

The Sample Complexity of Multiple Change Point Identification under Bandit Feedback

Maximilian Graf^{1,†}


Victor Thuot^{2,3,†}

¹Institut für Mathematik, Universität Potsdam

²MISTEA, INRAE, Institut Agro, Univ. Montpellier

³UMPA, ENS Lyon

[†]Equal contribution



Paper (arXiv)

Motivation


Change points are common: queueing • materials science • parameter tuning.
When **evaluations are costly** $\Rightarrow$ sample **adaptively** (bandit feedback).

Model and bandit feedback

Unknown **piecewise-constant** f with m change points:

$$f : [0, 1] \rightarrow \mathbb{R}, \quad f(x) := \mu_0 + \sum_{i=1}^m \Delta_i 1\{x \geq x_i^*\}$$

× past observations $(x_s, y_s)_{s < t}$
ⓘ t -th observation $y_t = f(x_t) + \varepsilon_t$



change points: $x_1^* < \dots < x_m^*$
jumps: $\Delta_i := f(x_i^*) - f(x_{i-1}^*)$
spacings: $\vartheta_i := x_{i+1}^* - x_i^*$
($\vartheta_0 = \vartheta_m = 1$)
local spacing: $s_i := \vartheta_{i-1} \wedge \vartheta_i$
energy: $\mathcal{E}_i^2 := s_i \Delta_i^2$

Protocol: sequential & adaptive sampling

for $t = 1, 2, \dots$ **do**
 choose $x_t \in [0, 1]$ $\triangleright$ based on past observations
 observe $y_t = f(x_t) + \varepsilon_t$ $\triangleright \varepsilon_t$ 1-sub-Gaussian noise
 if stop then $\mathcal{T} \leftarrow t$, **output** $\hat{x}_1, \dots, \hat{x}_N$ $\triangleright$ learner chooses when to stop

(δ, η, N) -correctness: for $N \leq m$, the output $\hat{x}_1, \dots, \hat{x}_N$ satisfies

$$\mathbb{P}(\exists i_1 < \dots < i_N \leq m : \forall l \in [N], |\hat{x}_l - x_{i_l}^*| \leq \eta) \geq 1 - \delta.$$

Goal: be (δ, η, N) -correct with a **sample complexity** $\mathcal{T}$ (number of queries) as small as possible.

Two difficulty measures

Localization $H_{\text{localize}} := \sum_{l=1}^N \frac{1}{\Delta_l^2}$ cost of *localizing* the N largest jumps
all jumps = $|\Delta|$; $H_{\text{localize}} = N/\Delta^2$

Detection $H_{\text{detect}} := \max_{i \leq m} \frac{1}{\mathcal{E}_i^2}$ cost of *detecting* the lowest-energy CP
all jumps = $|\Delta|$; $H_{\text{detect}} = \max_i s_i^{-1}/\Delta^2$

Remark: $H_{\text{detect}} \geq \frac{1}{2} H_{\text{localize}}$, **much larger** if CPs are close ($\min_i s_i \ll 1/N$).

MCPI [1]: $\mathbb{E}[\mathcal{T}] \underset{\delta \rightarrow 0}{\sim} 8 H_{\text{localize}} \log \frac{1}{\delta}$, jumps only.

Research question. What is the *optimal* $\mathbb{E}[\mathcal{T}]$ for any δ, η (even moderate), in terms of the jumps Δ_i and spacings s_i ?


Algorithm LCP

Algorithm: LocalizeChangePoints (LCP)

Input: target N , confidence δ , precision η $\triangleright m, \Delta_i, s_i$ unknown
 $k \leftarrow \lceil \log_2(2N) \rceil$
while true do $\triangleright$ doubling trick: no need to know $H_{\text{detect}}, H_{\text{localize}}$
 $(c_v, \text{ok}_v)_{v \leq N} \leftarrow \text{Fixed_Budget_LCP}(2^k, \delta^{(k)})$
 $\triangleright$ 1 detect, 2 estimate, 3 localize, 4 verify (see below)
 if all ok_v **then return** $c_1, \dots, c_N$ $\triangleright \delta^{(k)} \propto \delta/(Nk^2)$: error $\leq \delta$ (union bound)
 $k \leftarrow k + 1$ $\triangleright$ double the budget

Subroutine Fixed_Budget_LCP($T, \delta^{(k)}$), step by step


1 Detect: which intervals contain a CP? $\mathcal{I} \leftarrow \text{DetectIntervals}(T)$
How: multiscale tests; budget split evenly over depths $d \leq d_{\max} \approx \log_2 T$: each point of the grid $\{j/2^d\}$ gets $T_d = \lfloor T/(d_{\max}(2^d + 1)) \rfloor$ samples. **Flag** an interval if $|\hat{y}_{\text{right}} - \hat{y}_{\text{left}}| > \beta_d \approx \sqrt{\log(d_{\max} 2^d)/T_d}$, set so that w.p. $\geq 15/16$, **no false alarm**.



