Bandit Pareto Set Identification in a Multi-Output Linear Model

Cyrille Kone[1], **Emilie Kaufmann**[1], **Laura Richert**[2]

[1] Univ. Lille, Inria, CNRS, Centrale Lille, UMR 9198-CRIStAL, F-59000 Lille, France
[2] Univ. Bordeaux, Inserm, Inria, BPH, U1219, Sistm, F-33000 Bordeaux, France

Motivation

- Given K **multivariate** distributions (or arms), identify by adaptively sampling them the distributions whose average return is not uniformly worse than any other
- Each arm is associated with some observable features

Applications: clinical trials, large-scale recommender systems, software and hardware design, etc.

Problem setting

+ Sub-Gaussian distributions (or arms) over $\mathbb{R}^d$, $\nu_1, \dots, \nu_K$ with means (resp.) $\mu_1, \dots, \mu_K \in \mathbb{R}^d$ and **descriptive features** $x_1, \dots, x_K \in \mathbb{R}^h$
+ Linearity between vectors means and features i.e. $\mu_i = \Theta^\intercal x_i$ and $\Theta \in \mathbb{R}^{h \times d}$ **is unknown**
+ Large number of arms and a few descriptors: $h \ll K$.

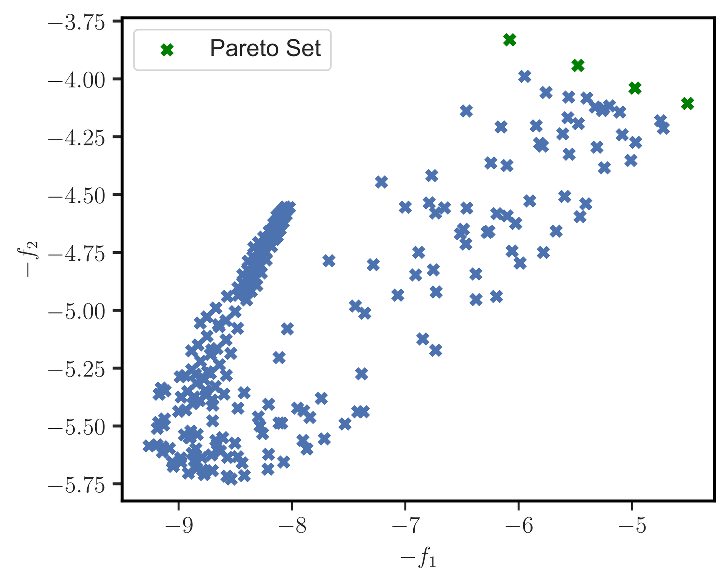

Figure 1: Energy consumption/runtime trade-off of synthesized application-specific networks on chip in Almer et al. [2011]. Each of the 259 architectures is described by 4 features: bus width, buffer depth, CPU clock, internal bandwidth.

Goal: Identify the Pareto Set

$$S^\star := \left\{ i \in [K] : \forall j \neq i, \exists c \in \{1, \dots, d\} : \mu_i(c) \geq \mu_j(c) \right\} ,$$

with a budget of T samples (**fixed-budget**, e.g., number of chips to synthesize or number of patients in a clinical trial) or with a maximum misidentification rate $\delta \in (0, 1)$ (**fixed-confidence**).

Contribution: we propose an elimination-based optimal-design algorithm for Pareto Set Identification (PSI) in a linear model with nearly-optimal guarantees in fixed-budget and fixed-confidence.

What is the difficulty of classifying an arm ?

For all arms i, j, introduce

$$\mathrm{M}(i,j) := \max_{c \in [d]} \left[\mu_i(c) - \mu_j(c)\right] \text{ and } \mathrm{m}(i,j) := \min_{c \in [d]} \left[\mu_j(c) - \mu_i(c)\right] ,$$

+ When $\mu_i \prec \mu_j$, $\mathrm{m}(i,j)$ is the minimum increase to a component of μ_i so that it is non-dominated by μ_j,
+ When $\mu_i \not\prec \mu_j$, $\mathrm{M}(i,j)$ is the smallest uniform increase of μ_j that makes it dominates μ_i.

