Clustering Items through Bandit Feedback

Finding the Right Feature out of Many

Authors

Maximilian Graf[1,a], Victor Thuot[2,b], Nicolas Verzelen[b]

[1] graf9@uni-potsdam.de, [2] victor.thuot@inrae.fr

[a] Universität Potsdam, Potsdam, Germany [b] INRAE, Mistea, Institut Agro, Univ Montpellier, Montpellier, France

Motivation

- forest patches, divided into two unknown groups
- want to use biodiversity sensors to recover groups
- which sensors are informative? how can we allocate them efficiently?

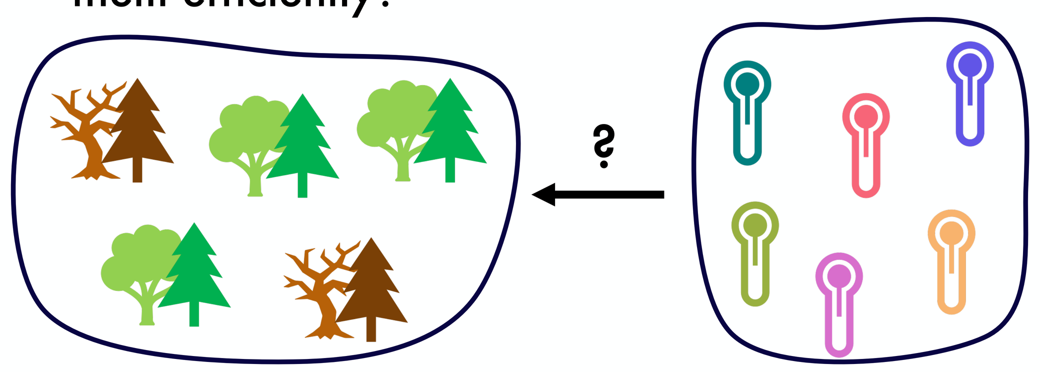

Main Results

Distribution-dependent control of the budget:

Let

$$H := \frac{d}{\theta}\frac{1}{\| \Delta \|^2} + \min_{s\in[d]}\left[\left(\frac{d}{s}+n\right)\frac{1}{\Delta^2_{(s)}}\right]$$

with $|\Delta_{(1)}| \geq |\Delta_{(2)}| \geq \cdots$. With probability $\geq 1-\delta$, `BanditClustering` recovers the true partition after $T \lesssim \log\left(\frac{1}{\delta}\right)\cdot H$ steps.

Matching minimax lower bound:

For any δ-PAC algorithm $\mathcal{A}$, we find a modification M' of M such that

$$\mathbb{P}_{M',\mathcal{A}}\left(T \geq \frac{2d}{\theta \| \Delta \|^2}\log\left(\frac{1}{6\delta}\right) \vee \frac{2(n-2)}{\| \Delta \|^2_\infty}\log\left(\frac{1}{4.8\delta}\right)\right) \geq \delta.$$

Mathematical Model

- n items, d features
- mean value of j-th feature on i-th item as matrix entry $M_{i,j}$
- **Assumption:** rows $M_{i,\cdot}$ either equal μ_0 or μ_1
- gap vector $\Delta := \mu_1 - \mu_0 \in \mathbb{R}^d \setminus \{0\}$
- minimal group proportion

$$\theta = \frac{|\{i: M_{i,\cdot} = \mu_0\}| \wedge |\{i: M_{i,\cdot} = \mu_1\}|}{n}$$

- Goal: cluster items ($M_{i,\cdot} = \mu_0$ vs $M_{i,\cdot} = \mu_1$)

	0	0	0	0	0	0
	0	1	0	0.5	0.05	0
	0	1	0	0.5	0.05	0
	0	0	0	0	0	0
	0	1	0	0.5	0.05	0

Example: Two Values

Δ being $h > 0$ in s entries, 0 in all others

with probability $\geq 1-\delta$, `BanditClustering` recovers true partition in $T \lesssim \log\left(\frac{1}{\delta}\right)\cdot H$ steps with

$$H = \frac{1}{\theta}\frac{d}{\| \Delta \|^2} + \frac{n}{h^2}$$

Recovering the Groups through Sequential Sampling

generating samples sequentially:

- can choose Indices $I_t \in [n]$, $J_t \in [d]$ at time $t = 1, 2, \ldots$
- observe feedback $X_t = M_{I_t, J_t}$ + random noise

Goal: given $\delta > 0$, want adaptive sampling strategy to recover partitions with probability $\geq 1-\delta$

