Learning Under Non-Stationarity:
Statistical Inference & Bandit Convex Optimization

Xiaoqi (Shirley) Liu, University of Oxford

Algorithms and Computationally Intensive Inference Seminar
Department of Statistics, University of Warwick

13 March 2026

High-dimensional changepoint localization

High-dimensional changepoint localization

Offline inference:

Observe patients' medical features & heart failure outcomes

Relationship between features & outcome change markedly at some age

? At what age does change happen?

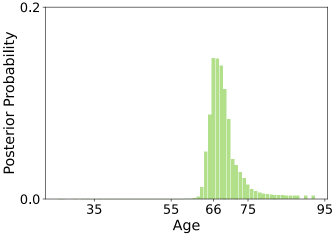

High-dimensional changepoint localization

Offline inference:

Observe patients' medical features & heart failure outcomes

Relationship between features & outcome change markedly at some age

? At what age does change happen?

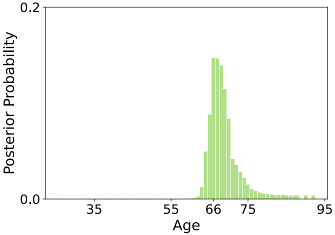

Non-stationary bandit convex optimization

High-dimensional changepoint localization

Offline inference:
Observe patients' medical features & heart failure outcomes
Relationship between features & outcome change markedly at some age
? At what age does change happen?

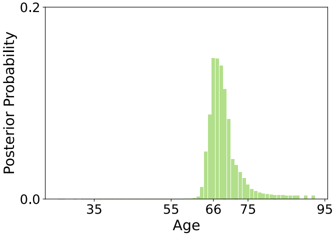

Non-stationary bandit convex optimization

Adaptive decision-making:
Observe outcomes *only* for the chosen treatment
Treatment effects change over time
? How to balance exploration-exploitation?

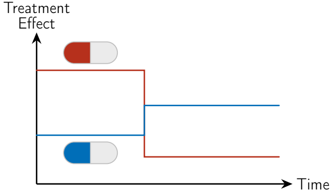


Statistical aim
Interactivity

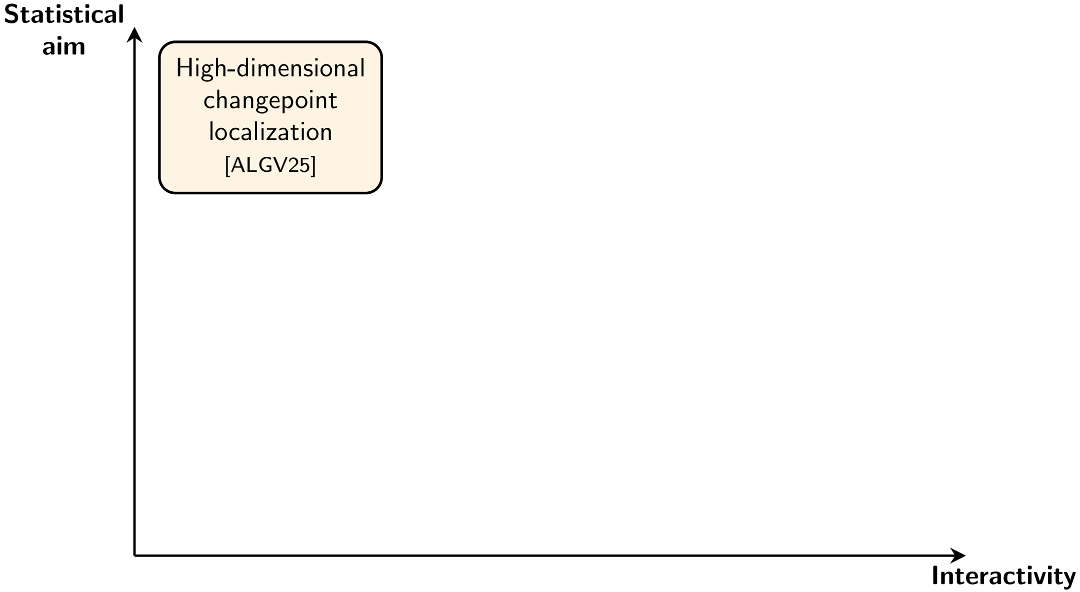
Statistical aim
High-dimensional changepoint localization [ALGV25]
Interactivity

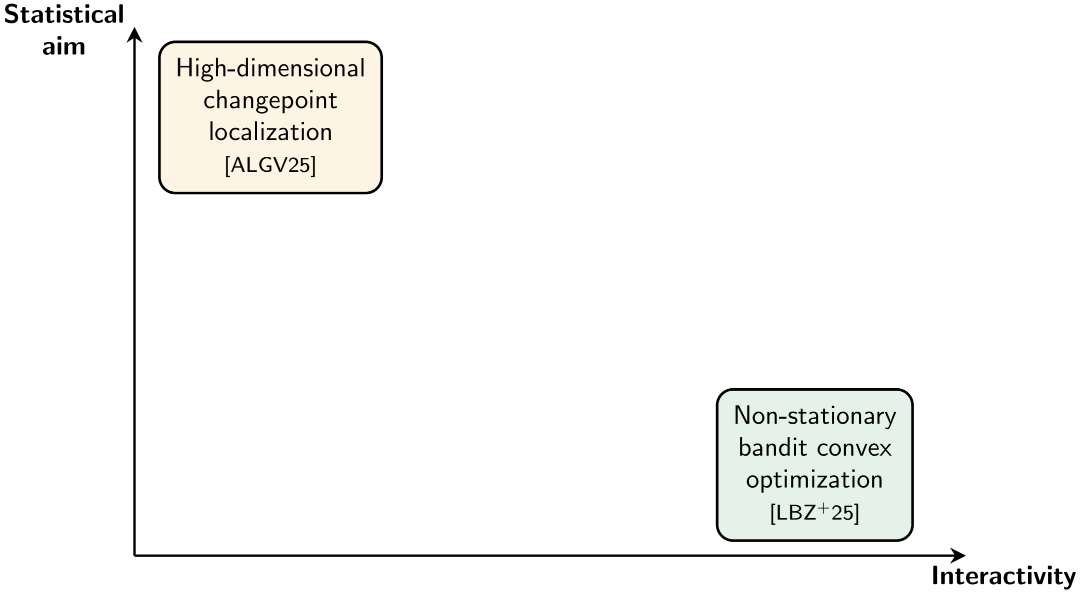
Statistical aim
High-dimensional changepoint localization [ALGV25]
Non-stationary bandit convex optimization [LBZ+25]
Interactivity

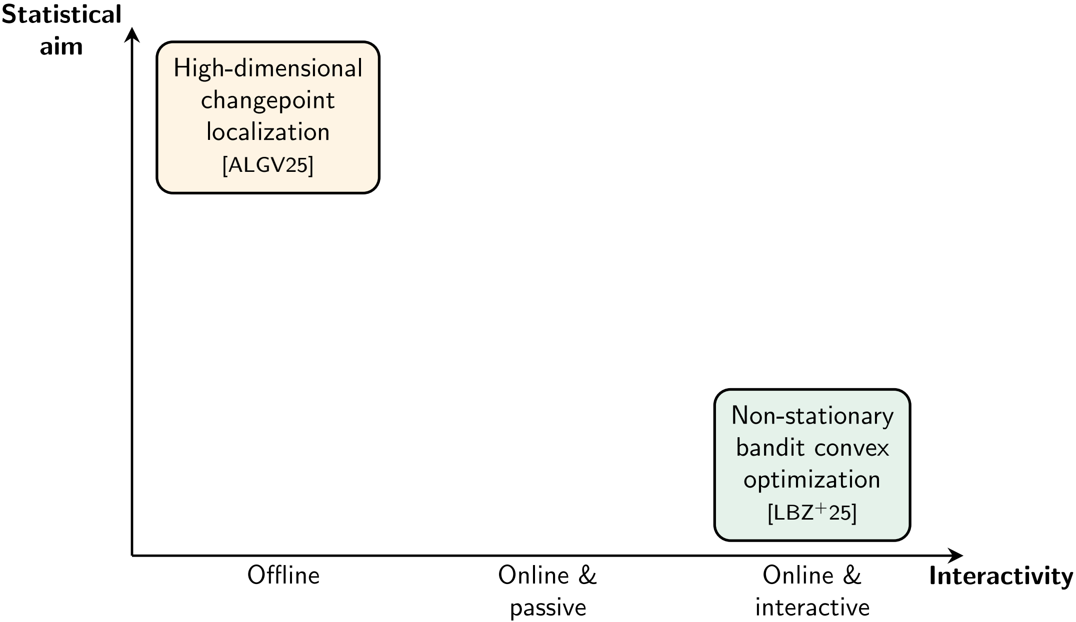
Statistical aim
High-dimensional changepoint localization [ALGV25]
Non-stationary bandit convex optimization [LBZ+25]
Offline
Online & passive
Online & interactive
Interactivity

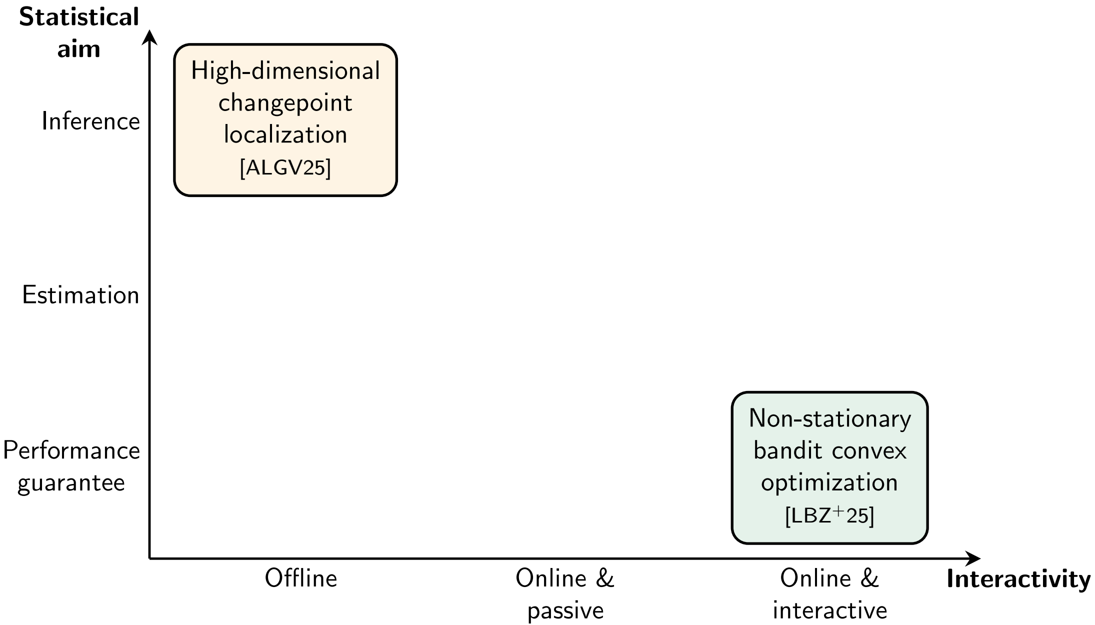
Statistical aim
Inference
Estimation
Performance guarantee
High-dimensional changepoint localization [ALGV25]
Non-stationary bandit convex optimization [LBZ+25]
Offline
Online & passive
Online & interactive
Interactivity

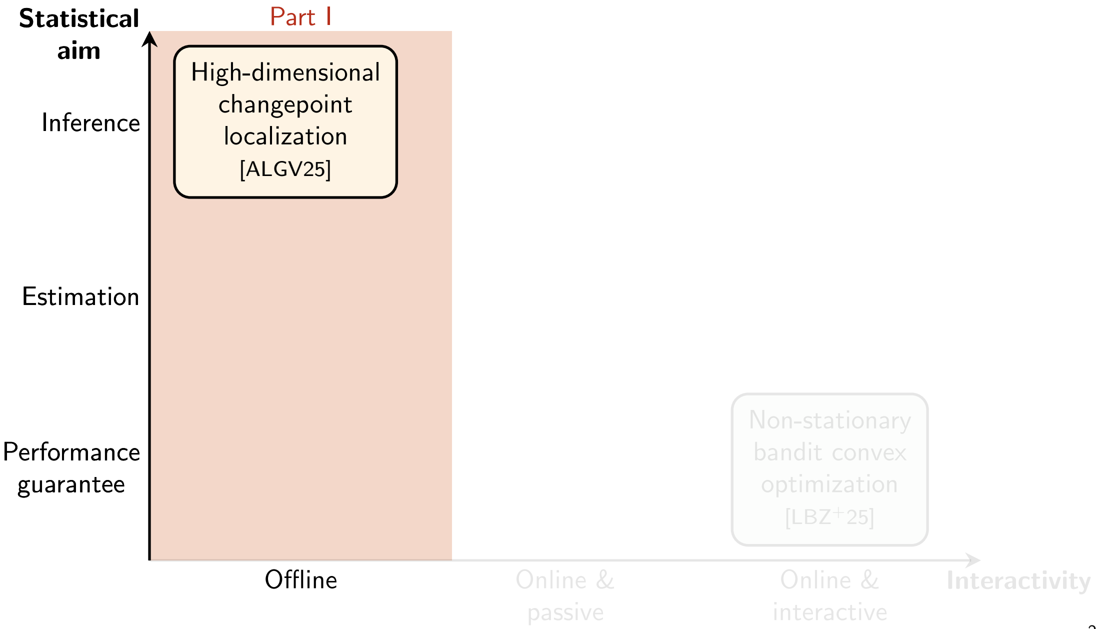

Part I
Statistical aim
Inference
Estimation
Performance guarantee
High-dimensional changepoint localization [ALGV25]
Non-stationary bandit convex optimization [LBZ+25]
Offline
Online & passive
Online & interactive
Interactivity

High-dimensional changepoint inference

High-dimensional changepoint inference

For $i = 1, 2, \ldots, n$, $\qquad y_i = \boldsymbol{x}_i^\top \boldsymbol{\beta}^{(1)} + \varepsilon_i$

$\boldsymbol{x}_i, \boldsymbol{\beta}^{(\ell)} \in \mathbb{R}^d$

High-dimensional changepoint inference

For $i = 1, 2, \ldots, n$,

$\boldsymbol{x}_i, \boldsymbol{\beta}^{(\ell)} \in \mathbb{R}^d$

$$y_i = \boldsymbol{x}_i^\top \boldsymbol{\beta}^{(1)} + \varepsilon_i$$

High-dimensional changepoint inference

For $i = 1, 2, \ldots, n$,

$\boldsymbol{x}_i, \boldsymbol{\beta}^{(\ell)} \in \mathbb{R}^d$

$$y_i = \boldsymbol{x}_i^\top \boldsymbol{\beta}^{(1)} + \varepsilon_i$$

$i = \eta_1$

High-dimensional changepoint inference

For $i = 1, 2, \ldots, n,$

$\boldsymbol{x}_i, \boldsymbol{\beta}^{(\ell)} \in \mathbb{R}^d$

$$y_i = \boldsymbol{x}_i^\top \boldsymbol{\beta}^{(1)} + \varepsilon_i$$

$$i = \eta_1$$

$$y_i = \boldsymbol{x}_i^\top \boldsymbol{\beta}^{(2)} + \varepsilon_i$$

High-dimensional changepoint inference

For $i = 1, 2, \ldots, n$,

$\boldsymbol{x}_i, \boldsymbol{\beta}^{(\ell)} \in \mathbb{R}^d$

$$y_i = \boldsymbol{x}_i^\top \boldsymbol{\beta}^{(1)} + \varepsilon_i$$

$i = \eta_1$

$$y_i = \boldsymbol{x}_i^\top \boldsymbol{\beta}^{(2)} + \varepsilon_i$$

$\vdots$

$i = \eta_{L^*-1}$

High-dimensional changepoint inference

For $i = 1, 2, \ldots, n$,
$\boldsymbol{x}_i, \boldsymbol{\beta}^{(\ell)} \in \mathbb{R}^d$

$$y_i = \boldsymbol{x}_i^\top \boldsymbol{\beta}^{(1)} + \varepsilon_i$$

$i = \eta_1$

$$y_i = \boldsymbol{x}_i^\top \boldsymbol{\beta}^{(2)} + \varepsilon_i$$

$\vdots$

$i = \eta_{L^*-1}$

L^* is unknown constant,
with $L^* < L$
$\boldsymbol{\eta} := [\eta_1, \ldots, \eta_{L^*-1}]$,
with $\eta_{\ell+1} - \eta_\ell \asymp n$

High-dimensional changepoint inference

For $i = 1, 2, \ldots, n$,
$\boldsymbol{x}_i, \boldsymbol{\beta}^{(\ell)} \in \mathbb{R}^d$

$$y_i = \boldsymbol{x}_i^\top \boldsymbol{\beta}^{(1)} + \varepsilon_i$$

$i = \eta_1$

$$y_i = \boldsymbol{x}_i^\top \boldsymbol{\beta}^{(2)} + \varepsilon_i$$

$\vdots$

$i = \eta_{L^*-1}$

L^* is unknown constant,
with $L^* < L$
$\boldsymbol{\eta} := [\eta_1, \ldots, \eta_{L^*-1}]$,
with $\eta_{\ell+1} - \eta_\ell \asymp n$

High-dimensional changepoint inference

For $i = 1, 2, \ldots, n$,
$\boldsymbol{x}_i, \boldsymbol{\beta}^{(\ell)} \in \mathbb{R}^d$

$$y_i = \boldsymbol{x}_i^\top \boldsymbol{\beta}^{(1)} + \varepsilon_i$$

$i = \eta_1$

$$y_i = \boldsymbol{x}_i^\top \boldsymbol{\beta}^{(2)} + \varepsilon_i$$

$\vdots$

$i = \eta_{L^*-1}$

L^* is unknown constant,
with $L^* < L$
$\boldsymbol{\eta} := [\eta_1, \ldots, \eta_{L^*-1}]$,
with $\eta_{\ell+1} - \eta_\ell \asymp n$

Model: entries of $\boldsymbol{x}_i \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, \frac{1}{n})$, entries of $\boldsymbol{\beta}^{(\ell)} \overset{\text{i.i.d.}}{\sim} p_{\bar{\beta}}$

High-dimensional changepoint inference

For $i = 1, 2, \dots, n$,
$\boldsymbol{x}_i, \boldsymbol{\beta}^{(\ell)} \in \mathbb{R}^d$

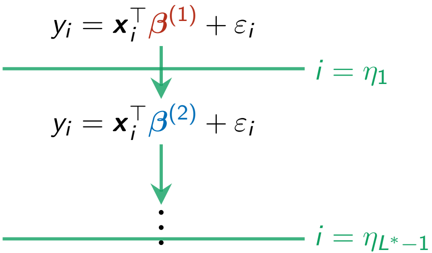

L^* is unknown constant,
with $L^* < L$
$\boldsymbol{\eta} := [\eta_1, \dots, \eta_{L^*-1}]$,
with $\eta_{\ell+1} - \eta_\ell \asymp n$