For an arm i fixed,

- the smaller $\mathrm{M}(i,j)$, the "more j is close to be dominating" i (or dominates it by a large margin): i.e for any criterion $c \in [d]$, $(\mu_i(c) - \mu_j(c))$ is small.
- the larger $\min_{j \neq i} \mathrm{M}(i,j)$, the "more i is optimal": for any j, there exists a criterion $c_{i,j}$ such that $\mu_i(c_{i,j}) \gg \mu_j(c_{i,j})$

The complexity of unstructured PSI scales as a sum of $1/\Delta_i^2$ terms. For a sub-optimal arm $i \notin S^\star$,

$$\Delta_i := \max_{j \in S^\star} \mathrm{m}(i,j) ,$$

which is the smallest quantity that should be added component-wise to μ_i to make i appear Pareto optimal w.r.t $\{\mu_i : i \in [K]\}$. For a Pareto-optimal arm $i \in S^\star$,

$$\Delta_i := \min(\delta_i^+, \delta_i^-) , \quad \text{where}$$

$$\delta_i^+ := \min_{j \in S^\star \setminus \{i\}} [\min(\mathrm{M}(i,j), \mathrm{M}(j,i))] \text{ and } \delta_i^- := \min_{j \in [K] \setminus S^\star} \left[(\mathrm{M}(j,i))_+ + \Delta_j\right] , \text{with } (x)_+ := \max(x, 0).$$

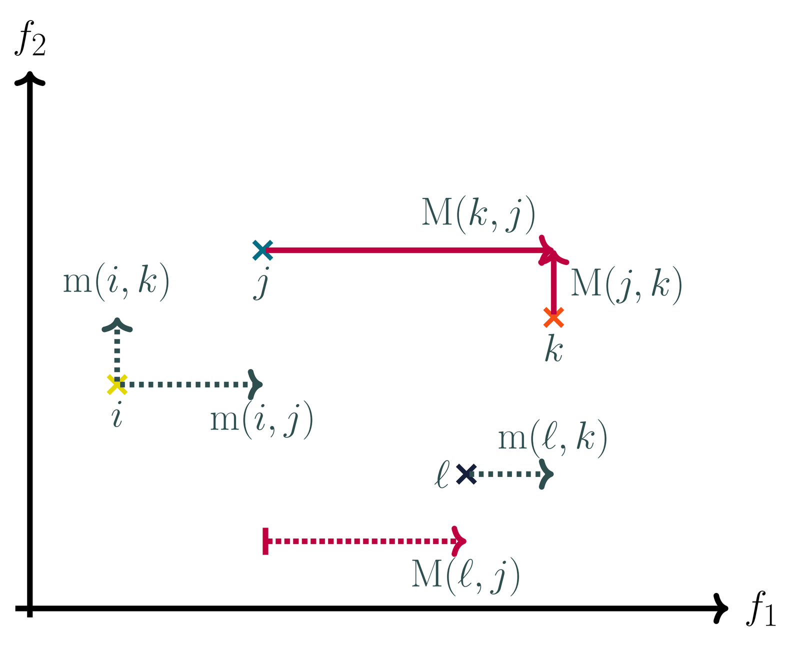

Figure 2: PSI "sub-optimality" gaps. Plain lines represent "distances" between Pareto optimal arms, and dashed lines are for margins from sub-optimal to optimal arms.

- The larger Δ_i the easier it is to classify i as optimal/sub-optimal
- The complexity of **unstrucured** PSI is characterized by

$$H_1(\boldsymbol{\mu}) = \sum_{i \in [K]} \frac{1}{\Delta_i^2} \text{ [fixed-confidence]} \quad ; \quad H_2(\boldsymbol{\mu}) := \max_{i \in [K]} \frac{i}{\Delta_i^2} \text{ [fixed-budget]} ,$$

where $\Delta_1 \leq \Delta_2 \leq \dots \leq \Delta_K$ (w.l.o.g).

Linear PSI

Lemma 1. *By deterministically sampling arms $a_1, a_2 \dots a_n \in [K]$, and collecting outcomes $y_1 := \Theta^\intercal x_{a_1} + \eta_1, \dots, y_n := \Theta^\intercal x_{a_n} + \eta_n \in \mathbb{R}^d$; introducing $X_n := (x_{a_1} \dots x_{a_n})^\intercal$, $Y_n := (y_1 \dots y_n)^\intercal$, the estimator $\widehat{\Theta}_n$ that minimizes $\mathbb{R}^{h \times d} \ni A \mapsto \|X_n A - Y_n\|_F^2$ satisfies for all $i \in \{a_1, \dots, a_n\}$, $\mathbb{E}[\widehat{\Theta}_n^\intercal x_i] = \mu_i$ and $Cov(\widehat{\Theta}_n^\intercal x_i) = \|x_i\|_{V_n^\dagger}^2 \Sigma$. Where $\eta, \eta_1, \dots, \eta_n$ are i.i.d centered, $Cov(\eta) = \Sigma$ and $V_n := X_n^\intercal X_n = \sum_{i \in [K]} N_n(i) x_i x_i^\intercal$; $N_n(i) = \sum_{s \in [n]} \mathbb{I}\{a_s = i\}$.*

G-optimal Empirical Gap Elimination [GEGE]

For $r = 1, 2, \dots, \tau$:

+ Comptute a G-optimal design i.e find

$$\mathbf{w}_r^\star \in \operatorname*{argmin}_{\mathbf{w} \in \boldsymbol{\Delta}_{A_r}} \max_{i \in A_r} \|x_i\|_{V_{r,\mathbf{w}}^{-1}}^2 \text{ with } \boldsymbol{\Delta}_{A_r} : \text{simplex on } A_r \text{ and } V_{r,\mathbf{w}} := \sum_{i \in A_r} \mathbf{w}(i) x_i x_i^\intercal , \qquad (1)$$

+ Use (1) and a **rounding** to collect n_r samples, compute $\widehat{\Theta}_{n_r}$ and **compute the empirical gaps**

$$\widehat{\Delta}_{i,r} := \begin{cases} \widehat{\Delta}_{i,r}^\star := \max_{j \in A_r \setminus \{i\}} \mathrm{m}(i,j;r) & \text{if } i \in A_r \setminus S_r, \\ \min_{j \in A_r \setminus \{i\}} \left[\mathrm{M}(i,j;r) \wedge (\mathrm{M}(j,i;r)_+ + (\widehat{\Delta}_{i,r}^\star)_+)\right] & \text{else} , \end{cases} \qquad (2)$$

where $\mathbf{A_r}$ is the set of active arms and $\mathbf{S_r}$, its empirical Pareto Set.
+ Discard a **proportion** p_r of the active arms
+ **Add the discarded empirically Pareto-optimal arms to B_τ**

Finally, the algorithm recommends B_τ as the Pareto Set and the goal is to upper-bound

$$\mathbb{P}_{\boldsymbol{\mu}}\left(B_R \neq S^\star\right) \text{ [fixed-budget]} \quad \text{and} \quad \sum_{r=1}^{\tau_\delta} n_r, \text{with } \mathbb{P}_{\boldsymbol{\mu}}(B_{\tau_\delta} \neq S^\star) \leq \delta \text{ [fixed-confidence]} .$$

Technical Challenge:
+ Ensure that the gaps defined in (2) properly estimate the PSI gaps
+ Avoid early discarding of some Pareto optimal arms; if not, sub-optimal arms may appear optimal

Theorem 1. *In the fixed-budget setting, given a budget $T \geq 45 h \log_2 h$,* GEGE *identifies the Pareto Set with a probability of error smaller than*

$$\mathcal{O}\left(\exp\left(-\frac{T}{H_{2,\mathrm{lin}}(\boldsymbol{\mu}) \lceil \log_2(h) \rceil}\right)\right) , \textit{with } H_{2,\mathrm{lin}}(\boldsymbol{\mu}) := \max_{i \in [h]} \frac{i}{\Delta_i^2} .$$

In the fixed-confidence setting, given $\delta \in (0,1)$, with probability $1 - \delta$, GEGE *correctly identifies the Pareto Set within $\log_2(2/\Delta_1)$ rounds, using at most*

$$\mathcal{O}(H_{1,\mathrm{lin}}(\boldsymbol{\mu}) \log(K d \log(H_{1,\mathrm{lin}})/\delta) \textit{ samples} , \textit{with } H_{1,\mathrm{lin}}(\boldsymbol{\mu}) = \sum_{i \in [h]} \frac{1}{\Delta_i^2} .$$

And these guarantees are minimax-optimal.

Some remarks :

- The complexity of Linear PSI only depends on the gaps of h arms
- $H_{q,\mathrm{lin}}(\boldsymbol{\mu}) \leq H_q(\boldsymbol{\mu})$ and in some instances, $H_q(\boldsymbol{\mu})/H_{q,\mathrm{lin}}(\boldsymbol{\mu}) = K/h$ for both $q = 1$ and $q = 2$
- When $h \ll K$, the guarantees of *GEGE* largely improves upon unstructured algorithms for PSI

Experimental results

- Sample complexity and misidentification rate vs number of arms for H_{lin}'s fixed