→ δ-PAC algorithm `BanditClustering`

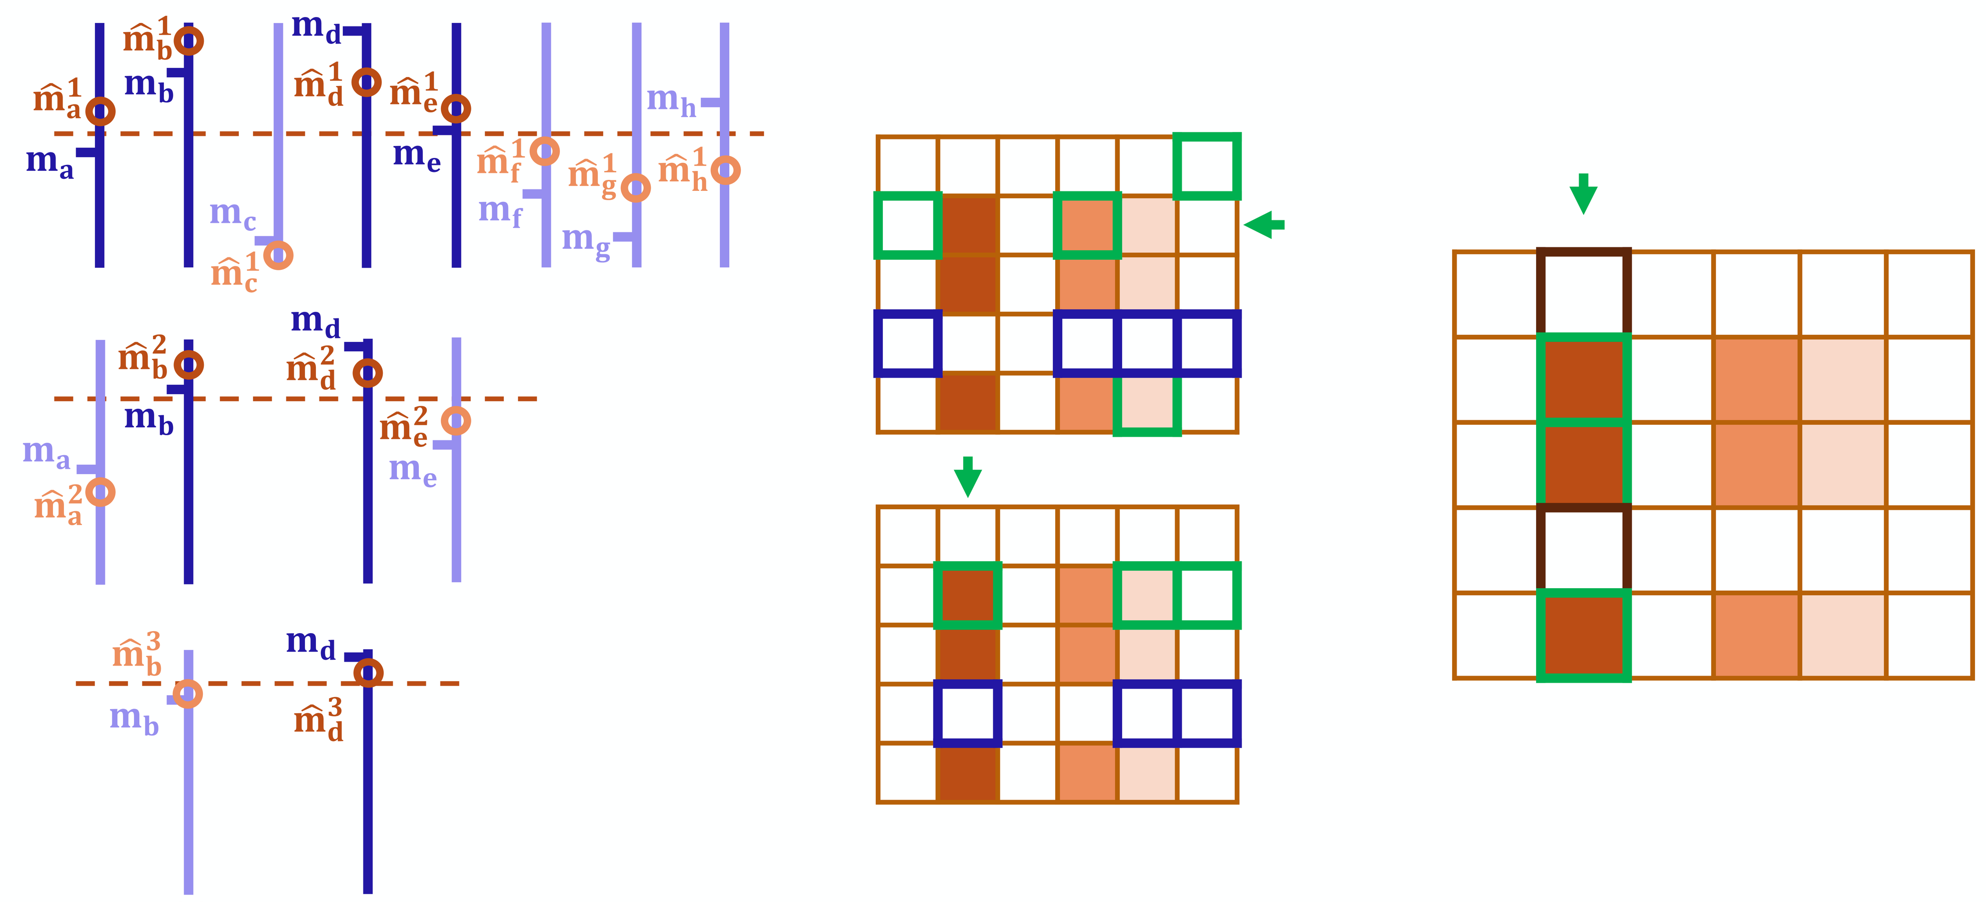

inspired by `Sequential Halving` for Best Arm Identification

subsample indices to find discriminative items/features

use informative feature for clustering step

Numerical Experiments

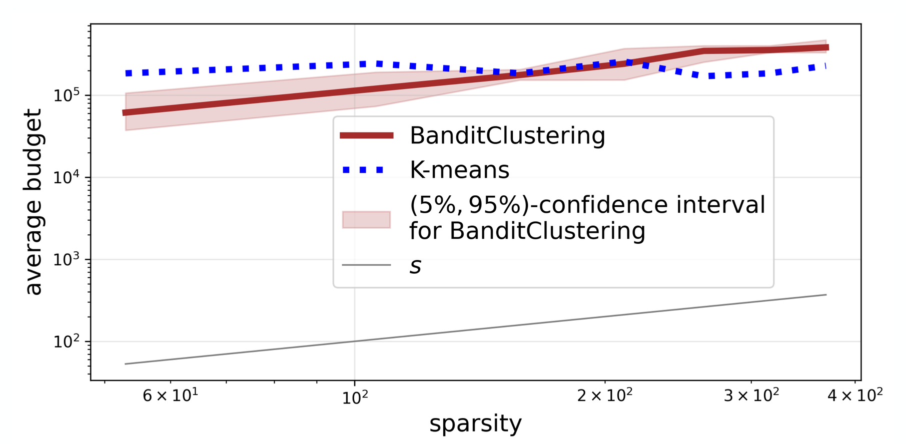

comparing `BanditClustering` and `KMeans` for varying s and fix $\| \Delta \|^2$, n, d, δ and θ

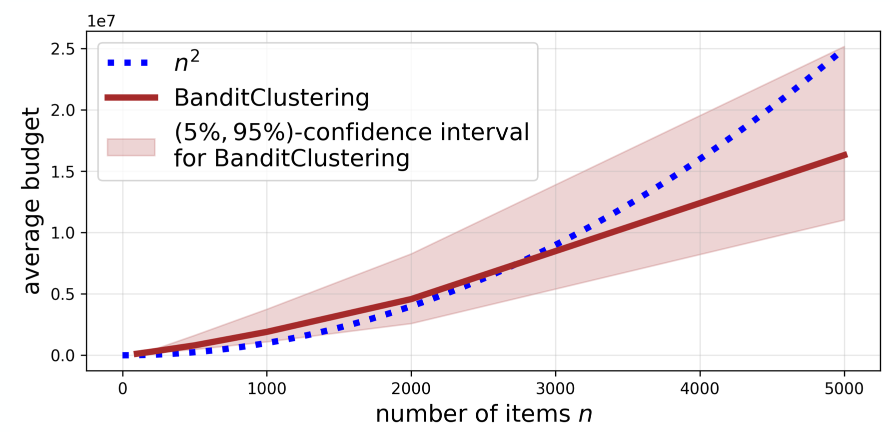

performance of `BanditClustering` for varying n, $d = 10 \cdot n$ and fix s, h, δ and θ

Some references:

ARIU, Kaito, et al. Optimal clustering from noisy binary feedback. *Machine Learning*, 2024

CASTRO, Rui M. Adaptive sensing performance lower bounds for sparse signal detection and support estimation. *Bernoulli*, 2014

SAAD, El Mehdi, et al. Active ranking of experts based on their performances in many tasks. In: *International Conference on Machine Learning*. PMLR, 2023

ZHAO, Yao, et al. Revisiting simple regret: Fast rates for returning a good arm. In: *International Conference on Machine Learning*. PMLR, 2023