Model: entries of $\boldsymbol{x}_i \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, \frac{1}{n})$, entries of $\boldsymbol{\beta}^{(\ell)} \overset{\text{i.i.d.}}{\sim} p_{\bar{\beta}}$

Goal: infer changepoint locations $(\eta_\ell)_{\ell=1}^{L^*-1}$ from $(\boldsymbol{x}_i, y_i)_{i=1}^n$, and estimate signals $(\boldsymbol{\beta}^{(\ell)})_{\ell=1}^{L^*}$

High-dimensional changepoint inference

For $i = 1, 2, \dots, n$,
$\boldsymbol{x}_i, \boldsymbol{\beta}^{(\ell)} \in \mathbb{R}^d$

$$y_i = \boldsymbol{x}_i^\top \boldsymbol{\beta}^{(1)} + \varepsilon_i$$

$i = \eta_1$

$$y_i = \boldsymbol{x}_i^\top \boldsymbol{\beta}^{(2)} + \varepsilon_i$$

$\vdots$

$i = \eta_{L^*-1}$

L^* is unknown constant,
with $L^* < L$
$\boldsymbol{\eta} := [\eta_1, \dots, \eta_{L^*-1}]$,
with $\eta_{\ell+1} - \eta_\ell \asymp n$

Model: entries of $\boldsymbol{x}_i \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, \frac{1}{n})$, entries of $\boldsymbol{\beta}^{(\ell)} \overset{\text{i.i.d.}}{\sim} p_{\bar{\beta}}$

Goal: infer changepoint locations $(\eta_\ell)_{\ell=1}^{L^*-1}$ from $(\boldsymbol{x}_i, y_i)_{i=1}^n$, and estimate signals $(\boldsymbol{\beta}^{(\ell)})_{\ell=1}^{L^*}$

Scaling regime: $n, d \to \infty$ with $n/d \to$ fixed constant δ

For $d \ll n$,

For $d \ll n$,

$$\hat{\eta}_1, \ldots, \hat{\eta}_{L^*-1} = \underset{\tilde{\eta}_1 \le \tilde{\eta}_2 \le \cdots \le \tilde{\eta}_{L^*-1}}{\arg\min} \sum_{\ell=1}^{L^*} \min_{\tilde{\boldsymbol{\beta}} \in \mathbb{R}^d} \sum_{i=\tilde{\eta}_{\ell-1}+1}^{\tilde{\eta}_\ell} \left(y_i - \boldsymbol{x}_i^\top \tilde{\boldsymbol{\beta}}\right)^2$$

For $d \ll n$,

$$\hat{\eta}_1, \ldots, \hat{\eta}_{L^*-1} = \underbrace{\operatorname*{arg\,min}_{\tilde{\eta}_1 \leq \tilde{\eta}_2 \leq \cdots \leq \tilde{\eta}_{L^*-1}} \sum_{\ell=1}^{L^*}}_{\substack{\text{search through} \\ \text{different partitioning}}} \min_{\tilde{\boldsymbol{\beta}} \in \mathbb{R}^d} \sum_{i=\tilde{\eta}_{\ell-1}+1}^{\tilde{\eta}_\ell} \left(y_i - \boldsymbol{x}_i^\top \tilde{\boldsymbol{\beta}}\right)^2$$

For $d \ll n$,

$$\hat{\eta}_1, \ldots, \hat{\eta}_{L^*-1} = \underbrace{\mathop{\arg\min}_{\tilde{\eta}_1 \leq \tilde{\eta}_2 \leq \cdots \leq \tilde{\eta}_{L^*-1}} \sum_{\ell=1}^{L^*}}_{\text{search through different partitioning}} \underbrace{\min_{\tilde{\boldsymbol{\beta}} \in \mathbb{R}^d} \sum_{i=\tilde{\eta}_{\ell-1}+1}^{\tilde{\eta}_\ell} \left(y_i - \boldsymbol{x}_i^\top \tilde{\boldsymbol{\beta}}\right)^2}_{\text{Residual Sum of Squares}}$$

[Bai97, BP98, Jul01]

For $d \ll n$,

$$\hat{\eta}_1, \ldots, \hat{\eta}_{L^*-1} = \underbrace{\operatorname*{arg\,min}_{\tilde{\eta}_1 \leq \tilde{\eta}_2 \leq \cdots \leq \tilde{\eta}_{L^*-1}} \sum_{\ell=1}^{L^*}}_{\text{search through different partitioning}} \underbrace{\min_{\tilde{\boldsymbol{\beta}} \in \mathbb{R}^d} \sum_{i=\tilde{\eta}_{\ell-1}+1}^{\tilde{\eta}_\ell} \left(y_i - \boldsymbol{x}_i^\top \tilde{\boldsymbol{\beta}}\right)^2}_{\text{Residual Sum of Squares}}$$

[Bai97, BP98, Jul01]

For $d \asymp n$,

For $d \ll n$,

$$\hat{\eta}_1, \ldots, \hat{\eta}_{L^*-1} = \underbrace{\mathop{\arg\min}_{\tilde{\eta}_1 \le \tilde{\eta}_2 \le \cdots \le \tilde{\eta}_{L^*-1}} \sum_{\ell=1}^{L^*}}_{\text{search through different partitioning}} \underbrace{\min_{\tilde{\boldsymbol{\beta}} \in \mathbb{R}^d} \sum_{i=\tilde{\eta}_{\ell-1}+1}^{\tilde{\eta}_\ell} \left(y_i - \boldsymbol{x}_i^\top \tilde{\boldsymbol{\beta}}\right)^2}_{\text{Residual Sum of Squares}}$$

[Bai97, BP98, Jul01]

For $d \asymp n$,

Existing estimation methods:

- For sparse $p_{\bar{\beta}}$: partition + penalized maximum likelihood [LB16, RWW+21, BS23]
- For sparse-difference $p_{\bar{\beta}}$: complementary sketching [GW22]
- For non-sparse $p_{\bar{\beta}}$: [CKL25, LQZL26]

For $d \ll n$,

$$\hat{\eta}_1, \ldots, \hat{\eta}_{L^*-1} = \underbrace{\operatorname*{arg\,min}_{\tilde{\eta}_1 \leq \tilde{\eta}_2 \leq \cdots \leq \tilde{\eta}_{L^*-1}} \sum_{\ell=1}^{L^*}}_{\text{search through different partitioning}} \underbrace{\min_{\tilde{\boldsymbol{\beta}} \in \mathbb{R}^d} \sum_{i=\tilde{\eta}_{\ell-1}+1}^{\tilde{\eta}_\ell} \left(y_i - \boldsymbol{x}_i^\top \tilde{\boldsymbol{\beta}}\right)^2}_{\text{Residual Sum of Squares [Bai97, BP98, Jul01]}}$$

For $d \asymp n$,

Existing estimation methods:

- For sparse $p_{\bar{\beta}}$: partition + penalized maximum likelihood [LB16, RWW+21, BS23]
- For sparse-difference $p_{\bar{\beta}}$: complementary sketching [GW22]
- For non-sparse $p_{\bar{\beta}}$: [CKL25, LQZL26]

Our proposed inference method:
Based on Approximate Message Passing [DMM09, BM11, BMN20, FVRS22, GB23]

For $d \ll n$,

$$\hat{\eta}_1, \ldots, \hat{\eta}_{L^*-1} = \underbrace{\operatorname*{arg\,min}_{\tilde{\eta}_1 \leq \tilde{\eta}_2 \leq \cdots \leq \tilde{\eta}_{L^*-1}} \sum_{\ell=1}^{L^*}}_{\text{search through different partitioning}} \underbrace{\min_{\tilde{\boldsymbol{\beta}} \in \mathbb{R}^d} \sum_{i=\tilde{\eta}_{\ell-1}+1}^{\tilde{\eta}_\ell} (y_i - \boldsymbol{x}_i^\top \tilde{\boldsymbol{\beta}})^2}_{\text{Residual Sum of Squares}}$$

Residual Sum of Squares [Bai97, BP98, Jul01]

For $d \asymp n$,

Existing estimation methods:

- For sparse $p_{\bar{\beta}}$: partition + penalized maximum likelihood [LB16, RWW+21, BS23]
- For sparse-difference $p_{\bar{\beta}}$: complementary sketching [GW22]
- For non-sparse $p_{\bar{\beta}}$: [CKL25, LQZL26]

Our proposed inference method:
Based on Approximate Message Passing [DMM09, BM11, BMN20, FVRS22, GB23]

👍 Allow general $p_{\bar{\beta}}$; credible sets with precise asymptotic characterization

For $d \ll n$,

$$\hat{\eta}_1, \ldots, \hat{\eta}_{L^*-1} = \underbrace{\operatorname*{arg\,min}_{\tilde{\eta}_1 \le \tilde{\eta}_2 \le \cdots \le \tilde{\eta}_{L^*-1}} \sum_{\ell=1}^{L^*}}_{\text{search through different partitioning}} \underbrace{\min_{\tilde{\boldsymbol{\beta}} \in \mathbb{R}^d} \sum_{i=\tilde{\eta}_{\ell-1}+1}^{\tilde{\eta}_\ell} \left(y_i - \boldsymbol{x}_i^\top \tilde{\boldsymbol{\beta}}\right)^2}_{\text{Residual Sum of Squares}}$$

search through different partitioning

Residual Sum of Squares [Bai97, BP98, Jul01]

For $d \asymp n$,

Existing estimation methods:

- For sparse $p_{\bar{\beta}}$: partition + penalized maximum likelihood [LB16, RWW+21, BS23]
- For sparse-difference $p_{\bar{\beta}}$: complementary sketching [GW22]
- For non-sparse $p_{\bar{\beta}}$: [CKL25, LQZL26]

Our proposed inference method:

Based on Approximate Message Passing [DMM09, BM11, BMN20, FVRS22, GB23]

👍 Allow general $p_{\bar{\beta}}$; credible sets with precise asymptotic characterization

👎 No minimax optimal localization error

Aside: AMP for sparse linear regression

Aside: AMP for sparse linear regression

$$\boldsymbol{y} = \boldsymbol{X}\boldsymbol{\beta} + \boldsymbol{\varepsilon}$$

entries of $\boldsymbol{X} \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, \frac{1}{n})$

entries of $\boldsymbol{\beta} \overset{\text{i.i.d.}}{\sim} p_{\bar{\beta}}$

Aside: AMP for sparse linear regression

$$\boldsymbol{y} = \boldsymbol{X}\boldsymbol{\beta} + \boldsymbol{\varepsilon} \qquad \text{entries of } \boldsymbol{X} \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, \tfrac{1}{n}) \qquad \text{entries of } \boldsymbol{\beta} \overset{\text{i.i.d.}}{\sim} p_{\bar{\beta}}$$

modified residual: $\boldsymbol{r}_t = \boldsymbol{y} - \boldsymbol{X}\boldsymbol{\beta}_t + b_t \boldsymbol{r}_{t-1}$

signal estimate: $\boldsymbol{\beta}_{t+1} = f_{\substack{\text{soft}\\\text{thres}}}\left(\boldsymbol{\beta}_t + \boldsymbol{X}^\top \boldsymbol{r}_t; \theta_t\right)$

Aside: AMP for sparse linear regression

$$\boldsymbol{y} = \boldsymbol{X}\boldsymbol{\beta} + \boldsymbol{\varepsilon}$$

entries of $\boldsymbol{X} \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, \frac{1}{n})$

entries of $\boldsymbol{\beta} \overset{\text{i.i.d.}}{\sim} p_{\bar{\beta}}$

modified residual: $\boldsymbol{r}_t = \boldsymbol{y} - \boldsymbol{X}\boldsymbol{\beta}_t + \boxed{b_t \boldsymbol{r}_{t-1}}$ debias term

signal estimate: $\boldsymbol{\beta}_{t+1} = f_{\substack{\text{soft}\\\text{thres}}}\left(\boldsymbol{\beta}_t + \boldsymbol{X}^\top \boldsymbol{r}_t; \theta_t\right)$

Aside: AMP for sparse linear regression

$$\boldsymbol{y} = \boldsymbol{X}\boldsymbol{\beta} + \boldsymbol{\varepsilon}$$

entries of $\boldsymbol{X} \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, \frac{1}{n})$

entries of $\boldsymbol{\beta} \overset{\text{i.i.d.}}{\sim} p_{\bar{\beta}}$

modified residual: $\boldsymbol{r}_t = \boldsymbol{y} - \boldsymbol{X}\boldsymbol{\beta}_t + \boxed{b_t \boldsymbol{r}_{t-1}}$ debias term

signal estimate: $\boldsymbol{\beta}_{t+1} = f_{\substack{\text{soft}\\\text{thres}}}\left(\boldsymbol{\beta}_t + \boldsymbol{X}^\top \boldsymbol{r}_t; \theta_t\right)$

entries of $\boldsymbol{\beta}_t + \boldsymbol{X}^\top \boldsymbol{r}_t - \boldsymbol{\beta} \approx \mathcal{N}\left(0, \kappa_t\right)$ as $n, d \to \infty$

Aside: AMP for sparse linear regression

$$\boldsymbol{y} = \boldsymbol{X}\boldsymbol{\beta} + \boldsymbol{\varepsilon}$$

entries of $\boldsymbol{X} \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, \frac{1}{n})$

entries of $\boldsymbol{\beta} \overset{\text{i.i.d.}}{\sim} p_{\bar{\beta}}$

modified residual: $\boldsymbol{r}_t = \boldsymbol{y} - \boldsymbol{X}\boldsymbol{\beta}_t + \boxed{b_t \boldsymbol{r}_{t-1}}$ debias term

signal estimate: $\boldsymbol{\beta}_{t+1} = f_{\substack{\text{soft}\\\text{thres}}}\left(\boldsymbol{\beta}_t + \boldsymbol{X}^\top \boldsymbol{r}_t; \theta_t\right)$

entries of $\boldsymbol{\beta}_t + \boldsymbol{X}^\top \boldsymbol{r}_t - \boldsymbol{\beta} \approx \mathcal{N}\left(0, \kappa_t\right)$ as $n, d \to \infty$

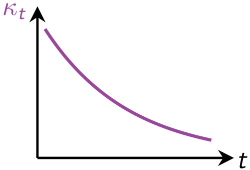

deterministic scalar recursion

AMP for Bayesian changepoint inference

AMP for Bayesian changepoint inference

Signal is $\boldsymbol{B} = \begin{bmatrix} \uparrow & & \uparrow \\ \beta^{(1)} & \cdots & \beta^{(L)} \\ \downarrow & & \downarrow \end{bmatrix} \overset{\text{rows}}{\underset{\text{i.i.d.}}{\sim}} p_{\bar{\boldsymbol{B}}}$ and define $\boldsymbol{\Theta} := \boldsymbol{X}\boldsymbol{B} \overset{\text{rows}}{\underset{\text{i.i.d.}}{\sim}} p_{\bar{\boldsymbol{\Theta}}}$.

AMP for Bayesian changepoint inference

Signal is $\boldsymbol{B} = \begin{bmatrix} \uparrow & & \uparrow \\ \beta^{(1)} & \cdots & \beta^{(L)} \\ \downarrow & & \downarrow \end{bmatrix} \overset{\text{rows i.i.d.}}{\sim} p_{\bar{\boldsymbol{B}}}$ and define $\boldsymbol{\Theta} := \boldsymbol{X}\boldsymbol{B} \overset{\text{rows i.i.d.}}{\sim} p_{\bar{\boldsymbol{\Theta}}}$.

$$\boldsymbol{\Theta}_t = \boldsymbol{X}\widehat{\boldsymbol{B}}_t - \boldsymbol{R}_{t-1}\boldsymbol{F}_t^\top$$

$$\boldsymbol{R}_t = g_t(\boldsymbol{\Theta}_t, \boldsymbol{y})$$

$$\boldsymbol{B}_{t+1} = \boldsymbol{X}^\top \boldsymbol{R}_t - \widehat{\boldsymbol{B}}_t \boldsymbol{C}_t^\top$$

$$\widehat{\boldsymbol{B}}_t = f_t(\boldsymbol{B}_t)$$

AMP for Bayesian changepoint inference

Signal is $\boldsymbol{B} = \begin{bmatrix} \uparrow & & \uparrow \\ \beta^{(1)} & \cdots & \beta^{(L)} \\ \downarrow & & \downarrow \end{bmatrix} \overset{\text{rows i.i.d.}}{\sim} p_{\bar{\boldsymbol{B}}}$ and define $\boldsymbol{\Theta} := \boldsymbol{X}\boldsymbol{B} \overset{\text{rows i.i.d.}}{\sim} p_{\bar{\boldsymbol{\Theta}}}$.

$$\boldsymbol{\Theta}_t = \boldsymbol{X}\widehat{\boldsymbol{B}}_t - \boldsymbol{R}_{t-1}\boldsymbol{F}_t^\top$$

$$\boldsymbol{R}_t = g_t(\boldsymbol{\Theta}_t, \boldsymbol{y})$$

$$\boldsymbol{B}_{t+1} = \boldsymbol{X}^\top \boldsymbol{R}_t - \widehat{\boldsymbol{B}}_t \boldsymbol{C}_t^\top$$

$$\widehat{\boldsymbol{B}}_t = f_t(\boldsymbol{B}_t)$$

produce signal estimate $\widehat{\boldsymbol{B}}_t$

AMP for Bayesian changepoint inference

Signal is $\boldsymbol{B} = \begin{bmatrix} \uparrow & & \uparrow \\ \beta^{(1)} & \cdots & \beta^{(L)} \\ \downarrow & & \downarrow \end{bmatrix} \overset{\text{rows i.i.d.}}{\sim} p_{\bar{\boldsymbol{B}}}$ and define $\boldsymbol{\Theta} := \boldsymbol{X}\boldsymbol{B} \overset{\text{rows i.i.d.}}{\sim} p_{\bar{\boldsymbol{\Theta}}}$.

$$\boldsymbol{\Theta}_t = \boldsymbol{X}\widehat{\boldsymbol{B}}_t - \boldsymbol{R}_{t-1}\boldsymbol{F}_t^\top$$

$$\boldsymbol{R}_t = g_t(\boldsymbol{\Theta}_t, \boldsymbol{y})$$

produce residual $\boldsymbol{R}_t$,
infer changepoints

$$\boldsymbol{B}_{t+1} = \boldsymbol{X}^\top \boldsymbol{R}_t - \widehat{\boldsymbol{B}}_t \boldsymbol{C}_t^\top$$

$$\widehat{\boldsymbol{B}}_t = f_t(\boldsymbol{B}_t)$$

produce signal estimate $\widehat{\boldsymbol{B}}_t$

AMP for Bayesian changepoint inference

Signal is $\boldsymbol{B} = \begin{bmatrix} \uparrow & & \uparrow \\ \beta^{(1)} & \cdots & \beta^{(L)} \\ \downarrow & & \downarrow \end{bmatrix} \overset{\text{rows i.i.d.}}{\sim} p_{\bar{B}}$ and define $\boldsymbol{\Theta} := \boldsymbol{X}\boldsymbol{B} \overset{\text{rows i.i.d.}}{\sim} p_{\bar{\Theta}}$.

