Adaptive Algorithms for Relaxed Pareto Set Identification

Cyrille Kone, Emilie Kaufmann, Laura Richert

Inria/CNRS/Univ-Lille/Univ-Bordeaux

Pareto Set Identification (PSI)

We are given K D-variates distributions (or arms) $\nu_1, \dots, \nu_K$ with means (resp.) $\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_K \in \mathbb{R}^D$

- i is dominated by j (or $\boldsymbol{\mu}_i \prec \boldsymbol{\mu}_j$) if: $\forall d \in [D],\ \mu_i^d < \mu_j^d$
- define $\mathcal{S}^\star$ the set of non-dominated (or Pareto-optimal) arms

Pareto Set Identification (PSI)

We are given K D-variates distributions (or arms) $\nu_1, \ldots, \nu_K$ with means (resp.) $\boldsymbol{\mu}_1, \ldots, \boldsymbol{\mu}_K \in \mathbb{R}^D$

- i is dominated by j (or $\boldsymbol{\mu}_i \prec \boldsymbol{\mu}_j$) if: $\forall d \in [D],\ \mu_i^d < \mu_j^d$
- define $\mathcal{S}^\star$ the set of non-dominated (or Pareto-optimal) arms

Goal: Given δ output $\widehat{S}_{\tau_\delta}$ s.t $\mathbb{P}(\tau_\delta < \infty$ and $\mathcal{S}^\star \neq \widehat{S}_{\tau_\delta}) \leq \delta$

Pareto Set Identification (PSI)

We are given K D-variates distributions (or arms) $\nu_1, \dots, \nu_K$ with means (resp.) $\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_K \in \mathbb{R}^D$

- i is dominated by j (or $\boldsymbol{\mu}_i \prec \boldsymbol{\mu}_j$) if: $\forall d \in [D],\ \mu_i^d < \mu_j^d$
- define $\mathcal{S}^\star$ the set of non-dominated (or Pareto-optimal) arms

Goal: Given δ output $\widehat{S}_{\tau_\delta}$ s.t $\mathbb{P}(\tau_\delta < \infty \text{ and } \mathcal{S}^\star \neq \widehat{S}_{\tau_\delta}) \leq \delta$

Possibly large sample complexity so we study relaxations

- PSI-k: find k Pareto-optimal arms
- (ε_1)-PSI (Auer et al. 2016): output $\widehat{S} \supset \mathcal{S}^\star$ and $\widehat{S}$ could contain some green points
- $(\varepsilon_1, \varepsilon_2)$-PSI: we are allowed to return one arm in each cluster

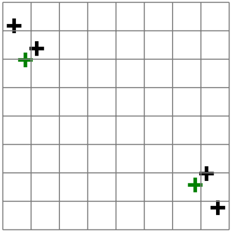

Pareto Set Identification (PSI)

We are given K D-variates distributions (or arms) $\nu_1, \ldots, \nu_K$ with means (resp.) $\boldsymbol{\mu}_1, \ldots, \boldsymbol{\mu}_K \in \mathbb{R}^D$

- i is dominated by j (or $\boldsymbol{\mu}_i \prec \boldsymbol{\mu}_j$) if: $\forall d \in [D],\ \mu_i^d < \mu_j^d$
- define $\mathcal{S}^\star$ the set of non-dominated (or Pareto-optimal) arms

Goal: Given δ output $\widehat{S}_{\tau_\delta}$ s.t $\mathbb{P}(\tau_\delta < \infty \text{ and } \mathcal{S}^\star \neq \widehat{S}_{\tau_\delta}) \leq \delta$

Possibly large sample complexity so we study relaxations

- PSI-k: find k Pareto-optimal arms
- (ε_1)-PSI (Auer et al. 2016): output $\widehat{S} \supset \mathcal{S}^\star$ and $\widehat{S}$ could contain some green points
- $(\varepsilon_1, \varepsilon_2)$-PSI: we are allowed to return one arm in each cluster

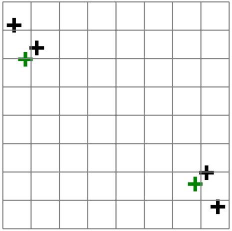

Our contribution: a single sampling strategy to tackle the three relaxations simultaneously.

Adaptive Pareto Exploration

Set $\Omega := [K]^2 \times [D]$ and define

- confidence intervals $[L^d_{i,j}(t,\delta), U^d_{i,j}(t,\delta)]$ s.t

$$\mathbb{P}\left(\forall t \geq 1,\ \forall (i,j,d) \in \Omega,\ (\mu^d_i - \mu^d_j) \in \left[L^d_{i,j}(t,\delta), U^d_{i,j}(t,\delta)\right]\right) \geq 1-\delta$$

- lower/upper CB on the "distance" between two arms i, j

$$\mathrm{M}^-(i,j,t) := \max_d L^d_{i,j}(t,\delta) \text{ and } \mathrm{M}^+(i,j,t) := \max_d U^d_{i,j}(t,\delta)$$

- nearly optimal arms at time t

$$\mathrm{OPT}^{\varepsilon_1}(t) := \left\{i \in [K] : \forall j \neq i, \mathrm{M}^-(i,j,t) + \varepsilon_1 > 0\right\}$$

Adaptive Pareto Exploration

Set $\Omega := [K]^2 \times [D]$ and define

- confidence intervals $[L^d_{i,j}(t,\delta), U^d_{i,j}(t,\delta)]$ s.t

$$\mathbb{P}\left(\forall t \geq 1,\ \forall (i,j,d) \in \Omega,\ (\mu_i^d - \mu_j^d) \in \left[L^d_{i,j}(t,\delta), U^d_{i,j}(t,\delta)\right]\right) \geq 1-\delta$$