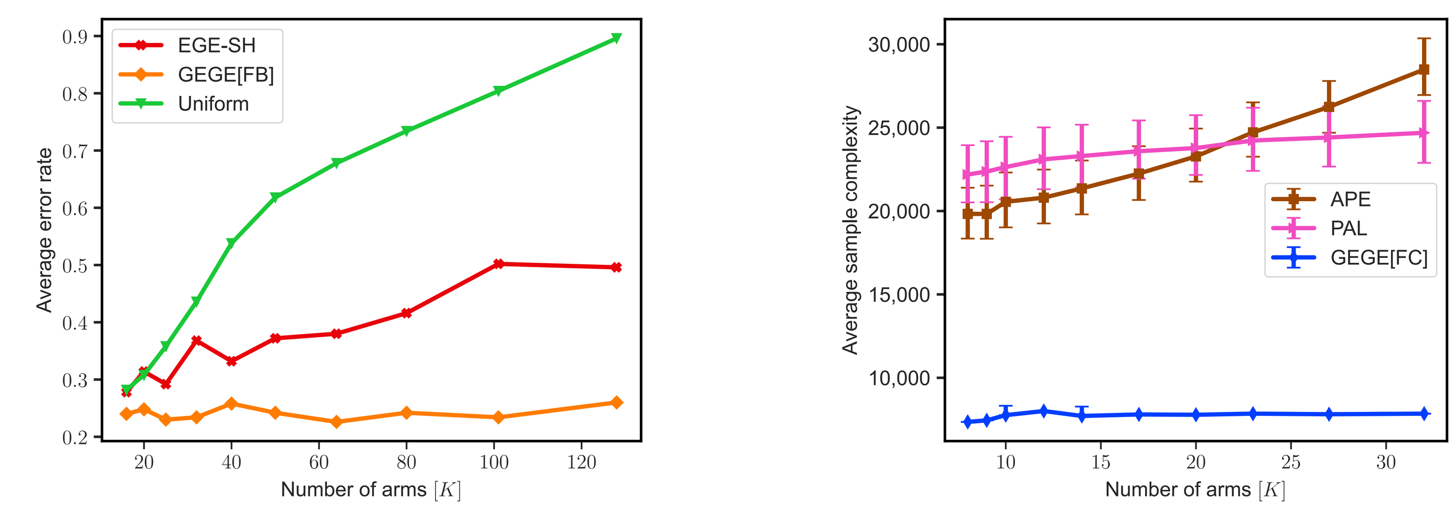

(a) Misidentification rate vs number of arms **(b)** Average sample complexity vs number of arms

Figure 3: Performance vs number of arms. $h = 8, d = 2, \delta = 0.01$(fixed-confidence). Statistics averaged over 500 runs.

- Performance on the NoC dataset

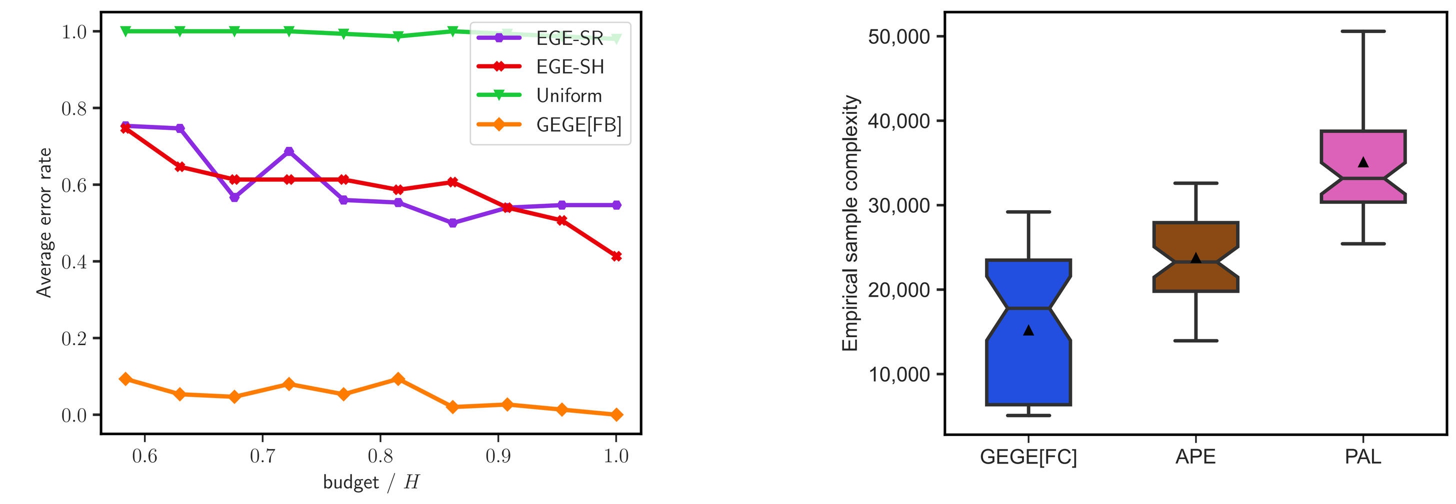

(a) Misidentification rate vs budget **(b)** Sample complexity for $\delta = 0.01$

Figure 4: Benchmark on the NoC dataset averaged over 500 runs. $h = 4, d = 2$ and $\delta = 0.01$ in fixed-confidence.

Conclusion and remarks

1) First optimal-design algorithm for PSI in fixed-budget and fixed-confidence regimes
2) $H_{1,\mathrm{lin}}(\boldsymbol{\mu})$ and $H_{2,\mathrm{lin}}(\boldsymbol{\mu})$ are good proxies to characterize the complexity of linear PSI in the fixed-confidence and fixed-budget settings