$$\boldsymbol{\Theta}_t = \boldsymbol{X}\widehat{\boldsymbol{B}}_t - \boldsymbol{R}_{t-1}\boldsymbol{F}_t^\top$$

$$\boldsymbol{R}_t = g_t(\boldsymbol{\Theta}_t, \boldsymbol{y})$$

produce residual $\boldsymbol{R}_t$,
infer changepoints

$$\boldsymbol{B}_{t+1} = \boldsymbol{X}^\top \boldsymbol{R}_t - \widehat{\boldsymbol{B}}_t \boldsymbol{C}_t^\top$$

$$\widehat{\boldsymbol{B}}_t = f_t(\boldsymbol{B}_t)$$

produce signal estimate $\widehat{\boldsymbol{B}}_t$

Theorem (informal) [ALGV25].

rows of $\boldsymbol{B}_t - \boldsymbol{B}\nu_{B,t} \approx \boldsymbol{G}_{B,t} \qquad \boldsymbol{G}_{B,t} \sim \mathcal{N}(\boldsymbol{0}, \kappa_{B,t})$

rows of $\boldsymbol{\Theta}_t - \boldsymbol{\Theta}\nu_{\Theta,t} \approx \boldsymbol{G}_{\Theta,t} \qquad \boldsymbol{G}_{\Theta,t} \sim \mathcal{N}(\boldsymbol{0}, \kappa_{\Theta,t})$

AMP for Bayesian changepoint inference

Signal is $\boldsymbol{B} = \begin{bmatrix} \uparrow & & \uparrow \\ \beta^{(1)} & \cdots & \beta^{(L)} \\ \downarrow & & \downarrow \end{bmatrix} \overset{\text{rows}}{\underset{}{\overset{\text{i.i.d.}}{\sim}}} p_{\bar{\boldsymbol{B}}}$ and define $\boldsymbol{\Theta} := \boldsymbol{X}\boldsymbol{B} \overset{\text{rows}}{\overset{\text{i.i.d.}}{\sim}} p_{\bar{\boldsymbol{\Theta}}}$.

$$\boldsymbol{\Theta}_t = \boldsymbol{X}\widehat{\boldsymbol{B}}_t - \boldsymbol{R}_{t-1}\boldsymbol{F}_t^\top$$

$$\boldsymbol{R}_t = g_t(\boldsymbol{\Theta}_t, \boldsymbol{y})$$

produce residual $\boldsymbol{R}_t$,
infer changepoints

$$\boldsymbol{B}_{t+1} = \boldsymbol{X}^\top \boldsymbol{R}_t - \widehat{\boldsymbol{B}}_t \boldsymbol{C}_t^\top$$

$$\widehat{\boldsymbol{B}}_t = f_t(\boldsymbol{B}_t) = \mathbb{E}\left[\bar{\boldsymbol{B}} \mid \bar{\boldsymbol{B}}\nu_{\boldsymbol{B},t} + \boldsymbol{G}_{\boldsymbol{B},t} = \boldsymbol{B}_t\right]$$

Bayes-optimal

produce signal estimate $\widehat{\boldsymbol{B}}_t$

Theorem (informal) [ALGV25].

rows of $\boldsymbol{B}_t - \boldsymbol{B}\nu_{\boldsymbol{B},t} \approx \boldsymbol{G}_{\boldsymbol{B},t}$ $\quad \boldsymbol{G}_{\boldsymbol{B},t} \sim \mathcal{N}(\boldsymbol{0}, \kappa_{\boldsymbol{B},t})$

rows of $\boldsymbol{\Theta}_t - \boldsymbol{\Theta}\nu_{\boldsymbol{\Theta},t} \approx \boldsymbol{G}_{\boldsymbol{\Theta},t}$ $\quad \boldsymbol{G}_{\boldsymbol{\Theta},t} \sim \mathcal{N}(\boldsymbol{0}, \kappa_{\boldsymbol{\Theta},t})$

AMP for Bayesian changepoint inference

Signal is $\boldsymbol{B} = \begin{bmatrix} \uparrow & & \uparrow \\ \beta^{(1)} & \cdots & \beta^{(L)} \\ \downarrow & & \downarrow \end{bmatrix} \overset{\text{rows}}{\underset{\text{i.i.d.}}{\sim}} p_{\bar{\boldsymbol{B}}}$ and define $\boldsymbol{\Theta} := \boldsymbol{X}\boldsymbol{B} \overset{\text{rows}}{\underset{\text{i.i.d.}}{\sim}} p_{\bar{\boldsymbol{\Theta}}}$.

g_t uses prior info about change-points $\eta \sim p_{\bar{\eta}}$

$$\boldsymbol{\Theta}_t = \boldsymbol{X}\widehat{\boldsymbol{B}}_t - \boldsymbol{R}_{t-1}\boldsymbol{F}_t^\top$$

$$\boldsymbol{R}_t = g_t(\boldsymbol{\Theta}_t, \boldsymbol{y})$$

produce residual $\boldsymbol{R}_t$, infer changepoints

$$\boldsymbol{B}_{t+1} = \boldsymbol{X}^\top \boldsymbol{R}_t - \widehat{\boldsymbol{B}}_t \boldsymbol{C}_t^\top$$

$$\widehat{\boldsymbol{B}}_t = f_t(\boldsymbol{B}_t) = \mathbb{E}[\bar{\boldsymbol{B}} \mid \bar{\boldsymbol{B}}\nu_{\boldsymbol{B},t} + \boldsymbol{G}_{\boldsymbol{B},t} = \boldsymbol{B}_t]$$

Bayes-optimal

produce signal estimate $\widehat{\boldsymbol{B}}_t$

Theorem (informal) [ALGV25].

rows of $\boldsymbol{B}_t - \boldsymbol{B}\nu_{\boldsymbol{B},t} \approx \boldsymbol{G}_{\boldsymbol{B},t}$ $\quad \boldsymbol{G}_{\boldsymbol{B},t} \sim \mathcal{N}(\boldsymbol{0}, \kappa_{\boldsymbol{B},t})$

rows of $\boldsymbol{\Theta}_t - \boldsymbol{\Theta}\nu_{\boldsymbol{\Theta},t} \approx \boldsymbol{G}_{\boldsymbol{\Theta},t}$ $\quad \boldsymbol{G}_{\boldsymbol{\Theta},t} \sim \mathcal{N}(\boldsymbol{0}, \kappa_{\boldsymbol{\Theta},t})$

AMP for Bayesian changepoint inference

Signal is $\boldsymbol{B} = \begin{bmatrix} \uparrow & & \uparrow \\ \beta^{(1)} & \cdots & \beta^{(L)} \\ \downarrow & & \downarrow \end{bmatrix} \overset{\text{rows i.i.d.}}{\sim} p_{\bar{\boldsymbol{B}}}$ and define $\boldsymbol{\Theta} := \boldsymbol{X}\boldsymbol{B} \overset{\text{rows i.i.d.}}{\sim} p_{\bar{\boldsymbol{\Theta}}}$.

g_t uses prior info about change-points $\eta \sim p_{\bar{\eta}}$

$$\boldsymbol{\Theta}_t = \boldsymbol{X}\widehat{\boldsymbol{B}}_t - \boldsymbol{R}_{t-1}\boldsymbol{F}_t^\top$$

$$\boldsymbol{R}_t = g_t(\boldsymbol{\Theta}_t, \boldsymbol{y})$$

produce residual $\boldsymbol{R}_t$, infer changepoints

$$\boldsymbol{B}_{t+1} = \boldsymbol{X}^\top \boldsymbol{R}_t - \widehat{\boldsymbol{B}}_t \boldsymbol{C}_t^\top$$

$$\widehat{\boldsymbol{B}}_t = f_t(\boldsymbol{B}_t) = \mathbb{E}\left[\bar{\boldsymbol{B}} \mid \bar{\boldsymbol{B}}\nu_{\boldsymbol{B},t} + \boldsymbol{G}_{\boldsymbol{B},t} = \boldsymbol{B}_t\right]$$

Bayes-optimal

produce signal estimate $\widehat{\boldsymbol{B}}_t$

Proposition (informal) [ALGV25].

$$p\left(\eta \mid \boldsymbol{\Theta}_t, \boldsymbol{y}\right) \overset{P}{\simeq} p\left(\eta \mid \bar{\boldsymbol{\Theta}}\nu_{\boldsymbol{\Theta},t} + \boldsymbol{G}_{\boldsymbol{\Theta},t}, \bar{y}(\bar{\boldsymbol{\Theta}}, \bar{\eta}, \bar{\varepsilon})\right)$$

Simulations

Simulations

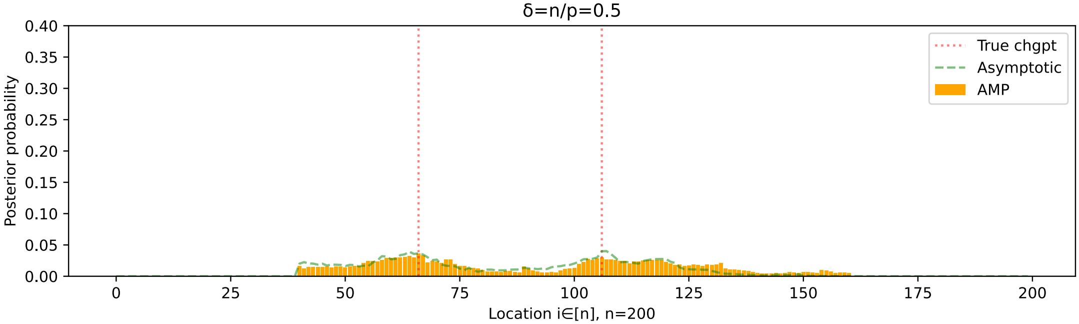

Simulations

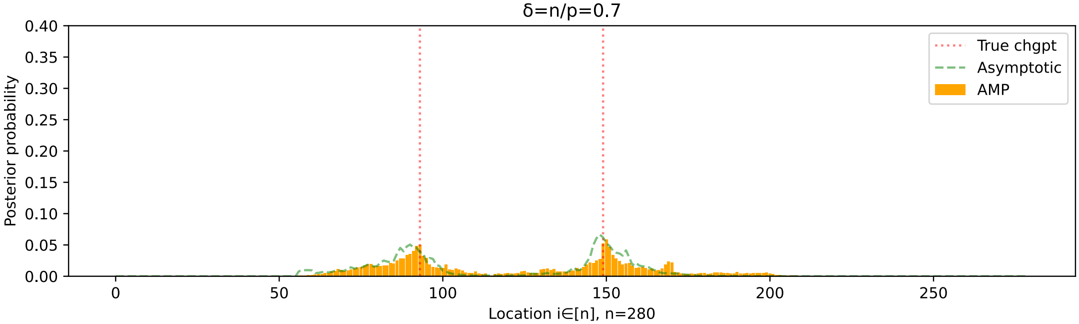

Simulations

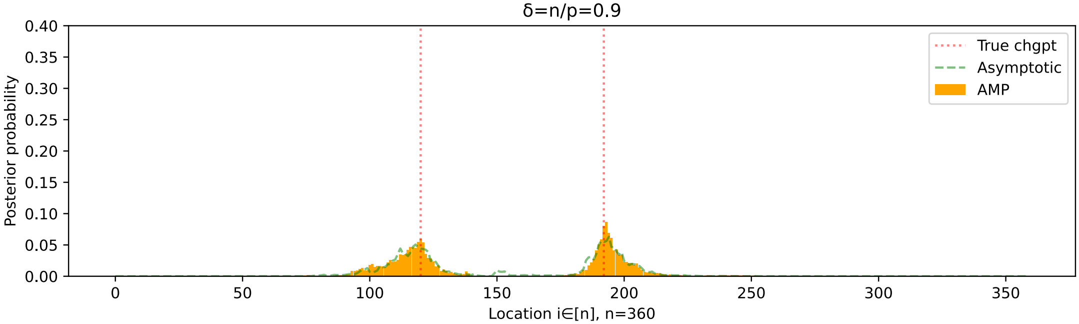

Simulations

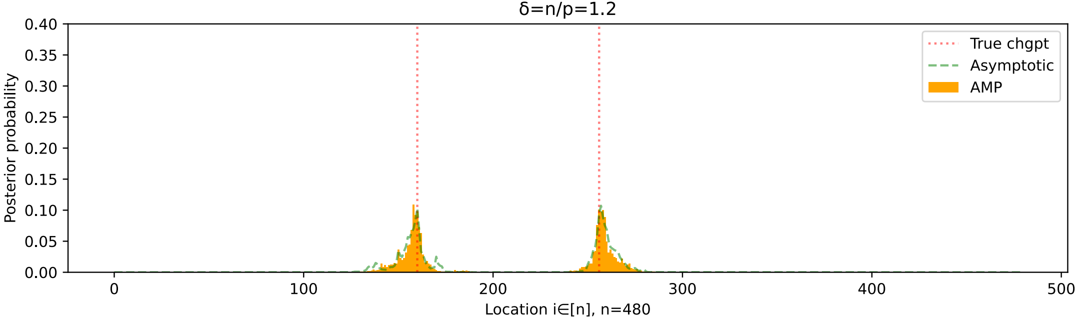

Simulations

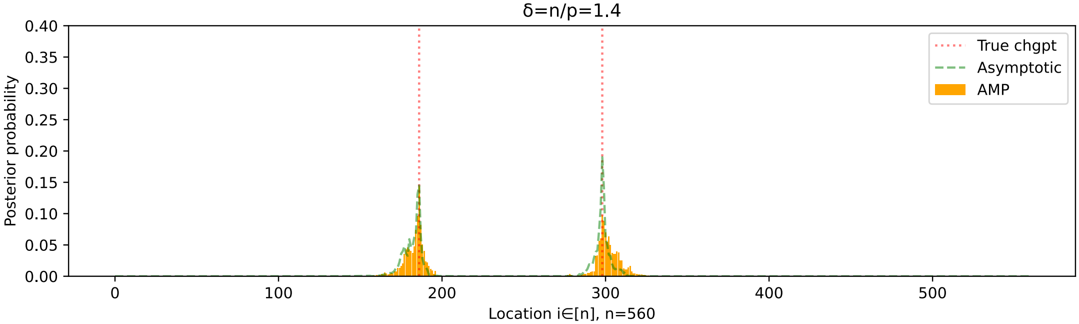

Simulations

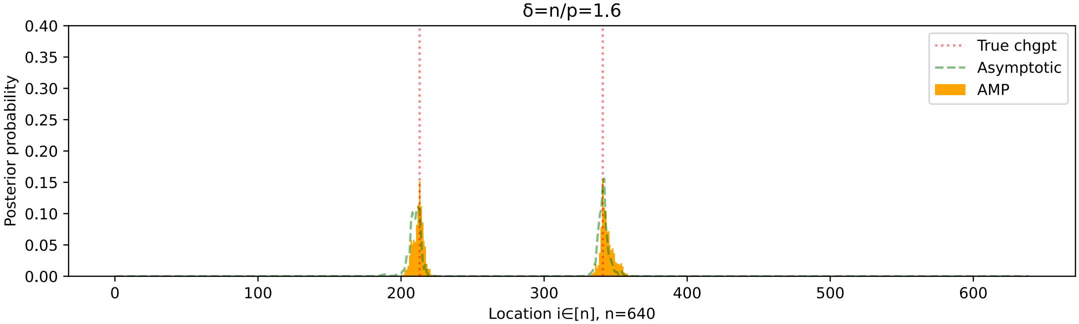

Simulations

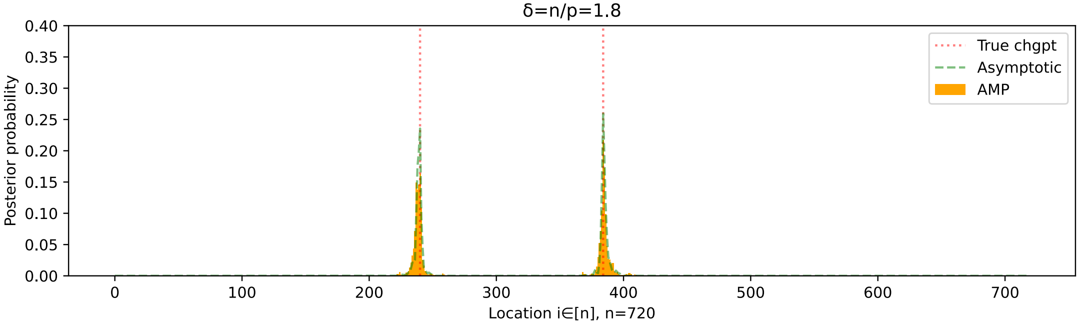

Simulations

Simulations

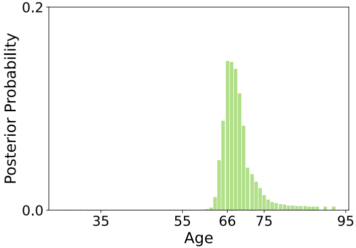

(a) Changepoint posterior vs. age.

Simulations

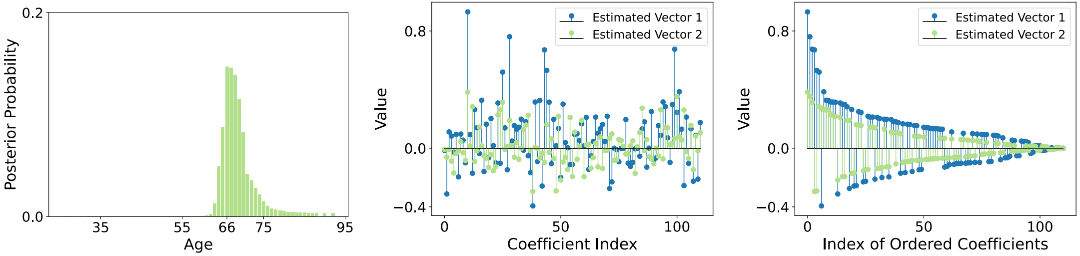

(a) Changepoint posterior vs. age. (b) Regression vectors. (c) Ordered regression vectors.

Figure: AMP applied to a Myocardial Infarction dataset from [GSR+20].

Simulations

Simulations

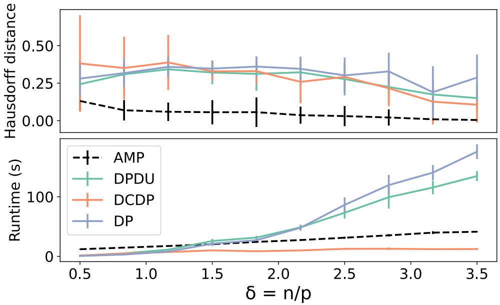

(a) Comparison with DPDU, DCDP and DP for **sparse** prior $p_{\bar{\boldsymbol{B}}} = 0.5\mathcal{N}(\boldsymbol{0}, \delta \boldsymbol{I}) + 0.5\delta_{\boldsymbol{0}}$.

