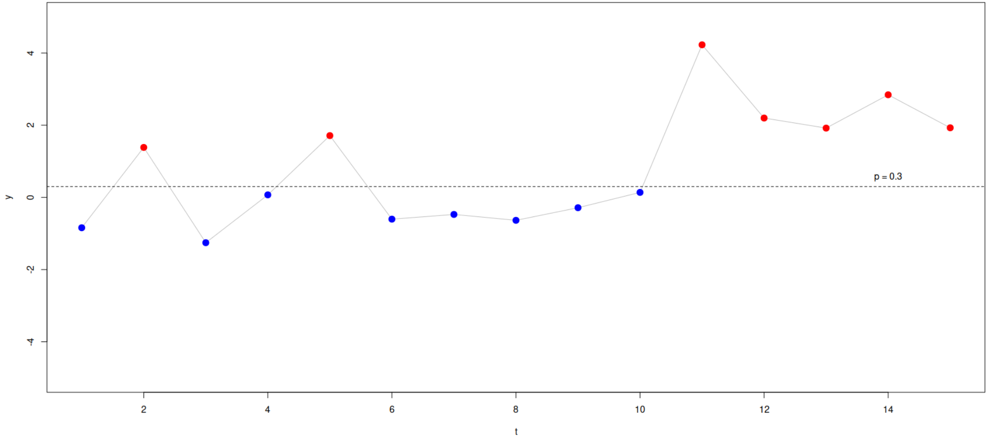

The Cost Function

For a given p the eCDF follows:

$$n\hat{F}_{1:n}(p) \sim \mathrm{Binom}(n, \theta).$$

The contribution of one point $x_t = \mathbb{I}(y_t < p)$ will be given by:

$$h(x_t, \theta) = x_t \log \theta + (1 - x_t) \log(1 - \theta).$$

A Likelihood Ratio test I

For a change in one component of our eCDF, we can build a likelihood ratio test of the form:

$$\mathcal{Q}_n = -\max_{\theta \in \mathbb{R}} \sum_{t=1}^{n} h(x_t, \theta) + \max_{\substack{\tau \in \{1,\ldots,n-1\} \\ \theta_0, \theta_1 \in \mathbb{R}}} \left\{ \sum_{t=1}^{\tau} h(x_t, \theta_0) + \sum_{t=\tau+1}^{n} h(x_t, \theta_1) \right\}, \quad (1)$$

A Likelihood Ratio test II

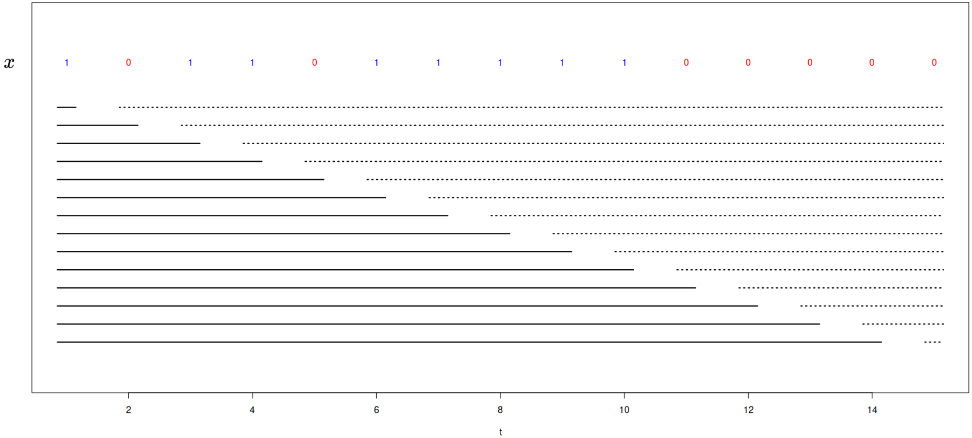

Computational Complexity

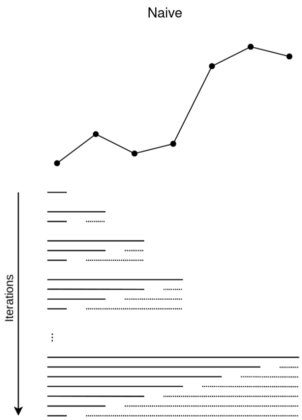

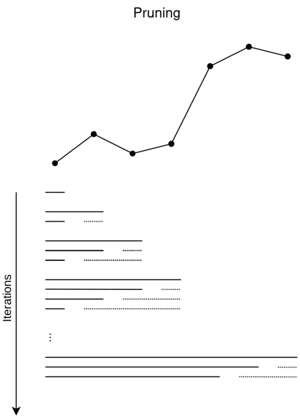

NP-FOCuS

Recall that:

$$h(x_t, \theta) = x_t \log \theta + (1 - x_t) \log(1 - \theta).$$

We can then solve our LR test via the recursion:

$$Q_0(\theta) = 0$$

and for $n = 1, 2, 3, \ldots$

$$Q_n(\theta) = \max\{0,\ Q_{n-1}(\theta)\} + h(x_n, \theta) - \max_{\theta} \sum_{t=1}^{n} h(x_t, \theta).$$

The recursion

Starting condition: our $Q_{t-1}(\theta)$

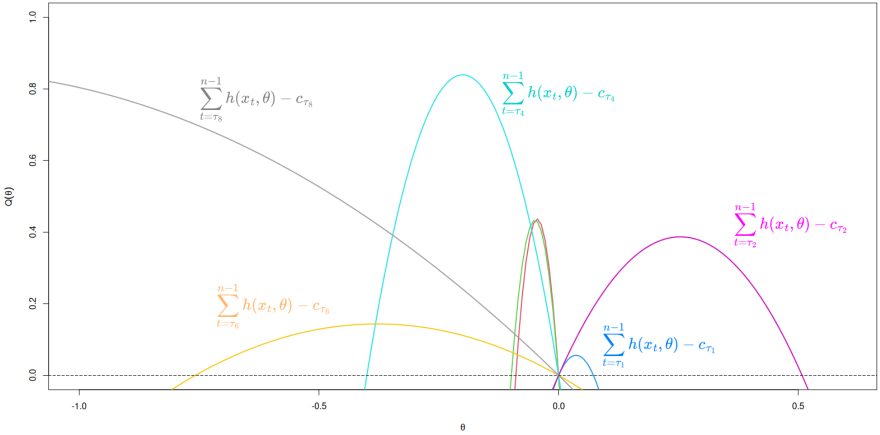

The recursion

$\max\{0,\ Q_{n-1}(\theta)\} + h(x_n, \theta).$

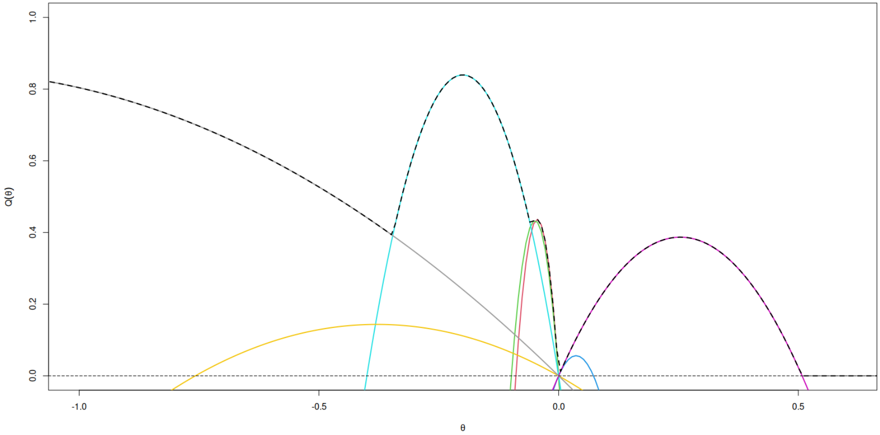

The recursion

$\max\{0,\ Q_{n-1}(\theta)\} + h(x_n, \theta).$

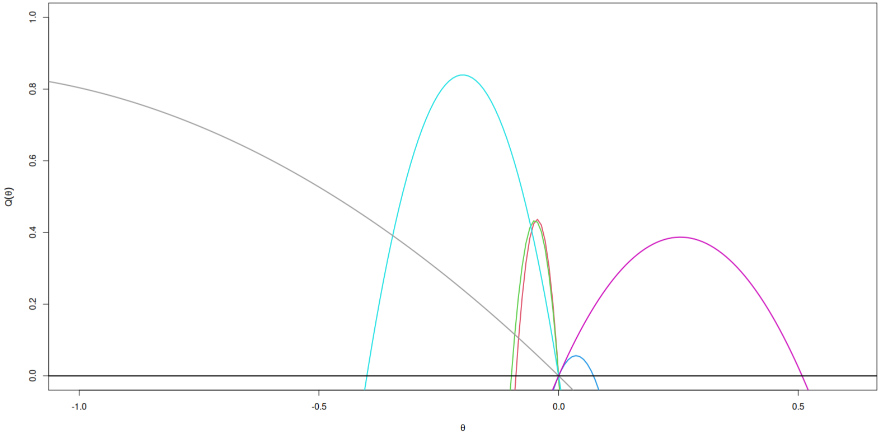

The recursion

$Q_n(\theta) = \max\{0,\ Q_{n-1}(\theta)\} + h(x_n, \theta).$

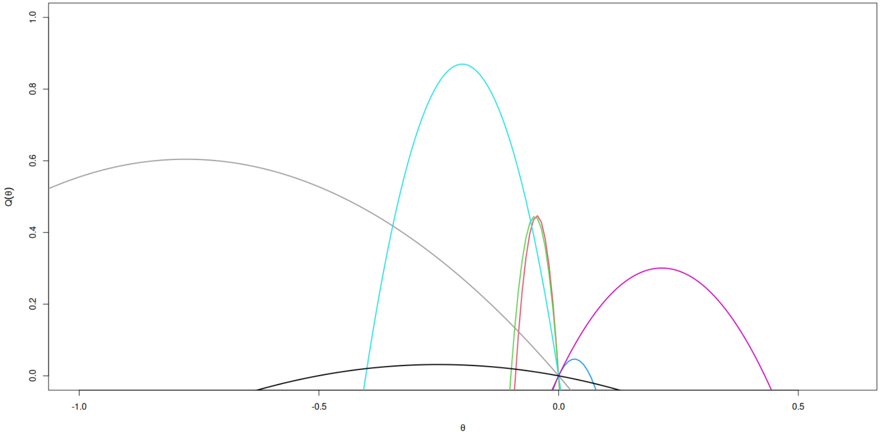

Aggregate the quantile traces
NP-FOCuS