Why energy: depth d flags x_i^* if it is **alone** ($2^{-d} \lesssim s_i$) and its jump beats β_d ($\Delta_i^2 \gtrsim 2^d/T$, up to logs); some depth works if $T \gtrsim \mathcal{E}_i^{-2}$.

2 Estimate: how large is each jump? $(J_v, \hat{\Delta}_v)_{v \leq N} \leftarrow \text{EstimateJumps}(\mathcal{I}, T)$
How: sample both endpoints of each flagged interval $J_v = [l_v, r_v]$ in doubling batches, until the estimate $\hat{\Delta}_v$ of its **net jump** $\Delta_v := f(r_v) - f(l_v)$ is significant; keep the N largest.
Guarantee: w.p. $\geq 15/16$, $\frac{1}{2} \hat{\Delta}_v \leq |\Delta_v| \leq \frac{3}{2} \hat{\Delta}_v$ (accurate up to a factor 2).
Why: split the budget as $T_v \propto \hat{\Delta}_v^{-2} T$: the **oracle allocation**, up to constants.

3 Localize: binary search with backtracking [2] $c_v \leftarrow \text{SHB}(J_v, T_v, \eta)$
How: $R = \lceil 6 \log(|J_v|/\eta) \rceil = O(\log \frac{1}{\eta})$ rounds, **same budget** T_v/R each: keep the half holding the jump, **backtrack** after a wrong turn; c_v = final midpoint.
Why: if J_v holds a single CP, precision η costs only $\hat{\Delta}_v^{-2} \log \frac{1}{\delta \eta}$, **logarithmic** in $1/\eta$.



4 Verify: is c_v really η -close to a CP? $\text{ok}_v \leftarrow \text{VerifyCP}(c_v \pm \eta, \delta^{(k)})$
How: two-sample test between $c_v - \eta$ and $c_v + \eta$, at level $\delta^{(k)}$, with budget T_v .
Why: small CPs can mimic one large jump; a positive test certifies a CP η -close to c_v .
✓ all $\text{ok}_v \rightarrow$ LCP returns $c_1, \dots, c_N$
✗ else: $k \leftarrow k + 1$ (double the budget), back to step 1

Upper bound: LCP is non-asymptotic

Theorem 1. For all $\delta, \eta < \frac{1}{4}$, LCP is (δ, η, N) -correct, and, up to $\log(H_{\text{detect}})$ factors:

$$\mathbb{E}[\mathcal{T}] \lesssim H_{\text{detect}} + H_{\text{localize}} \log \frac{1}{\delta \eta}$$

Lower bound: energy is unavoidable

Theorem 2. Let $N = m$, $\delta < \frac{1}{16}$, $\eta < \frac{1}{8}$, change points 2η -apart. For **any** (δ, η, m) -correct algorithm and any instance, there is an instance with comparable $H_{\text{detect}}, H_{\text{localize}}$ on which:


$$\mathbb{E}[\mathcal{T}] \gtrsim H_{\text{detect}} + H_{\text{localize}} \log \frac{1}{\delta} + \sum_i \frac{1}{\Delta_i^2} \log_+ \frac{s_i}{\eta}$$

Key tool: a new reduction from bandit to **batch signal detection**.


Bounds match ($N = m$, up to logs): LCP is **rate-optimal**.

Experiments: LCP vs. MCPI [1]


$\Delta_1 \uparrow, \Delta_2 \downarrow$ $\Delta_1 = -\Delta_2 = 1$, Gaussian noise, 1000 runs.
MCPI: discretization with $\asymp 1/\eta$ arms.



Exp. 1: spacing s
LCP $\propto 1/s$
($\eta = 2^{-11}$, $\delta = 0.05$)



Exp. 2: confidence δ
LCP cheaper until
 $\delta \approx e^{-110}$
($\eta = 2^{-8}$, $s = 1/4$)



Exp. 3: precision η
LCP flat, MCPI $\propto 1/\eta$
($\delta = 0.05$, $s = 1/4$)

Main message

Optimal expected sample complexity ($N = m$, up to log factors):

$$\mathbb{E}[\mathcal{T}] \asymp H_{\text{detect}} + H_{\text{localize}} \log \frac{1}{\delta \eta}$$

🔧 Beyond $\delta \rightarrow 0$: jumps **and** spacings matter.

⚡ The **energy** $s_i \Delta_i^2$ drives detection, as in batch [3].

🔍 Precision is cheap: $\log(1/\eta)$ (vs. $1/\eta$ for MCPI).

🔄 LCP adapts to unknown m and s_i .

📌 **Remark:** for the $(1-\delta)$ -quantile of $\mathcal{T}$, the optimal scaling is also characterized, and **differs**:

$$H_{\text{detect}} \log \frac{1}{\delta} + H_{\text{localize}} \log \frac{1}{\delta \eta}$$