Simulations

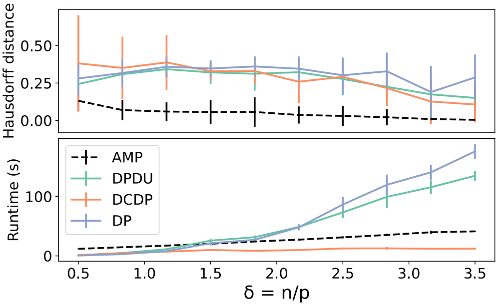

(a) Comparison with DPDU, DCDP and DP for **sparse** prior $p_{\bar{B}} = 0.5\mathcal{N}(\mathbf{0}, \delta \boldsymbol{I}) + 0.5\delta_{\mathbf{0}}$.

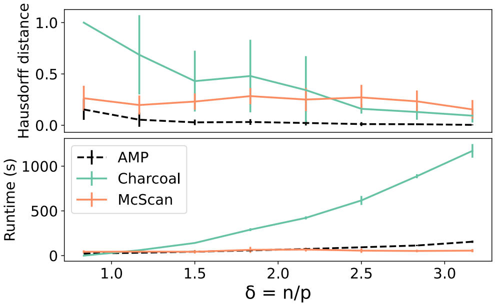

(b) Comparison with charcoal and McScan for a **sparse-difference** prior with sparsity level 0.5.

Figure: Linear model, $L^* = L = 3$.

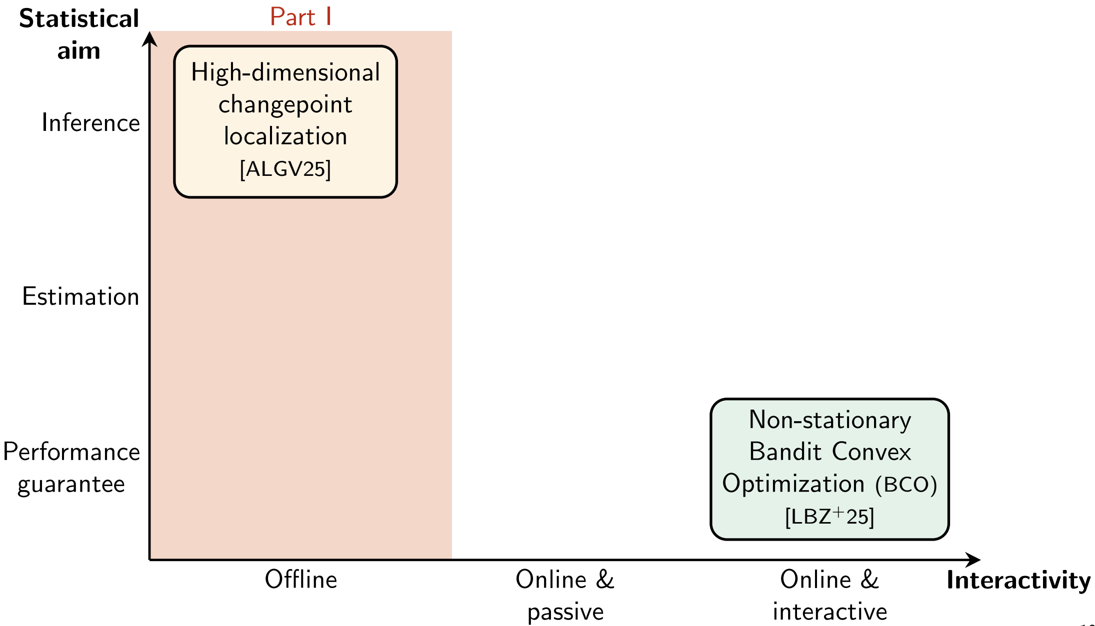
Statistical aim
Part I
High-dimensional changepoint localization [ALGV25]
Inference
Estimation
Performance guarantee
Non-stationary Bandit Convex Optimization (BCO) [LBZ+25]
Offline
Online & passive
Online & interactive
Interactivity

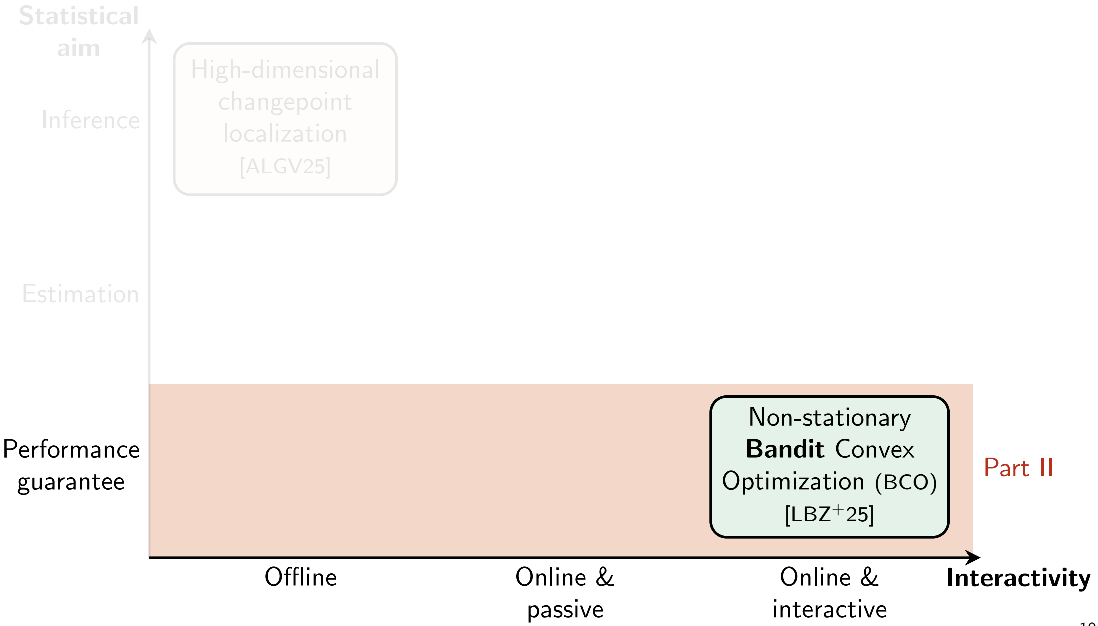
Statistical aim
Inference
Estimation
Performance guarantee
High-dimensional changepoint localization [ALGV25]
Non-stationary Bandit Convex Optimization (BCO) [LBZ+25]
Part II
Offline
Online & passive
Online & interactive
Interactivity

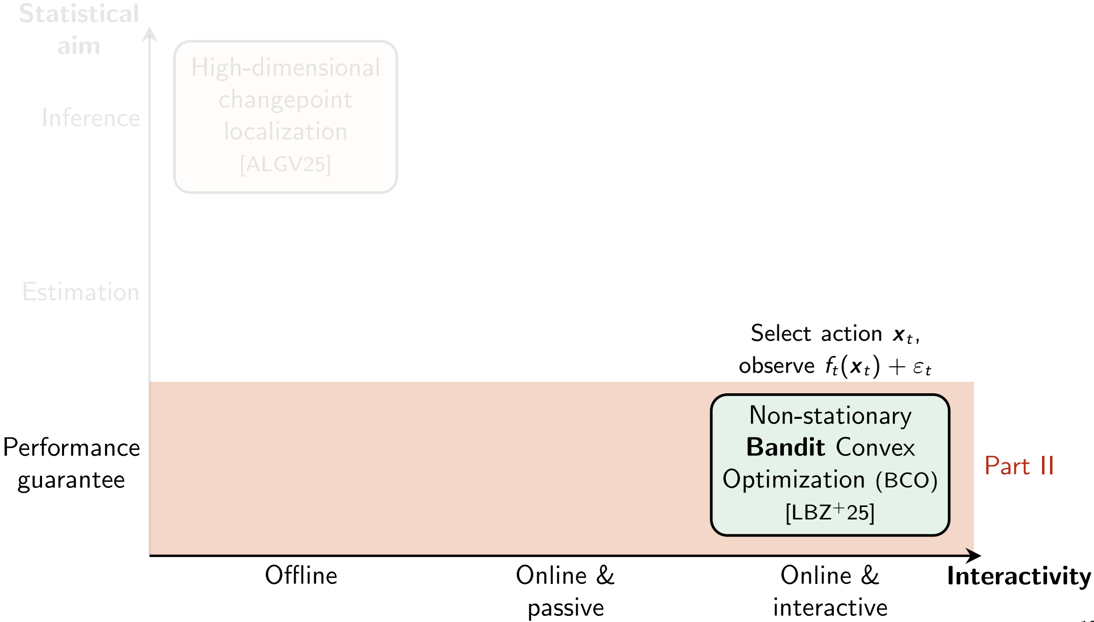
Statistical aim
Inference
Estimation
Performance guarantee
High-dimensional changepoint localization [ALGV25]
Select action $\boldsymbol{x}_t$, observe $f_t(\boldsymbol{x}_t) + \varepsilon_t$
Non-stationary **Bandit** Convex Optimization (BCO) [LBZ+25]
Part II
Offline
Online & passive
Online & interactive
Interactivity

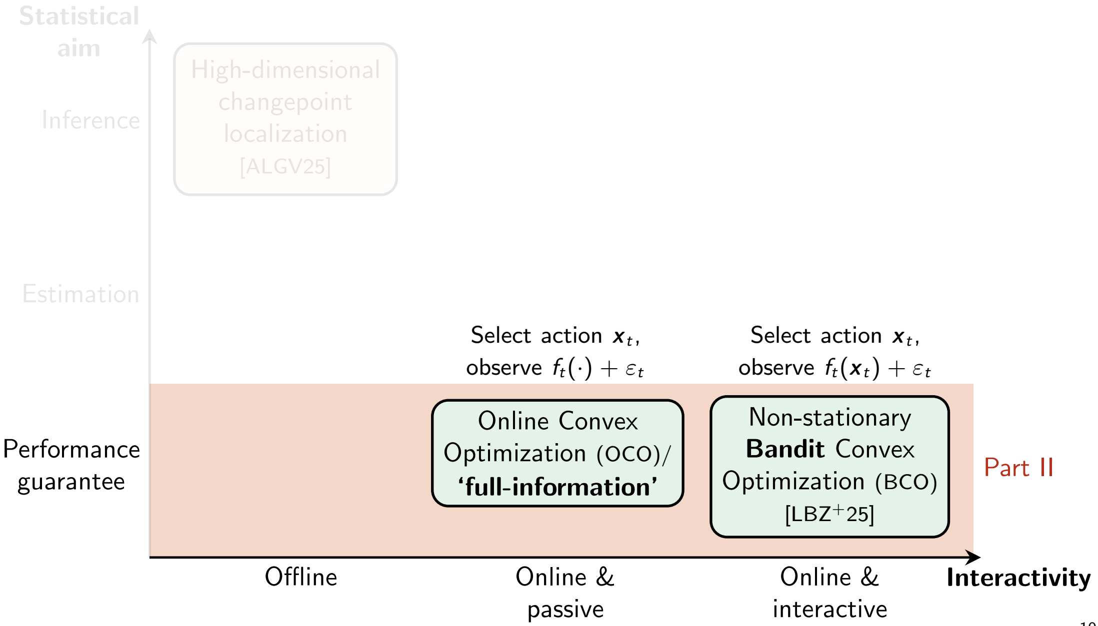
Statistical aim
High-dimensional changepoint localization [ALGV25]
Inference
Estimation
Select action $\boldsymbol{x}_t$, observe $f_t(\cdot) + \varepsilon_t$
Select action $\boldsymbol{x}_t$, observe $f_t(\boldsymbol{x}_t) + \varepsilon_t$
Performance guarantee
Online Convex Optimization (OCO)/ **'full-information'**
Non-stationary **Bandit** Convex Optimization (BCO) [LBZ+25]
Part II
Offline
Online & passive
Online & interactive
Interactivity

Bandit Convex Optimization (BCO)

Bandit Convex Optimization (BCO)

- Adversary fixes convex loss functions $f_1, f_2, \ldots, f_T : \mathbb{R}^d \to \mathbb{R}$
- For $t \geq 1$, learner
  - Selects action $\boldsymbol{x}_t$ from continuous arm set $\Theta \subseteq \mathbb{R}^d$
  - Incurs loss $f_t(\boldsymbol{x}_t)$, observes bandit feedback

$$f_t(\boldsymbol{x}_t) + \varepsilon_t$$

Bandit Convex Optimization (BCO)

- Adversary fixes convex loss functions $f_1, f_2, \ldots, f_T : \mathbb{R}^d \to \mathbb{R}$
- For $t \geq 1$, learner
  - Selects action $\boldsymbol{x}_t$ from continuous arm set $\Theta \subseteq \mathbb{R}^d$
  - Incurs loss $f_t(\boldsymbol{x}_t)$, observes bandit feedback

$$f_t(\boldsymbol{x}_t) + \varepsilon_t$$

- Regret against comparators $\boldsymbol{u}_{1:T} = \{\boldsymbol{u}_1, \ldots, \boldsymbol{u}_T\}$ chosen by adversary:

$$R(T, \boldsymbol{u}_{1:T}) := \mathbb{E}\left[\sum_{t=1}^{T} f_t(\boldsymbol{x}_t) - \sum_{t=1}^{T} f_t(\boldsymbol{u}_t)\right]$$

Bandit Convex Optimization (BCO)

- Adversary fixes convex loss functions $f_1, f_2, \ldots, f_T : \mathbb{R}^d \to \mathbb{R}$
- For $t \geq 1$, learner
  - Selects action $\boldsymbol{x}_t$ from continuous arm set $\Theta \subseteq \mathbb{R}^d$
  - Incurs loss $f_t(\boldsymbol{x}_t)$, observes bandit feedback

$$f_t(\boldsymbol{x}_t) + \varepsilon_t$$

- Regret against comparators $\boldsymbol{u}_{1:T} = \{\boldsymbol{u}_1, \ldots, \boldsymbol{u}_T\}$ chosen by adversary:

$$R(T, \boldsymbol{u}_{1:T}) := \mathbb{E}\left[\sum_{t=1}^{T} f_t(\boldsymbol{x}_t) - \sum_{t=1}^{T} f_t(\boldsymbol{u}_t)\right]$$

  - Static regret:

$$\max_{\boldsymbol{u}_{1:T}:\boldsymbol{u}_1=\cdots=\boldsymbol{u}_T} R(T, \boldsymbol{u}_{1:T}) = \mathbb{E}\left[\sum_{t=1}^{T} f_t(\boldsymbol{x}_t) - \min_{\boldsymbol{u}\in\Theta} \sum_{t=1}^{T} f_t(\boldsymbol{u})\right]$$

Bandit Convex Optimization (BCO)

- Adversary fixes convex loss functions $f_1, f_2, \ldots, f_T : \mathbb{R}^d \to \mathbb{R}$
- For $t \geq 1$, learner
  - Selects action $\boldsymbol{x}_t$ from continuous arm set $\Theta \subseteq \mathbb{R}^d$
  - Incurs loss $f_t(\boldsymbol{x}_t)$, observes bandit feedback

$$f_t(\boldsymbol{x}_t) + \varepsilon_t$$

- Regret against comparators $\boldsymbol{u}_{1:T} = \{\boldsymbol{u}_1, \ldots, \boldsymbol{u}_T\}$ chosen by adversary:

$$R(T, \boldsymbol{u}_{1:T}) := \mathbb{E}\left[\sum_{t=1}^{T} f_t(\boldsymbol{x}_t) - \sum_{t=1}^{T} f_t(\boldsymbol{u}_t)\right]$$

  - Static regret:

$$\max_{\boldsymbol{u}_{1:T}:\boldsymbol{u}_1=\cdots=\boldsymbol{u}_T} R(T, \boldsymbol{u}_{1:T}) = \mathbb{E}\left[\sum_{t=1}^{T} f_t(\boldsymbol{x}_t) - \min_{\boldsymbol{u}\in\Theta} \sum_{t=1}^{T} f_t(\boldsymbol{u})\right]$$

  - Non-stationary regrets?