At each time we have m piecewise costs, $Q_n^1(\theta), \dots, Q_n^M(\theta)$, one for each $p_1, \dots, p_m$. By computing the global maximum for each cost:

$$\mathcal{Q}_n^m = \max_\theta Q_n^m(\theta).$$

Then, we will detect a changepoint at time n whether:

$$\sum_{m=1}^{M} \mathcal{Q}_n^m \geq \xi^{sum} \text{ or } \max_{m \in \{1,\dots,M\}} \mathcal{Q}_n^m \geq \xi^{max},$$

with $\xi^{sum}, \ \xi^{max} \in \mathbb{R}$.

NP-FOCuS

Some remarks

- The LR test is solved exactly
- We avoid additional grid parameters
- Low computational complexity, $M \times \mathcal{O}(log(n))$ per iteration
- Taking max and sum for sparse and dense changes in the eCDF.

Empirical results

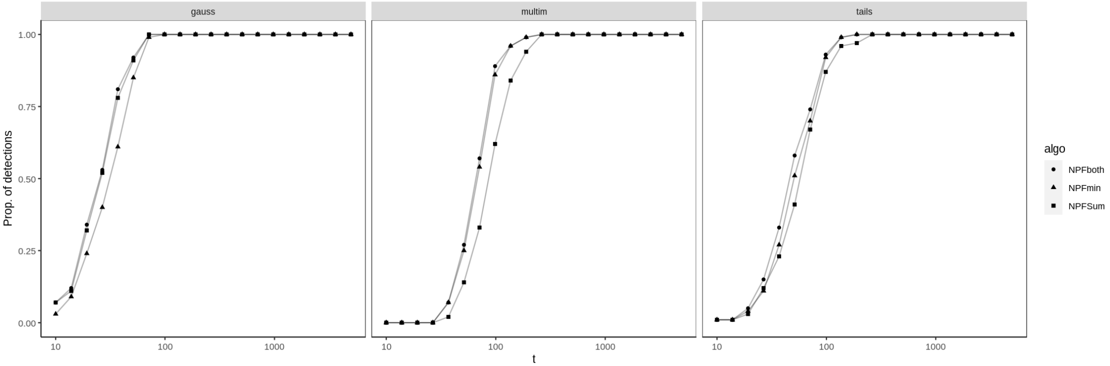

Thank you for your attention