Non-stationary regrets

Non-stationary regrets

Focus on switching regret:

$$R^{\text{swi}}(T, S) = \max_{\boldsymbol{u}_{1:T} : S(\boldsymbol{u}_{1:T}) \le S} R(T, \boldsymbol{u}_{1:T}), \quad \text{where} \quad S(\boldsymbol{u}_{1:T}) := 1 + \sum_{t=2}^{T} \mathbf{1}\{\boldsymbol{u}_t \neq \boldsymbol{u}_{t-1}\}$$

Non-stationary regrets

Focus on switching regret:

$$R^{\mathrm{swi}}(T,S) = \max_{\boldsymbol{u}_{1:T}: S(\boldsymbol{u}_{1:T}) \le S} R(T, \boldsymbol{u}_{1:T}), \quad \text{where} \quad S(\boldsymbol{u}_{1:T}) := 1 + \sum_{t=2}^{T} \mathbf{1}\{\boldsymbol{u}_t \neq \boldsymbol{u}_{t-1}\}$$

Example:

Non-stationary regrets

Focus on switching regret:

$$R^{\text{swi}}(T,S) = \max_{\boldsymbol{u}_{1:T} : S(\boldsymbol{u}_{1:T}) \le S} R(T, \boldsymbol{u}_{1:T}), \quad \text{where} \quad S(\boldsymbol{u}_{1:T}) := 1 + \sum_{t=2}^{T} \mathbf{1}\{\boldsymbol{u}_t \neq \boldsymbol{u}_{t-1}\}$$

Example:

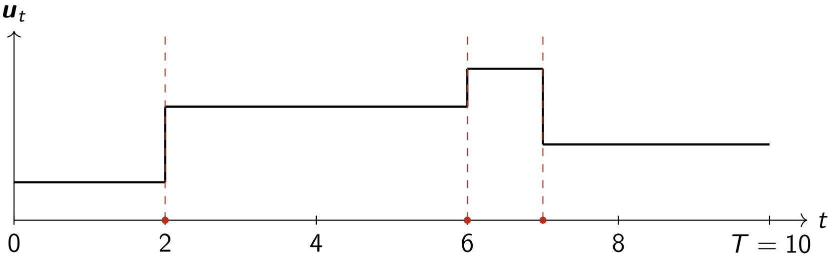

Non-stationary regrets

Focus on switching regret:

$$R^{\text{swi}}(T,S) = \max_{\boldsymbol{u}_{1:T}:S(\boldsymbol{u}_{1:T})\leq S} R(T,\boldsymbol{u}_{1:T}), \quad \text{where} \quad S(\boldsymbol{u}_{1:T}) := 1 + \sum_{t=2}^{T} \mathbf{1}\{\boldsymbol{u}_t \neq \boldsymbol{u}_{t-1}\}$$

Example:

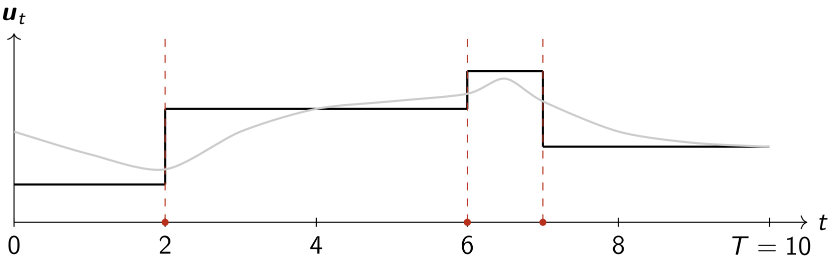

Non-stationary regrets

Focus on switching regret:

$$R^{\text{swi}}(T,S) = \max_{\boldsymbol{u}_{1:T} : S(\boldsymbol{u}_{1:T}) \leq S} R(T, \boldsymbol{u}_{1:T}), \quad \text{where} \quad S(\boldsymbol{u}_{1:T}) := 1 + \sum_{t=2}^{T} \mathbf{1}\{\boldsymbol{u}_t \neq \boldsymbol{u}_{t-1}\}$$

Example:

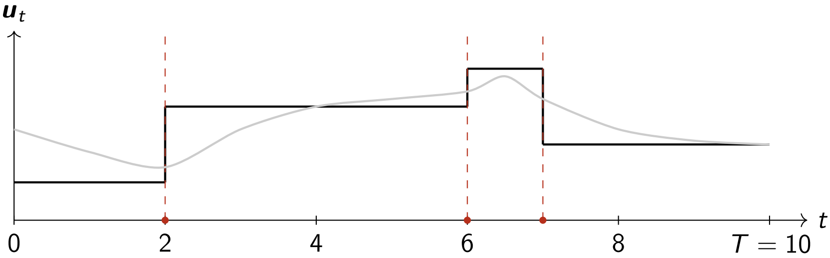

Reduction to adaptive regret:

$$R^{\text{ada}}(\mathrm{B}, T) = \max_{\substack{p,q \in [T] \\ 0 < q-p \leq \mathrm{B}}} \max_{\boldsymbol{u} \in \Theta} \sum_{t=p}^{q} \mathbb{E}[f_t(\boldsymbol{x}_t) - f_t(\boldsymbol{u})] .$$

[HS09, DGSS15, AKCV16, JOWW17, Cut20, WZZ18, ZXZZ22, ZWT^{+}21, ZLZ18, YWZZ24]

Non-stationary regrets

Focus on switching regret:

$$R^{\text{swi}}(T,S) = \max_{\boldsymbol{u}_{1:T}:S(\boldsymbol{u}_{1:T})\le S} R(T,\boldsymbol{u}_{1:T}), \quad \text{where} \quad S(\boldsymbol{u}_{1:T}) := 1 + \sum_{t=2}^{T} \mathbf{1}\{\boldsymbol{u}_t \neq \boldsymbol{u}_{t-1}\}$$

Example:

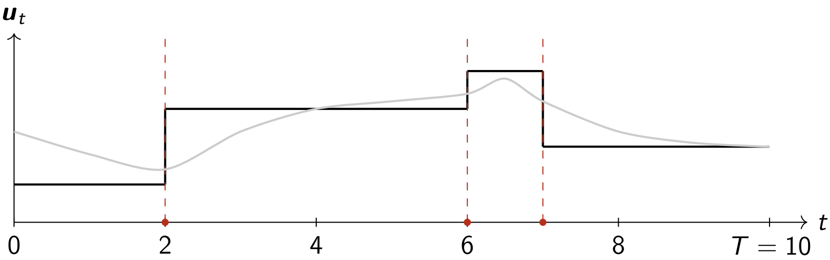

Reduction to adaptive regret:

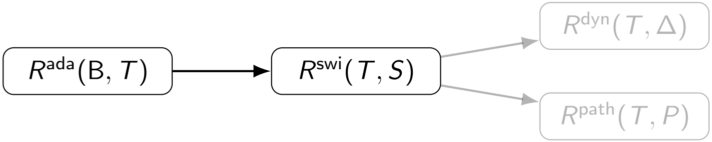

Main results

Main results

Main results

- Unified treatment for non-stationary **Bandit** Convex Optimization (BCO); previous work handled only isolated aspects [BGZ15, Wan25, ZWZZ21, CG18]

Main results

- Unified treatment for non-stationary **Bandit** Convex Optimization (BCO); previous work handled only isolated aspects [BGZ15, Wan25, ZWZZ21, CG18]
- Build on Online Convex Optimization (OCO)/**'full-information'** setting, i.e., observe $f_t(\cdot) + \varepsilon_t$

Main results

- Unified treatment for non-stationary **Bandit** Convex Optimization (BCO); previous work handled only isolated aspects [BGZ15, Wan25, ZWZZ21, CG18]
- Build on Online Convex Optimization (OCO)/**'full-information'** setting, i.e., observe $f_t(\cdot) + \varepsilon_t$
  - Expert tracking methods: [HW98, LW94, CBFH$^+$97, Vov98, FSSW97]

Main results

- Unified treatment for non-stationary **Bandit** Convex Optimization (BCO); previous work handled only isolated aspects [BGZ15, Wan25, ZWZZ21, CG18]
- Build on Online Convex Optimization (OCO)/**'full-information'** setting, i.e., observe $f_t(\cdot) + \varepsilon_t$
  - Expert tracking methods: [HW98, LW94, CBFH+97, Vov98, FSSW97]
  - Study of adaptive regret: [HS09, DGSS15, WZZ18, ZXZZ22, ZWT+21, ZLZ18, YWZZ24]

Main results

- Unified treatment for non-stationary **Bandit** Convex Optimization (BCO); previous work handled only isolated aspects [BGZ15, Wan25, ZWZZ21, CG18]
- Build on Online Convex Optimization (OCO)/**'full-information'** setting, i.e., observe $f_t(\cdot) + \varepsilon_t$
  - Expert tracking methods: [HW98, LW94, CBFH+97, Vov98, FSSW97]
  - Study of adaptive regret: [HS09, DGSS15, WZZ18, ZXZZ22, ZWT+21, ZLZ18, YWZZ24]

Algorithm	**clipped Exploration by Optimization (cExO)**	**Tilted-EW Average with Sleeping Experts (TEWA-SE)**
General convex (GC)	$d^{\frac{5}{2}}\underline{\sqrt{ST}}$	$\sqrt{d}\, S^{\frac{1}{4}}\, T^{\frac{3}{4}}$
Strongly convex (SC)		$\underline{d\sqrt{ST}}$

Underline: minimax-optimal rates [LBZ+25].

Main results

- Unified treatment for non-stationary **Bandit** Convex Optimization (BCO); previous work handled only isolated aspects [BGZ15, Wan25, ZWZZ21, CG18]
- Build on Online Convex Optimization (OCO)/**'full-information'** setting, i.e., observe $f_t(\cdot) + \varepsilon_t$
  - Expert tracking methods: [HW98, LW94, CBFH+97, Vov98, FSSW97]
  - Study of adaptive regret: [HS09, DGSS15, WZZ18, ZXZZ22, ZWT+21, ZLZ18, YWZZ24]

Algorithm	**clipped Exploration by Optimization (cExO)**	**Tilted-EW Average with Sleeping Experts (TEWA-SE)**
General convex (GC)	$d^{\frac{5}{2}}\underline{\sqrt{ST}}$	$\sqrt{d}\, S^{\frac{1}{4}} T^{\frac{3}{4}}$
Strongly convex (SC)		$\underline{d\sqrt{ST}}$
Features	• Conceptual design • Minimax-optimal in S, T • Not poly-time computable	

Underline: minimax-optimal rates [LBZ+25].

Main results

- Unified treatment for non-stationary **Bandit** Convex Optimization (BCO); previous work handled only isolated aspects [BGZ15, Wan25, ZWZZ21, CG18]
- Build on Online Convex Optimization (OCO)/**'full-information'** setting, i.e., observe $f_t(\cdot) + \varepsilon_t$
  - Expert tracking methods: [HW98, LW94, CBFH+97, Vov98, FSSW97]
  - Study of adaptive regret: [HS09, DGSS15, WZZ18, ZXZZ22, ZWT+21, ZLZ18, YWZZ24]

Algorithm	**clipped Exploration by Optimization (cExO)**	**Tilted-EW Average with Sleeping Experts (TEWA-SE)**
General convex (GC)	$\underline{d^{\frac{5}{2}}\sqrt{ST}}$	$\sqrt{d}\, S^{\frac{1}{4}} T^{\frac{3}{4}}$
Strongly convex (SC)	$\underline{d^{\frac{5}{2}}\sqrt{ST}}$	$\underline{d\sqrt{ST}}$
Features	• Conceptual design • Minimax-optimal in S, T • Not poly-time computable	• Instantiates design principles • Adaptive to curvature α • Poly-time

Underline: minimax-optimal rates [LBZ+25].

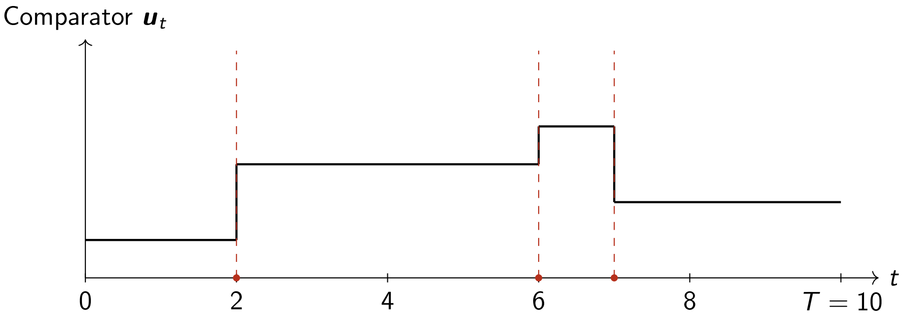
Comparator $\boldsymbol{u}_t$
t
0
2
4
6
8
$T = 10$

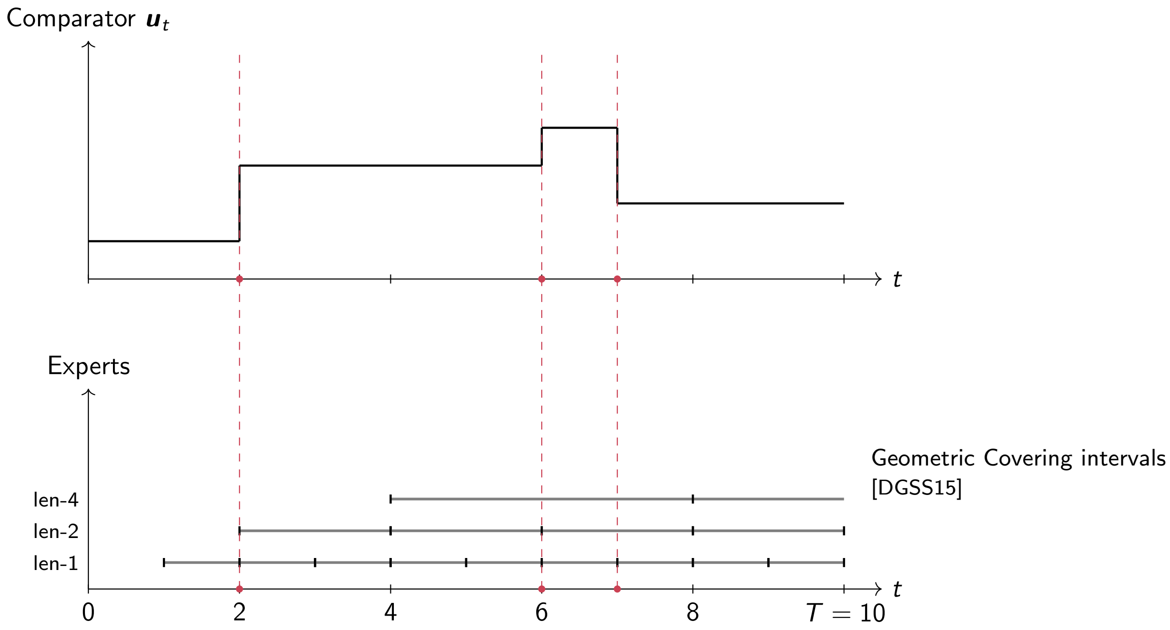
Comparator $\boldsymbol{u}_t$
t
Experts
Geometric Covering intervals
[DGSS15]
len-4
len-2
len-1
t
0
2
4
6
8
$T = 10$

Sleeping experts (full-information)

Sleeping experts (full-information)

for $t = 1, 2, \ldots, T$ **do**

Let $\{E_1, \ldots, E_{n_t}\}$ be the active experts at t with $E_i \equiv E_i(I_i)$.

for $i = 1, 2, \ldots, n_t$ **do**

Expert E_i outputs action $\boldsymbol{x}_{t,I_i}$

end for

Play meta-action using exponential-weights (EW):

$$\boldsymbol{x}_t = \sum_{i=1}^{n_t} \frac{e^{-L_{t-1,I_i}}}{\sum_{j=1}^{n_t} e^{-L_{t-1,I_j}}} \boldsymbol{x}_{t,I_i}$$

Observe entire loss function $f_t(\cdot)$

for $i = 1, 2, \ldots, n_t$ **do**

Send $f_t(\cdot)$ to E_i

Increment cumulative loss $L_{t,I_i} = L_{t-1,I_i} + f_t(\boldsymbol{x}_{t,I_i})$

end for

end for

[DGSS15, JOWW17, ZYZ+18]

Sleeping experts (full-information)

for $t = 1, 2, \ldots, T$ **do**

Let $\{E_1, \ldots, E_{n_t}\}$ be the active experts at t with $E_i \equiv E_i(I_i)$.

for $i = 1, 2, \ldots, n_t$ **do**

Expert E_i outputs action $\boldsymbol{x}_{t,I_i}$

end for

Play meta-action using exponential-weights (EW):

$$\boldsymbol{x}_t = \sum_{i=1}^{n_t} \frac{e^{-L_{t-1,I_i}}}{\sum_{j=1}^{n_t} e^{-L_{t-1,I_j}}} \boldsymbol{x}_{t,I_i}$$

Observe entire loss function $f_t(\cdot)$

for $i = 1, 2, \ldots, n_t$ **do**

Send $f_t(\cdot)$ to E_i

Increment cumulative loss $L_{t,I_i} = L_{t-1,I_i} + f_t(\boldsymbol{x}_{t,I_i})$

end for

end for

For any covering interval I,

$$\sum_{t \in I} \left(f_t(\boldsymbol{x}_t) - f_t(\boldsymbol{u})\right) = \underbrace{\sum_{t \in I} \left(f_t(\boldsymbol{x}_t) - f_t(\boldsymbol{x}_{t,I})\right)}_{\text{meta-regret}} + \underbrace{\sum_{t \in I} \left(f_t(\boldsymbol{x}_{t,I}) - f_t(\boldsymbol{u})\right)}_{\text{expert-regret}}$$

[DGSS15, JOWW17, ZYZ[+]18]

Sleeping experts (full-information)

for $t = 1, 2, \ldots, T$ **do**

Let $\{E_1, \ldots, E_{n_t}\}$ be the active experts at t with $E_i \equiv E_i(I_i)$.

for $i = 1, 2, \ldots, n_t$ **do**

Expert E_i outputs action $\boldsymbol{x}_{t,I_i}$

end for

Play meta-action using exponential-weights (EW):

$$\boldsymbol{x}_t = \sum_{i=1}^{n_t} \frac{e^{-L_{t-1,I_i}}}{\sum_{j=1}^{n_t} e^{-L_{t-1,I_j}}} \boldsymbol{x}_{t,I_i}$$

Observe entire loss function $f_t(\cdot)$

for $i = 1, 2, \ldots, n_t$ **do**

Send $f_t(\cdot)$ to E_i

Increment cumulative loss $L_{t,I_i} = L_{t-1,I_i} + f_t(\boldsymbol{x}_{t,I_i})$

end for

end for

For any covering interval I,

$$\sum_{t \in I} \big(f_t(\boldsymbol{x}_t) - f_t(\boldsymbol{u})\big) = \underbrace{\sum_{t \in I} \big(f_t(\boldsymbol{x}_t) - f_t(\boldsymbol{x}_{t,I})\big)}_{\text{meta-regret}} + \underbrace{\sum_{t \in I} \big(f_t(\boldsymbol{x}_{t,I}) - f_t(\boldsymbol{u})\big)}_{\text{expert-regret}}$$

$$\Longrightarrow \; R^{\text{ada}}(\mathrm{B}, T) \lesssim \begin{cases} \sqrt{\mathrm{B}} & \text{(GC)} \\ \frac{1}{\alpha} \log \mathrm{B} & \text{(SC)} \end{cases}.$$

[DGSS15, JOWW17, ZYZ+18]

From full-information to bandit feedback

From full-information to bandit feedback

for $t = 1, 2, \ldots, T$ **do**

Let $\{E_1, \ldots, E_{n_t}\}$ be the active experts at t with $E_i \equiv E_i(I_i)$.

for $i = 1, 2, \ldots, n_t$ **do**

Expert E_i outputs action $\boldsymbol{x}_{t,I_i}$

end for

Play meta-action using EW:

$$\boldsymbol{x}_t = \sum_{i=1}^{n_t} \frac{e^{-L_{t-1,I_i}}}{\sum_{j=1}^{n_t} e^{-L_{t-1,I_j}}} \boldsymbol{x}_{t,I_i}$$

Observe bandit feedback $f_t(\boldsymbol{x}_t) + \varepsilon_t$

for $i = 1, 2, \ldots, n_t$ **do**

Increment surrogate loss $L_{t,I_i} = L_{t-1,I_i} + \ell_t(\boldsymbol{x}_{t,I_i})$

end for

end for

From full-information to bandit feedback

for $t = 1, 2, \dots, T$ **do**

Let $\{E_1, \dots, E_{n_t}\}$ be the active experts at t with $E_i \equiv E_i(I_i)$.

for $i = 1, 2, \dots, n_t$ **do**

Expert E_i outputs action $\boldsymbol{x}_{t,I_i}$

end for

Play meta-action using EW: $\zeta_t \sim$ uniform on unit sphere

$$\boldsymbol{x}_t = \tilde{\boldsymbol{x}}_t + h\zeta_t, \quad \tilde{\boldsymbol{x}}_t = \sum_{i=1}^{n_t} \frac{e^{-L_{t-1,I_i}}}{\sum_{j=1}^{n_t} e^{-L_{t-1,I_j}}} \boldsymbol{x}_{t,I_i}$$

Observe bandit feedback $f_t(\boldsymbol{x}_t) + \varepsilon_t$

Construct gradient estimate $\boldsymbol{g}_t = (d/h)(f_t(\boldsymbol{x}_t) + \varepsilon_t)\zeta_t$

for $i = 1, 2, \dots, n_t$ **do**

Increment surrogate loss $L_{t,I_i} = L_{t-1,I_i} + \ell_t(\boldsymbol{x}_{t,I_i})$

end for

end for

From full-information to bandit feedback

for $t = 1, 2, \ldots, T$ **do**

Let $\{E_1, \ldots, E_{n_t}\}$ be the active experts at t with $E_i \equiv E_i(I_i)$.

for $i = 1, 2, \ldots, n_t$ **do**

Expert E_i outputs action $\boldsymbol{x}_{t,I_i}$

end for

Play meta-action using EW: $\zeta_t \sim$ uniform on unit sphere

$$\boldsymbol{x}_t = \tilde{\boldsymbol{x}}_t + h\zeta_t, \quad \tilde{\boldsymbol{x}}_t = \sum_{i=1}^{n_t} \frac{e^{-L_{t-1,I_i}}}{\sum_{j=1}^{n_t} e^{-L_{t-1,I_j}}} \boldsymbol{x}_{t,I_i}$$

Observe bandit feedback $f_t(\boldsymbol{x}_t) + \varepsilon_t$

Construct gradient estimate $\boldsymbol{g}_t = (d/h)(f_t(\boldsymbol{x}_t) + \varepsilon_t)\zeta_t$

for $i = 1, 2, \ldots, n_t$ **do**

Send $\boldsymbol{x}_t$ and $\boldsymbol{g}_t$ to E_i

Increment surrogate loss $L_{t,I_i} = L_{t-1,I_i} + \ell_t(\boldsymbol{x}_{t,I_i})$

end for

end for

From full-information to bandit feedback

for $t = 1, 2, \ldots, T$ **do**

Let $\{E_1, \ldots, E_{n_t}\}$ be the active experts at t with $E_i \equiv E_i(I_i)$.

for $i = 1, 2, \ldots, n_t$ **do**

Expert E_i outputs action $\boldsymbol{x}_{t,I_i}$

end for

Play meta-action using EW: $\zeta_t \sim$ uniform on unit sphere

$$\boldsymbol{x}_t = \tilde{\boldsymbol{x}}_t + h\zeta_t, \quad \tilde{\boldsymbol{x}}_t = \sum_{i=1}^{n_t} \frac{e^{-L_{t-1,I_i}}}{\sum_{j=1}^{n_t} e^{-L_{t-1,I_j}}} \boldsymbol{x}_{t,I_i}$$

Observe bandit feedback $f_t(\boldsymbol{x}_t) + \varepsilon_t$

Construct gradient estimate $\boldsymbol{g}_t = (d/h)(f_t(\boldsymbol{x}_t) + \varepsilon_t)\zeta_t$

for $i = 1, 2, \ldots, n_t$ **do**

Send $\boldsymbol{x}_t$ and $\boldsymbol{g}_t$ to E_i

Increment surrogate loss $L_{t,I_i} = L_{t-1,I_i} + \ell_t(\boldsymbol{x}_{t,I_i})$

end for

end for

- $\mathbb{E}[\boldsymbol{g}_t | \tilde{\boldsymbol{x}}_t] = \nabla \hat{f}_t(\tilde{\boldsymbol{x}}_t)$ where $\hat{f}_t(\cdot)$ is a smoothed version of $f_t(\cdot)$, also convex
- Bias $= \mathbb{E}[\hat{f}_t(\boldsymbol{u}) - f_t(\boldsymbol{u})] \lesssim h$, Variance-proxy $\|\boldsymbol{g}_t\| \leq G \asymp d/h$

[FKM05, Kle04]

From full-information to bandit feedback

for $t = 1, 2, \ldots, T$ **do**

Let $\{E_1, \ldots, E_{n_t}\}$ be the active experts at t with $E_i \equiv E_i(I_i)$.

for $i = 1, 2, \ldots, n_t$ **do**

Expert E_i outputs action $\boldsymbol{x}_{t,I_i}$

end for

Play meta-action using EW: $\zeta_t \sim$ uniform on unit sphere

$$\boldsymbol{x}_t = \tilde{\boldsymbol{x}}_t + h\zeta_t, \quad \tilde{\boldsymbol{x}}_t = \sum_{i=1}^{n_t} \frac{e^{-L_{t-1,I_i}}}{\sum_{j=1}^{n_t} e^{-L_{t-1,I_j}}} \boldsymbol{x}_{t,I_i}$$

Observe bandit feedback $f_t(\boldsymbol{x}_t) + \varepsilon_t$

Construct gradient estimate $\boldsymbol{g}_t = (d/h)(f_t(\boldsymbol{x}_t) + \varepsilon_t)\zeta_t$

for $i = 1, 2, \ldots, n_t$ **do**

Send $\boldsymbol{x}_t$ and $\boldsymbol{g}_t$ to E_i

Increment surrogate loss $L_{t,I_i} = L_{t-1,I_i} + \ell_t(\boldsymbol{x}_{t,I_i})$

end for

end for

- $\mathbb{E}[\boldsymbol{g}_t | \tilde{\boldsymbol{x}}_t] = \nabla \hat{f}_t(\tilde{\boldsymbol{x}}_t)$ where $\hat{f}_t(\cdot)$ is a smoothed version of $f_t(\cdot)$, also convex
- Bias $= \mathbb{E}[\hat{f}_t(\boldsymbol{u}) - f_t(\boldsymbol{u})] \lesssim h$, Variance-proxy $\|\boldsymbol{g}_t\| \leq G \asymp d/h$

How to design $\ell_t(\cdot)$ using $\boldsymbol{g}_t$?

[FKM05, Kle04]

Designing surrogate loss $\ell_t(\cdot)$

- Linear loss $\ell_t(\boldsymbol{x}) = -\boldsymbol{g}_t^\top(\tilde{\boldsymbol{x}}_t - \boldsymbol{x})$

Designing surrogate loss $\ell_t(\cdot)$

- Linear loss $\ell_t(\boldsymbol{x}) = -\boldsymbol{g}_t^\top(\tilde{\boldsymbol{x}}_t - \boldsymbol{x})$ ▶ $R^{\text{ada}}(\text{B}, T) \lesssim \sqrt{\text{B}}$ in OCO [WZZ18]

Designing surrogate loss $\ell_t(\cdot)$

- Linear loss $\ell_t(\boldsymbol{x}) = -\boldsymbol{g}_t^\top(\tilde{\boldsymbol{x}}_t - \boldsymbol{x})$
  - $R^{\text{ada}}(\mathrm{B}, T) \lesssim \sqrt{\mathrm{B}}$ in OCO [WZZ18]
  - expert-regret for $E_I \lesssim \sqrt{|I|} \Rightarrow R^{\text{ada}}(\mathrm{B}, T) \lesssim B$ in BCO

Designing surrogate loss $\ell_t(\cdot)$

- Linear loss $\ell_t(\mathbf{x}) = -\mathbf{g}_t^\top(\tilde{\mathbf{x}}_t - \mathbf{x})$
  - ▶ $R^{\mathrm{ada}}(\mathrm{B}, T) \lesssim \sqrt{\mathrm{B}}$ in OCO [WZZ18]
  - ▶ expert-regret for $E_I \lesssim \sqrt{|I|} \Rightarrow R^{\mathrm{ada}}(\mathrm{B}, T) \lesssim B$ in BCO
- Inspired by [CV09, vEKvdH21, WLZ20, ZWT+21], propose quadratic surrogate loss

$$\ell_t^\eta(\mathbf{x}) = -\eta \mathbf{g}_t^\top(\tilde{\mathbf{x}}_t - \mathbf{x}) + \eta^2 G^2 \|\tilde{\mathbf{x}}_t - \mathbf{x}\|^2, \quad \forall \mathbf{x} \in \mathbb{R}^d$$

Designing surrogate loss $\ell_t(\cdot)$

- Linear loss $\ell_t(\boldsymbol{x}) = -\boldsymbol{g}_t^\top(\tilde{\boldsymbol{x}}_t - \boldsymbol{x})$
  - ▶ $R^{\text{ada}}(\text{B}, T) \lesssim \sqrt{\text{B}}$ in OCO [WZZ18]
  - ▶ expert-regret for $E_I \lesssim \sqrt{|I|} \Rightarrow R^{\text{ada}}(\text{B}, T) \lesssim B$ in BCO
- Inspired by [CV09, vEKvdH21, WLZ20, ZWT+21], propose quadratic surrogate loss

$$\ell_t^\eta(\boldsymbol{x}) = -\eta\boldsymbol{g}_t^\top(\tilde{\boldsymbol{x}}_t - \boldsymbol{x}) + \eta^2 G^2\|\tilde{\boldsymbol{x}}_t - \boldsymbol{x}\|^2, \quad \forall \boldsymbol{x} \in \mathbb{R}^d$$

Tilted by learning rate η

Designing surrogate loss $\ell_t(\cdot)$

- Linear loss $\ell_t(\mathbf{x}) = -\mathbf{g}_t^\top(\tilde{\mathbf{x}}_t - \mathbf{x})$
  - ▶ $R^{\text{ada}}(\mathrm{B}, T) \lesssim \sqrt{\mathrm{B}}$ in OCO [WZZ18]
  - ▶ expert-regret for $E_I \lesssim \sqrt{|I|} \Rightarrow R^{\text{ada}}(\mathrm{B}, T) \lesssim B$ in BCO
- Inspired by [CV09, vEKvdH21, WLZ20, ZWT+21], propose quadratic surrogate loss

$$\ell_t^\eta(\mathbf{x}) = -\eta\mathbf{g}_t^\top(\tilde{\mathbf{x}}_t - \mathbf{x}) + \underbrace{\eta^2 G^2\|\tilde{\mathbf{x}}_t - \mathbf{x}\|^2}, \quad \forall \mathbf{x} \in \mathbb{R}^d$$

Tilted by learning rate η

(i) Absorbs the variance of $\mathbf{g}_t$

(ii) Together with η, adapts to unknown curvature α

Designing surrogate loss $\ell_t(\cdot)$

- Linear loss $\ell_t(\mathbf{x}) = -\mathbf{g}_t^\top(\tilde{\mathbf{x}}_t - \mathbf{x})$ ▶ $R^{\text{ada}}(\mathrm{B}, T) \lesssim \sqrt{\mathrm{B}}$ in OCO [WZZ18]
  ▶ expert-regret for $E_I \lesssim \sqrt{|I|} \Rightarrow R^{\text{ada}}(\mathrm{B}, T) \lesssim B$ in BCO
- Inspired by [CV09, vEKvdH21, WLZ20, ZWT+21], propose quadratic surrogate loss

$$\ell_t^\eta(\mathbf{x}) = -\eta \mathbf{g}_t^\top(\tilde{\mathbf{x}}_t - \mathbf{x}) + \eta^2 G^2 \|\tilde{\mathbf{x}}_t - \mathbf{x}\|^2, \quad \forall \mathbf{x} \in \mathbb{R}^d$$

Designing surrogate loss $\ell_t(\cdot)$

- Linear loss $\ell_t(\mathbf{x}) = -\mathbf{g}_t^\top(\tilde{\mathbf{x}}_t - \mathbf{x})$ ▶ $R^{\text{ada}}(\mathrm{B}, T) \lesssim \sqrt{\mathrm{B}}$ in OCO [WZZ18]
  ▶ expert-regret for $E_I \lesssim \sqrt{|I|} \Rightarrow R^{\text{ada}}(\mathrm{B}, T) \lesssim B$ in BCO
- Inspired by [CV09, vEKvdH21, WLZ20, ZWT+21], propose quadratic surrogate loss

$$\ell_t^\eta(\mathbf{x}) = -\eta \mathbf{g}_t^\top(\tilde{\mathbf{x}}_t - \mathbf{x}) + \eta^2 G^2 \|\tilde{\mathbf{x}}_t - \mathbf{x}\|^2, \quad \forall \mathbf{x} \in \mathbb{R}^d$$

$\Downarrow$ by definition

$$\underbrace{\sum_{t\in I} \langle \mathbb{E}[\mathbf{g}_t | \tilde{\mathbf{x}}_t], \tilde{\mathbf{x}}_t - \mathbf{u} \rangle}_{\text{linearized regret associated with } \hat{f}_t} \leq \frac{1}{\eta} \underbrace{\sum_{t\in I} \mathbb{E}\left[\ell_t^\eta(\tilde{\mathbf{x}}_t) - \ell_t^\eta(\mathbf{u}) | \tilde{\mathbf{x}}_t\right]}_{:=\bigstar \lesssim \log |I|} + \eta G^2 \sum_{t\in I} \|\tilde{\mathbf{x}}_t - \mathbf{u}\|^2$$

Designing surrogate loss $\ell_t(\cdot)$

- Linear loss $\ell_t(\boldsymbol{x}) = -\boldsymbol{g}_t^\top(\tilde{\boldsymbol{x}}_t - \boldsymbol{x})$ ▶ $R^{\text{ada}}(\text{B}, T) \lesssim \sqrt{\text{B}}$ in OCO [WZZ18]
  ▶ expert-regret for $E_I \lesssim \sqrt{|I|} \Rightarrow R^{\text{ada}}(\text{B}, T) \lesssim B$ in BCO
- Inspired by [CV09, vEKvdH21, WLZ20, ZWT+21], propose quadratic surrogate loss

$$\ell_t^\eta(\boldsymbol{x}) = -\eta \boldsymbol{g}_t^\top(\tilde{\boldsymbol{x}}_t - \boldsymbol{x}) + \eta^2 G^2 \|\tilde{\boldsymbol{x}}_t - \boldsymbol{x}\|^2, \quad \forall \boldsymbol{x} \in \mathbb{R}^d$$

$\Downarrow$ by definition

$$\underbrace{\sum_{t \in I} \langle \mathbb{E}[\boldsymbol{g}_t | \tilde{\boldsymbol{x}}_t], \tilde{\boldsymbol{x}}_t - \boldsymbol{u}\rangle}_{\text{linearized regret associated with } \hat{f}_t} \le \frac{1}{\eta} \underbrace{\sum_{t \in I} \mathbb{E}\left[\ell_t^\eta(\tilde{\boldsymbol{x}}_t) - \ell_t^\eta(\boldsymbol{u}) | \tilde{\boldsymbol{x}}_t\right]}_{:= \bigstar \lesssim \log |I|} + \eta G^2 \sum_{t \in I} \|\tilde{\boldsymbol{x}}_t - \boldsymbol{u}\|^2$$

$\Downarrow$ by convexity of $\hat{f}_t$

$$\underbrace{\sum_{t \in I} \mathbb{E}\left[\hat{f}_t(\tilde{\boldsymbol{x}}_t) - \hat{f}_t(\boldsymbol{u})\right]}_{\text{regret associated with } \hat{f}_t} \le \underbrace{\mathbb{E}\left[\frac{1}{\eta}\bigstar + \left(\eta G^2 - \frac{\alpha}{2}\right) \sum_{t \in I} \|\tilde{\boldsymbol{x}}_t - \boldsymbol{u}\|^2\right]}_{\substack{\text{grid over } \eta \text{ to hedge} \\ \text{against unknown curvature } \alpha}}$$

Main results: Tilted-EW Average with Sleeping Experts (TEWA-SE)

Main results: Tilted-EW Average with Sleeping Experts (TEWA-SE)

Theorem (informal) [LBZ+25]. For any $T \in \mathbb{N}^+$ and $\mathrm{B} \in [T]$, TEWA-SE with $h = \sqrt{d}\mathrm{B}^{-\frac{1}{4}}$ & geometric grid for η, satisfies

$$R^{\mathrm{ada}}(\mathrm{B}, T) \lesssim \begin{cases} \sqrt{d}\mathrm{B}^{\frac{3}{4}} & \text{(GC)} \\ \frac{d}{\alpha}\sqrt{\mathrm{B}} & \text{(SC)} \end{cases}$$

Main results: Tilted-EW Average with Sleeping Experts (TEWA-SE)

Theorem (informal) [LBZ+25]. For any $T \in \mathbb{N}^+$ and $\mathrm{B} \in [T]$, TEWA-SE with $h = \sqrt{d}\mathrm{B}^{-\frac{1}{4}}$ & geometric grid for η, satisfies

$$R^{\mathrm{ada}}(\mathrm{B}, T) \lesssim \begin{cases} \sqrt{d}\mathrm{B}^{\frac{3}{4}} & \text{(GC)} \\ \frac{d}{\alpha}\sqrt{\mathrm{B}} & \text{(SC)} \end{cases}$$

Corollary. For known S, setting $\mathrm{B} = \frac{T}{S}$ yields

$$R^{\mathrm{swi}}(T, S) \lesssim \begin{cases} \sqrt{d}S^{\frac{1}{4}}T^{\frac{3}{4}} & \text{(GC)} \\ \underline{d\sqrt{ST}} & \text{(SC)} \end{cases}$$

Main results: Tilted-EW Average with Sleeping Experts (TEWA-SE)

Theorem (informal) [LBZ+25]. For any $T \in \mathbb{N}^+$ and $\mathrm{B} \in [T]$, TEWA-SE with $h = \sqrt{d}\mathrm{B}^{-\frac{1}{4}}$ & geometric grid for η, satisfies

$$R^{\text{ada}}(\mathrm{B}, T) \lesssim \begin{cases} \sqrt{d}\mathrm{B}^{\frac{3}{4}} & \text{(GC)} \\ \frac{d}{\alpha}\sqrt{\mathrm{B}} & \text{(SC)} \end{cases}$$

Corollary. For known S, setting $\mathrm{B} = \frac{T}{S}$ yields

$$R^{\text{swi}}(T, S) \lesssim \begin{cases} \sqrt{d}S^{\frac{1}{4}}T^{\frac{3}{4}} & \text{(GC)} \\ \underline{d\sqrt{ST}} & \text{(SC)} \end{cases}$$

Open directions:

Main results: Tilted-EW Average with Sleeping Experts (TEWA-SE)

Theorem (informal) [LBZ+25]. For any $T \in \mathbb{N}^+$ and $\mathrm{B} \in [T]$, TEWA-SE with $h = \sqrt{d}\mathrm{B}^{-\frac{1}{4}}$ & geometric grid for η, satisfies

$$R^{\mathrm{ada}}(\mathrm{B}, T) \lesssim \begin{cases} \sqrt{d}\mathrm{B}^{\frac{3}{4}} & \text{(GC)} \\ \frac{d}{\alpha}\sqrt{\mathrm{B}} & \text{(SC)} \end{cases}$$

Corollary. For known S, setting $\mathrm{B} = \frac{T}{S}$ yields

$$R^{\mathrm{swi}}(T, S) \lesssim \begin{cases} \sqrt{d}S^{\frac{1}{4}}T^{\frac{3}{4}} & \text{(GC)} \\ \underline{d\sqrt{ST}} & \text{(SC)} \end{cases}$$

Open directions: Towards minimax-optimal, efficient, parameter-free algorithms.

Main results: Tilted-EW Average with Sleeping Experts (TEWA-SE)

Theorem (informal) [LBZ+25]. For any $T \in \mathbb{N}^+$ and $\mathrm{B} \in [T]$, TEWA-SE with $h = \sqrt{d}\mathrm{B}^{-\frac{1}{4}}$ & geometric grid for η, satisfies

$$R^{\mathrm{ada}}(\mathrm{B}, T) \lesssim \begin{cases} \sqrt{d}\mathrm{B}^{\frac{3}{4}} & \text{(GC)} \\ \frac{d}{\alpha}\sqrt{\mathrm{B}} & \text{(SC)} \end{cases}$$

Corollary. For known S, setting $\mathrm{B} = \frac{T}{S}$ yields

$$R^{\mathrm{swi}}(T, S) \lesssim \begin{cases} \sqrt{d}S^{\frac{1}{4}}T^{\frac{3}{4}} & \text{(GC)} \\ \underline{d\sqrt{ST}} & \text{(SC)} \end{cases}$$

Open directions: Towards minimax-optimal, efficient, parameter-free algorithms.

- Improve TEWA-SE through coin betting [JOWW17]?

Main results: Tilted-EW Average with Sleeping Experts (TEWA-SE)

Theorem (informal) [LBZ+25]. For any $T \in \mathbb{N}^+$ and $\mathrm{B} \in [T]$, TEWA-SE with $h = \sqrt{d}\mathrm{B}^{-\frac{1}{4}}$ & geometric grid for η, satisfies

$$R^{\text{ada}}(\mathrm{B}, T) \lesssim \begin{cases} \sqrt{d}\mathrm{B}^{\frac{3}{4}} & \text{(GC)} \\ \frac{d}{\alpha}\sqrt{\mathrm{B}} & \text{(SC)} \end{cases}$$

Corollary. For known S, setting $\mathrm{B} = \frac{T}{S}$ yields

$$R^{\text{swi}}(T, S) \lesssim \begin{cases} \sqrt{d}S^{\frac{1}{4}}T^{\frac{3}{4}} & \text{(GC)} \\ \underline{d\sqrt{ST}} & \text{(SC)} \end{cases}$$

Open directions: Towards minimax-optimal, efficient, parameter-free algorithms.

- Improve TEWA-SE through coin betting [JOWW17]?
- Leverage second-order information like online Newton methods from [FvdHLM24, SSNH24] that achieve state-of-the-art $\sqrt{T}$ static regrets?

Summary & ongoing work

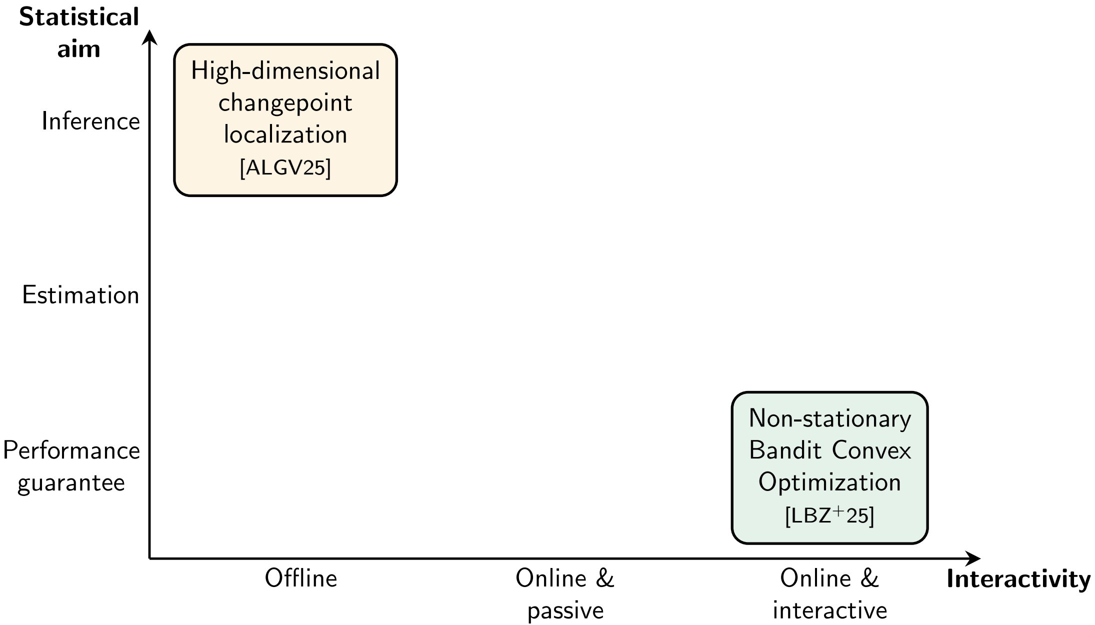

Statistical aim
Inference
Estimation
Performance guarantee
High-dimensional changepoint localization [ALGV25]
Non-stationary Bandit Convex Optimization [LBZ+25]
Offline
Online & passive
Online & interactive
Interactivity

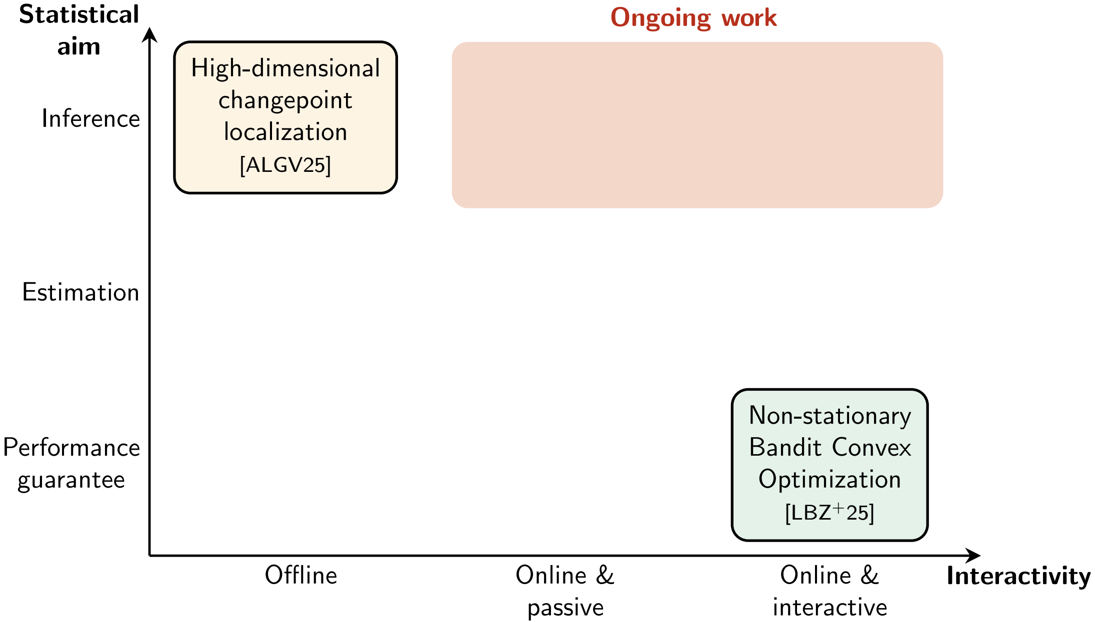
Statistical aim
Ongoing work
High-dimensional changepoint localization [ALGV25]
Inference
Estimation
Performance guarantee
Non-stationary Bandit Convex Optimization [LBZ+25]
Offline
Online & passive
Online & interactive
Interactivity

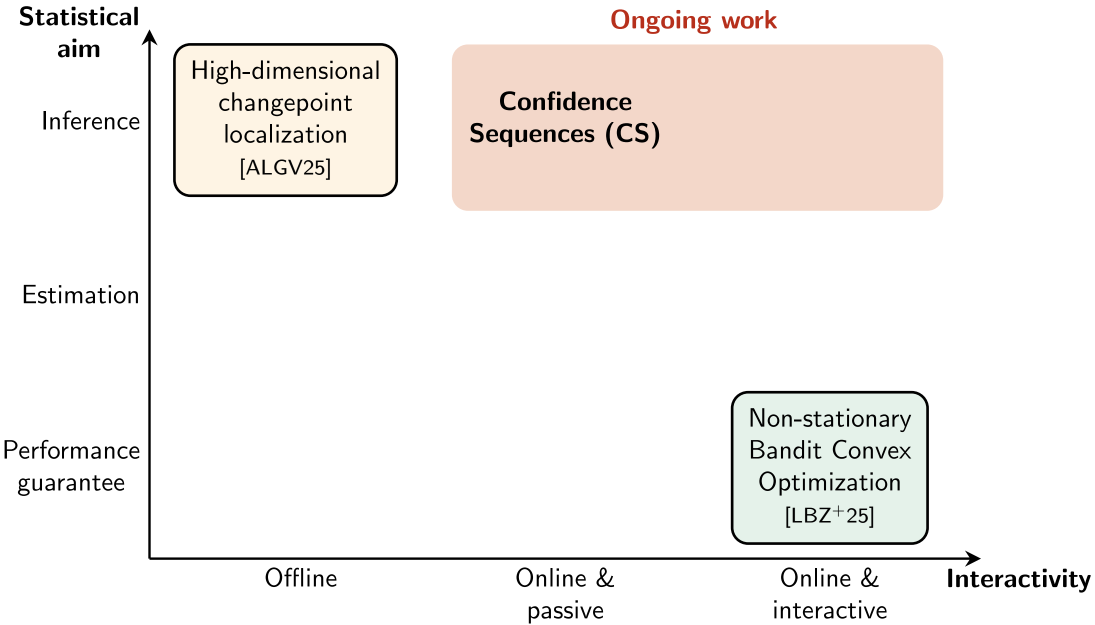
Statistical aim
Ongoing work
High-dimensional changepoint localization [ALGV25]
Confidence Sequences (CS)
Inference
Estimation
Non-stationary Bandit Convex Optimization [LBZ+25]
Performance guarantee
Offline
Online & passive
Online & interactive
Interactivity

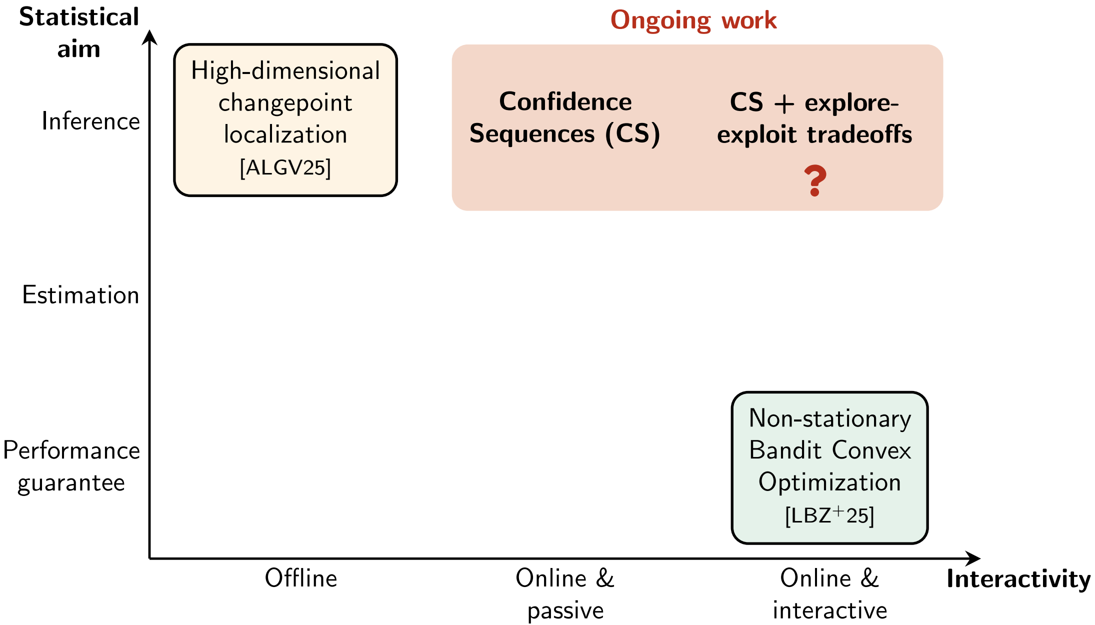
Statistical aim
Ongoing work
High-dimensional changepoint localization [ALGV25]
Confidence Sequences (CS)
CS + explore-exploit tradeoffs
?
Inference
Estimation
Performance guarantee
Non-stationary Bandit Convex Optimization [LBZ+25]
Offline
Online & passive
Online & interactive
Interactivity

Anytime-valid inference for interactive decision-making

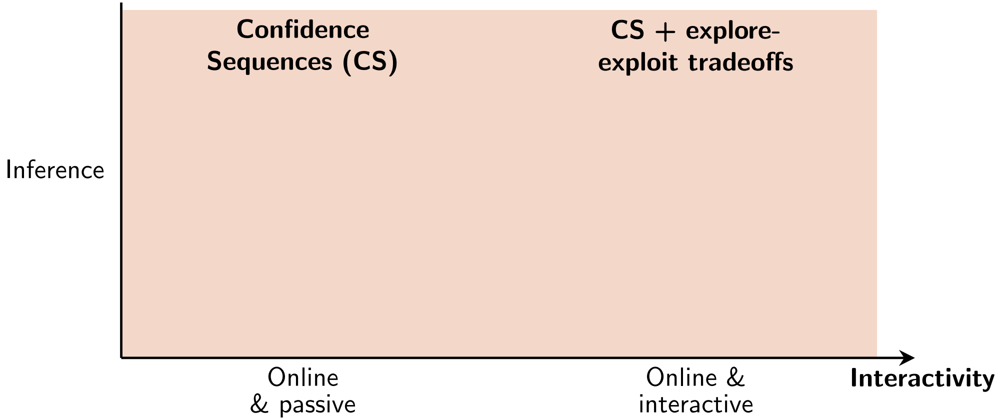

Anytime-valid inference for interactive decision-making

Particularly useful for monitoring generative AI

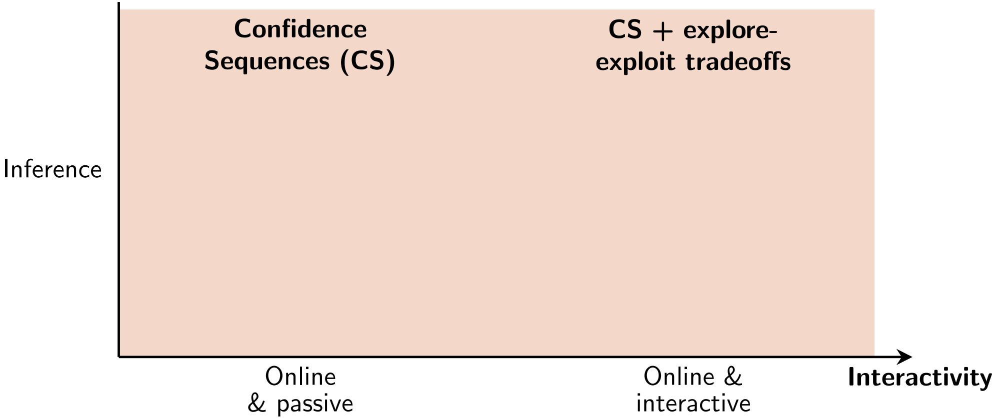

Anytime-valid **inference** for **interactive** decision-making

Particularly useful for monitoring generative AI

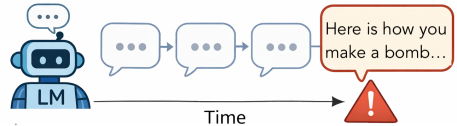

Confidence Sequences (CS)

CS + explore-exploit tradeoffs

Inference

Online & passive

Online & interactive

Interactivity

Anytime-valid inference for interactive decision-making

Particularly useful for monitoring generative AI

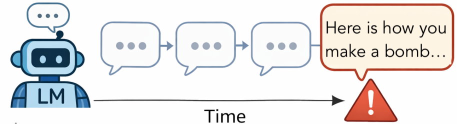

Inference

Online & passive	Online & interactive
Confidence Sequences (CS) • Time-to-unsafe sampling; correctness: [DGPC25, SPF+25, DFFR26]	**CS + explore-exploit tradeoffs**

Interactivity

Anytime-valid inference for interactive decision-making

Particularly useful for monitoring generative AI

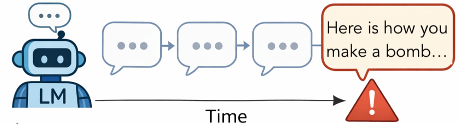

Confidence Sequences (CS)

- Time-to-unsafe sampling; correctness: [DGPC25, SPF+25, DFFR26]
- Routing & cascades

CS + explore-exploit tradeoffs

Inference

Online & passive

Online & interactive

Interactivity

Anytime-valid inference for interactive decision-making

Particularly useful for monitoring generative AI

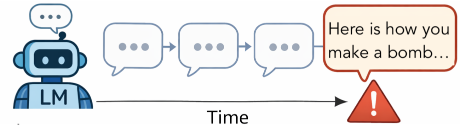

Confidence Sequences (CS)

- Time-to-unsafe sampling; correctness: [DGPC25, SPF+25, DFFR26]
- Routing & cascades
- Self-consistency: [CED25]

CS + explore-exploit tradeoffs

Inference

Online & passive

Online & interactive

Interactivity

Anytime-valid inference for interactive decision-making

Particularly useful for monitoring generative AI

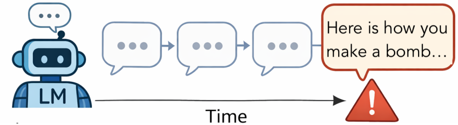

	Online & passive	Online & interactive
Inference	**Confidence Sequences (CS)** • Time-to-unsafe sampling; correctness: [DGPC25, SPF+25, DFFR26] • Routing & cascades • Self-consistency: [CED25] • Watermarking: [SWZ26, HXR+26]	**CS + explore-exploit tradeoffs**

Interactivity

Anytime-valid inference for interactive decision-making

Particularly useful for monitoring generative AI

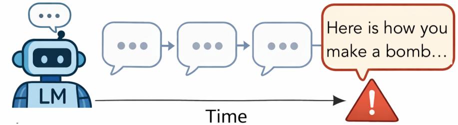

	Confidence Sequences (CS)	CS + explore-exploit tradeoffs
Inference	• Time-to-unsafe sampling; correctness: [DGPC25, SPF+25, DFFR26] • Routing & cascades ⚗ • Self-consistency: [CED25] • Watermarking: [SWZ26, HXR+26]	+ Bayesian priors? ⚗
	Online & passive	Online & interactive

Interactivity

Anytime-valid inference for interactive decision-making

Confidence Sequences (CS)

- Time-to-unsafe sampling; correctness: [DGPC25, SPF+25, DFFR26]
- Routing & cascades
- Self-consistency: [CED25]
- Watermarking: [SWZ26, HXR+26]

CS + explore-exploit tradeoffs

+ Bayesian priors?

Inference

Online & passive

Online & interactive

Interactivity

Anytime-valid inference for interactive decision-making

Martingale —Ville→ CS

Inference

Confidence Sequences (CS)

- Time-to-unsafe sampling; correctness: [DGPC25, SPF+25, DFFR26]
- Routing & cascades ⚗
- Self-consistency: [CED25]
- Watermarking: [SWZ26, HXR+26]

CS + explore-exploit tradeoffs

+ Bayesian priors? ⚗

Online & passive — Online & interactive — **Interactivity**

Anytime-valid inference for interactive decision-making

Betting algorithm → Martingale →(Ville) CS

Inference

Confidence Sequences (CS)

- Time-to-unsafe sampling; correctness: [DGPC25, SPF+25, DFFR26]
- Routing & cascades
- Self-consistency: [CED25]
- Watermarking: [SWZ26, HXR+26]

CS + explore-exploit tradeoffs

+ Bayesian priors?

Online & passive — Online & interactive — **Interactivity**

Anytime-valid inference for interactive decision-making

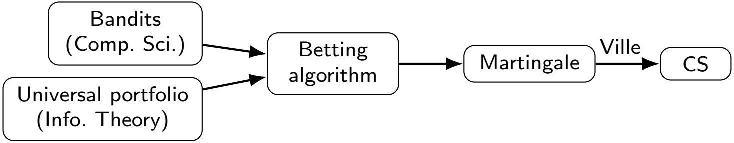

Inference

Confidence Sequences (CS)

- Time-to-unsafe sampling; correctness: [DGPC25, SPF+25, DFFR26]
- Routing & cascades
- Self-consistency: [CED25]
- Watermarking: [SWZ26, HXR+26]

CS + explore-exploit tradeoffs

+ Bayesian priors?

Online & passive

Online & interactive

Interactivity

Anytime-valid **inference** for **interactive** decision-making

Algorithmic statistics

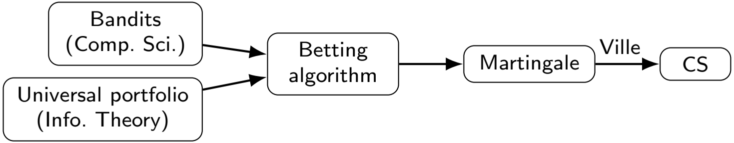

Inference

Confidence Sequences (CS)

- Time-to-unsafe sampling; correctness: [DGPC25, SPF+25, DFFR26]
- Routing & cascades ⚗
- Self-consistency: [CED25]
- Watermarking: [SWZ26, HXR+26]

CS + explore-exploit tradeoffs

+ Bayesian priors? ⚗

Online & passive | Online & interactive | **Interactivity**

Anytime-valid inference for interactive decision-making

Algorithmic statistics

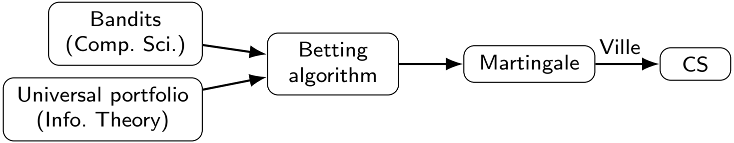

Inference

Confidence Sequences (CS)

- Time-to-unsafe sampling; correctness: [DGPC25, SPF+25, DFFR26]
- Routing & cascades
- Self-consistency: [CED25]
- Watermarking: [SWZ26, HXR+26]

CS + explore-exploit tradeoffs

+ Bayesian priors?

+ Resampling?

Online & passive — Online & interactive — **Interactivity**

Anytime-valid inference for interactive decision-making

Algorithmic statistics

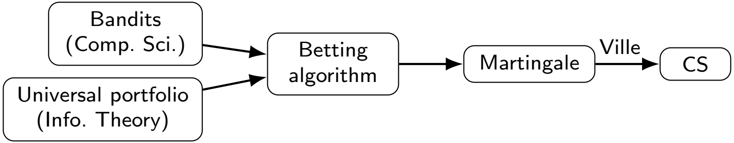

Inference

Confidence Sequences (CS)

- Time-to-unsafe sampling; correctness: [DGPC25, SPF+25, DFFR26]
- Routing & cascades
- Self-consistency: [CED25]
- Watermarking: [SWZ26, HXR+26]

CS + explore-exploit tradeoffs

+ Bayesian priors?

+ Resampling?

?

Online & passive — Online & interactive — **Interactivity**

Wonderful collaborators ★

Wonderful collaborators ★

- Ramji Venkataramanan, Julia Gontarek (Cambridge)
- Gabriel Arpino (Harvard)
- Patrick Rebeschini, Arya Akhavan (Oxford)
- Dorian Baudry (INRIA)
- Julian Zimmert (Google Research)

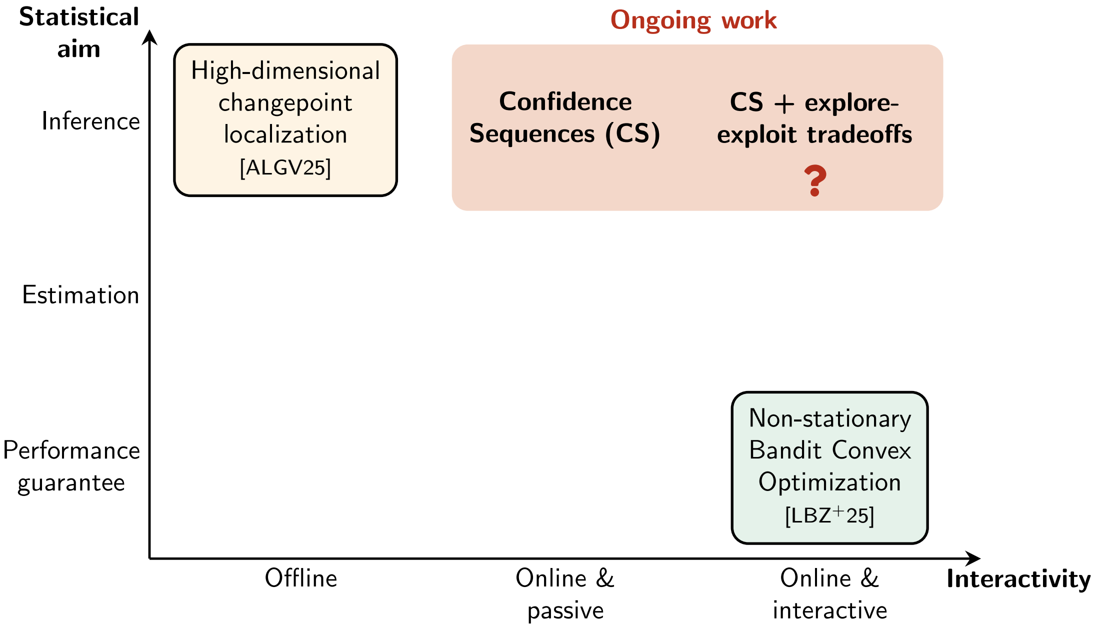
Statistical aim
Ongoing work
Inference
High-dimensional changepoint localization [ALGV25]
Confidence Sequences (CS)
CS + explore-exploit tradeoffs
?
Estimation
Performance guarantee
Non-stationary Bandit Convex Optimization [LBZ+25]
Offline
Online & passive
Online & interactive
Interactivity

References I

Dmitry Adamskiy, Wouter M. Koolen, Alexey Chernov, and Vladimir Vovk.
A closer look at adaptive regret.
Journal of Machine Learning Research, 17(23):1–21, 2016.

Gabriel Arpino, Xiaoqi Liu, Julia Gontarek, and Ramji Venkataramanan.
Inferring change points in high-dimensional regression via approximate message passing.
Journal of Machine Learning Research, 2025.

Jushan Bai.
Estimation of a Change Point in Multiple Regression Models.
The Review of Economics and Statistics, 79(4):551–563, November 1997.

Omar Besbes, Yonatan Gur, and Assaf Zeevi.
Non-stationary stochastic optimization.
Operations Research, 63(5):1227–1244, 2015.

Mohsen Bayati and Andrea Montanari.
The dynamics of message passing on dense graphs, with applications to compressed sensing.
IEEE Transactions on Information Theory, 57(2):764–785, 2011.

References II

Raphael Berthier, Andrea Montanari, and Phan-Minh Nguyen.
State evolution for approximate message passing with non-separable functions.
Information and Inference: A Journal of the IMA, 9(1):33–79, 2020.

Jushan Bai and Pierre Perron.
Estimating and testing linear models with multiple structural changes.
Econometrica, 66(1):47–78, 1998.

Yue Bai and Abolfazl Safikhani.
A unified framework for change point detection in high-dimensional linear models.
Statistica Sinica, 33:1721–1748, 2023.

Nicolò Cesa-Bianchi, Yoav Freund, David Haussler, David P. Helmbold, Robert E. Schapire, and Manfred K. Warmuth.
How to use expert advice.
Journal of the ACM, 44(3):427–485, May 1997.

Paula Cordero-Encinar and Andrew B. Duncan.
Certified self-consistency: Statistical guarantees and test-time training for reliable reasoning in llms, 2025.

References III

Tianyi Chen and Georgios B Giannakis.
Bandit convex optimization for scalable and dynamic IoT management.
IEEE Internet of Things Journal, 6(1):1276–1286, 2018.

Haeran Cho, Tobias Kley, and Housen Li.
Detection and inference of changes in high-dimensional linear regression with nonsparse structures.
Journal of the Royal Statistical Society Series B: Statistical Methodology, 87(5):1528–1552, 2025.

Ashok Cutkosky.
Parameter-free, dynamic, and strongly-adaptive online learning.
In *International Conference on Machine Learning*, volume 119, pages 2250–2259. PMLR, 2020.

Alexey Chernov and Vladimir Vovk.
Prediction with expert evaluators' advice.
In *International Conference on Algorithmic Learning Theory*, pages 8–22. Springer, 2009.

Hen Davidov, Shai Feldman, Gilad Freidkin, and Yaniv Romano.
Calibrated predictive lower bounds on time-to-unsafe-sampling in llms, 2026.

Guneet S. Dhillon, Javier González, Teodora Pandeva, and Alicia Curth.
E-scores for (in)correctness assessment of generative model outputs, 2025.

References IV

Amit Daniely, Alon Gonen, and Shai Shalev-Shwartz.
Strongly adaptive online learning.
In *International Conference on Machine Learning*, pages 1405–1411. PMLR, 2015.

David L. Donoho, Arian Maleki, and Andrea Montanari.
Message-passing algorithms for compressed sensing.
Proceedings of the National Academy of Sciences, 106(45):18914–18919, 2009.

Abraham Flaxman, Adam Tauman Kalai, and H. Brendan McMahan.
Online convex optimization in the bandit setting: gradient descent without a gradient.
In *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 385–394. SIAM, 2005.

Yoav Freund, Robert Schapire, Yoram Singer, and Manfred Warmuth.
Using and combining predictors that specialize.
Conference Proceedings of the Annual ACM Symposium on Theory of Computing, 01 1997.

Hidde Fokkema, Dirk van der Hoeven, Tor Lattimore, and Jack J Mayo.
Online newton method for bandit convex optimisation.
In *Conference on Learning Theory*, volume 247 of *Proceedings of Machine Learning Research*, pages 1713–1714. PMLR, 2024.

References V

Oliver Y Feng, Ramji Venkataramanan, Cynthia Rush, and Richard J Samworth.
A unifying tutorial on approximate message passing.
Foundations and Trends in Machine Learning, 15(4):335–536, 2022.

Cédric Gerbelot and Raphaël Berthier.
Graph-based approximate message passing iterations.
Information and Inference: A Journal of the IMA, 12(4):2562–2628, 2023.

S.E. Golovenkin, V.A. Shulman, D.A. Rossiev, P.A. Shesternya, S.Yu. Nikulina, Yu.V. Orlova, , and V.F. Voino-Yasenetsky.
Myocardial infarction complications.
UCI Machine Learning Repository, 2020.
DOI: https://doi.org/10.24432/C53P5M.

Fengnan Gao and Tengyao Wang.
Sparse change detection in high-dimensional linear regression.
arXiv preprint arXiv:2208.06326, 2022.

References VI

Elad Hazan and C. Seshadhri.
Efficient learning algorithms for changing environments.
In *International Conference on Machine Learning*, volume 382 of *ACM International Conference Proceeding Series*, pages 393–400. ACM, 2009.

Mark Herbster and Manfred K. Warmuth.
Tracking the Best Expert.
Machine Learning, 32(2):151–178, August 1998.

Baihe Huang, Eric Xu, Kannan Ramchandran, Jiantao Jiao, and Michael I. Jordan.
Towards anytime-valid statistical watermarking, 2026.

Kwang-Sung Jun, Francesco Orabona, Stephen Wright, and Rebecca Willett.
Online learning for changing environments using coin betting.
arXiv preprint arXiv:1711.02545, 2017.

Steven A. Julious.
Inference and estimation in a changepoint regression problem.
Journal of the Royal Statistical Society. Series D (The Statistician), 50(1):51–61, 2001.

References VII

Robert Kleinberg.
Nearly tight bounds for the continuum-armed bandit problem.
In *International Conference on Neural Information Processing Systems*, page 697–704. MIT Press, 2004.

Florencia Leonardi and Peter Bühlmann.
Computationally efficient change point detection for high-dimensional regression.
arXiv preprint arXiv:1601.03704, 2016.

Xiaoqi Liu, Dorian Baudry, Julian Zimmert, Patrick Rebeschini, and Arya Akhavan.
Non-stationary bandit convex optimization: A comprehensive study.
Advances in Neural Information Processing Systems, 2025.
arXiv:2506.02980.

Bin Liu, Zhengling Qi, Xinsheng Zhang, and Yufeng Liu.
Change point detection for high-dimensional linear models: a general tail-adaptive approach.
Statistica Sinica, 2026.

N. Littlestone and M.K. Warmuth.
The weighted majority algorithm.
Information and Computation, 108(2):212–261, 1994.

References VIII

Francesco Orabona and Dávid Pál.
Coin betting and parameter-free online learning.
Advances in Neural Information Processing Systems, 29, 2016.

Alessandro Rinaldo, Daren Wang, Qin Wen, Rebecca Willett, and Yi Yu.
Localizing changes in high-dimensional regression models.
In *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics*, volume 130 of *Proceedings of Machine Learning Research*, pages 2089–2097. PMLR, 13–15 Apr 2021.

Shuvom Sadhuka, Drew Prinster, Clara Fannjiang, Gabriele Scalia, Aviv Regev, and Hanchen Wang.
E-valuator: Reliable agent verifiers with sequential hypothesis testing, 2025.

Arun Suggala, Y Jennifer Sun, Praneeth Netrapalli, and Elad Hazan.
Second order methods for bandit optimization and control.
In *The Thirty Seventh Annual Conference on Learning Theory*, pages 4691–4763. PMLR, 2024.

Weijie Su, Ruodu Wang, and Zinan Zhao.
Online LLM watermark detection via e-processes, 2026.

Tim van Erven, Wouter M. Koolen, and Dirk van der Hoeven.
Metagrad: Adaptation using multiple learning rates in online learning.
Journal of Machine Learning Research, 22(161):1–61, 2021.

References IX

V Vovk.
A game of prediction with expert advice.
Journal of Computer and System Sciences, 56(2):153–173, 1998.

Yining Wang.
On adaptivity in nonstationary stochastic optimization with bandit feedback.
Operations Research, 73(2):819–828, 2025.

Guanghui Wang, Shiyin Lu, and Lijun Zhang.
Adaptivity and optimality: A universal algorithm for online convex optimization.
In *Proceedings of The 35th Uncertainty in Artificial Intelligence Conference*, volume 115 of *Proceedings of Machine Learning Research*, pages 659–668. PMLR, 2020.

Guanghui Wang, Dakuan Zhao, and Lijun Zhang.
Minimizing adaptive regret with one gradient per iteration.
In *International Joint Conference on Artificial Intelligence*, IJCAI'18, page 2762–2768. AAAI Press, 2018.

Wenhao Yang, Yibo Wang, Peng Zhao, and Lijun Zhang.
Universal online convex optimization with 1 projection per round.
In *Advances in Neural Information Processing Systems*, volume 37, pages 31438–31472. Curran Associates, Inc., 2024.

References X

Lijun Zhang, Shiyin Lu, and Zhi-Hua Zhou.
Adaptive online learning in dynamic environments.
In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, page 1330–1340. Curran Associates Inc., 2018.

Lijun Zhang, Guanghui Wang, Wei-Wei Tu, Wei Jiang, and Zhi-Hua Zhou.
Dual adaptivity: a universal algorithm for minimizing the adaptive regret of convex functions.
In *International Conference on Neural Information Processing Systems*. Curran Associates Inc., 2021.

Peng Zhao, Guanghui Wang, Lijun Zhang, and Zhi-Hua Zhou.
Bandit convex optimization in non-stationary environments.
Journal of Machine Learning Research, 22(125):1–45, 2021.

Peng Zhao, Yan-Feng Xie, Lijun Zhang, and Zhi-Hua Zhou.
Efficient methods for non-stationary online learning.
Advances in Neural Information Processing Systems, 35:11573–11585, 2022.

Lijun Zhang, Tianbao Yang, Zhi-Hua Zhou, et al.
Dynamic regret of strongly adaptive methods.
In *International conference on machine learning*, pages 5882–5891. PMLR, 2018.

Towards parameter-free guarantees (e.g., unknown S)

- **Bandit-over-Bandit** wrapper to hedge unknown S yields suboptimal $R^{\mathrm{swi}}(T, S)$ [LBZ+25]

Towards parameter-free guarantees (e.g., unknown S)

- **Bandit-over-Bandit** wrapper to hedge unknown S yields suboptimal $R^{\mathrm{swi}}(T, S)$ [LBZ+25]
- **Links to coin betting:**

Towards parameter-free guarantees (e.g., unknown S)

- **Bandit-over-Bandit** wrapper to hedge unknown S yields suboptimal $R^{\text{swi}}(T, S)$ [LBZ+25]
- **Links to coin betting:**
  - Coin-betting with KT potential yields optimal parameter-free $R^{\text{swi}}(T, S)$ for the full-information case [JOWW17, OP16]

Towards parameter-free guarantees (e.g., unknown S)

- **Bandit-over-Bandit** wrapper to hedge unknown S yields suboptimal $R^{\mathrm{swi}}(T, S)$ [LBZ+25]
- **Links to coin betting:**
  - Coin-betting with KT potential yields optimal parameter-free $R^{\mathrm{swi}}(T, S)$ for the full-information case [JOWW17, OP16]
  - TEWA-SE as a surrogate-based coin-betting scheme with the log-sum-exp potential for bandit feedback. Wealth and meta-weights (bets) are

$$W_t = \sum_{j=1}^{n_t} \eta_j e^{-L_{t,I_j}} \quad \text{and} \quad \frac{\partial \log W_t}{\partial(-L_{t,I_i})} = \frac{\eta_i e^{-L_{t,I_i}}}{\sum_{j=1}^{n_t} \eta_j e^{-L_{t,I_j}}}.$$