- lower/upper CB on the "distance" between two arms i, j

$$\mathrm{M}^-(i,j,t) := \max_d L^d_{i,j}(t,\delta) \text{ and } \mathrm{M}^+(i,j,t) := \max_d U^d_{i,j}(t,\delta)$$

- nearly optimal arms at time t

$$\mathrm{OPT}^{\varepsilon_1}(t) := \left\{i \in [K] : \forall j \neq i, \mathrm{M}^-(i,j,t) + \varepsilon_1 > 0\right\}$$

Sampling rule: pull the least explored arm among b_t and c_t:

$$\mathbf{b_t} := \operatorname*{argmax}_{i \in [K] \setminus \mathrm{OPT}^{\varepsilon_1}(t)} \min_{j \neq i} \mathrm{M}^+(i,j,t),$$

$$\mathbf{c_t} := \operatorname*{argmin}_{j \neq b_t} \mathrm{M}^-(b_t,j,t)$$

Adaptive Pareto Exploration

Set $\Omega := [K]^2 \times [D]$ and define

- confidence intervals $[L^d_{i,j}(t,\delta), U^d_{i,j}(t,\delta)]$ s.t

$$\mathbb{P}\left(\forall t \geq 1,\ \forall (i,j,d) \in \Omega,\ (\mu_i^d - \mu_j^d) \in \left[L^d_{i,j}(t,\delta), U^d_{i,j}(t,\delta)\right]\right) \geq 1 - \delta$$

- lower/upper CB on the "distance" between two arms i, j

$$\mathrm{M}^-(i,j,t) := \max_d L^d_{i,j}(t,\delta) \text{ and } \mathrm{M}^+(i,j,t) := \max_d U^d_{i,j}(t,\delta)$$

- nearly optimal arms at time t

$$\mathrm{OPT}^{\varepsilon_1}(t) := \left\{ i \in [K] : \forall j \neq i, \mathrm{M}^-(i,j,t) + \varepsilon_1 > 0 \right\}$$

Sampling rule: pull the least explored arm among b_t and c_t:

$$\mathbf{b_t} := \operatorname*{argmax}_{i \in [K] \setminus \mathrm{OPT}^{\varepsilon_1}(t)} \min_{j \neq i} \mathrm{M}^+(i,j,t),$$

$$\mathbf{c_t} := \operatorname*{argmin}_{j \neq b_t} \mathrm{M}^-(b_t, j, t)$$

💡 b_t is the “most likely to be Pareto-optimal" in $[K] \setminus \mathrm{OPT}^{\varepsilon_1}(t)$

💡 c_t is the “most likely to be dominating (or close to)" b_t

Stopping and Recommendation

Let for all $i \in [K]$, ε_1

$$g_i(t) := \max_{j \neq i} [-\mathrm{M}^+(i,j,t)] \text{ and } h_i^{\varepsilon_1}(t) := \min_{j \neq i} \mathrm{M}^-(i,j,t) + \varepsilon_1$$

💡 if $g_i(t) > 0$ then i is not Pareto-optimal (w.h.p)
💡 if $h_i^{\varepsilon_1}(t) > 0$ then i is nearly Pareto-optimal (w.h.p)

Introduce

$$Z_1^{\varepsilon_1}(t) := \min_{i \in S(t)} h_i^{\varepsilon_1}(t) \quad \text{and} \quad Z_2^{\varepsilon_1}(t) := \min_{i \in S(t)^{\mathrm{c}}} \max(g_i(t), h_i^{\varepsilon_1}(t))$$

and let $S(t)$ be the current empirical Pareto set

Stopping and Recommendation

Let for all $i \in [K]$, ε_1

$$g_i(t) := \max_{j \neq i} [-\mathrm{M}^+(i,j,t)] \text{ and } h_i^{\varepsilon_1}(t) := \min_{j \neq i} \mathrm{M}^-(i,j,t) + \varepsilon_1$$

- if $g_i(t) > 0$ then i is not Pareto-optimal (w.h.p)
- if $h_i^{\varepsilon_1}(t) > 0$ then i is nearly Pareto-optimal (w.h.p)

Introduce

$$Z_1^{\varepsilon_1}(t) := \min_{i \in S(t)} h_i^{\varepsilon_1}(t) \text{ and } Z_2^{\varepsilon_1}(t) := \min_{i \in S(t)^c} \max(g_i(t), h_i^{\varepsilon_1}(t))$$

and let $S(t)$ be the current empirical Pareto set

	Stopping condition	Recommendation	Objective
τ_{ε_1}	$Z_1^{\varepsilon_1}(t) > 0 \wedge Z_2^{\varepsilon_1}(t) > 0$	$S(\tau_{\varepsilon_1}) \cup W(\tau_{\varepsilon_1})$	ε_1-PSI
$\tau_{\varepsilon_1,\varepsilon_2}$	$Z_1^{\varepsilon_1,\varepsilon_2}(t) > 0 \wedge Z_2^{\varepsilon_1,\varepsilon_2}(t) > 0$	$\mathrm{OPT}^{\varepsilon_1}(\tau_{\varepsilon_1,\varepsilon_2})$	$(\varepsilon_1,\varepsilon_2)$-PSI
τ^k	$\lvert \mathrm{OPT}^{\varepsilon_1}(t) \rvert \geq k$	$\mathrm{OPT}^{\varepsilon_1}(\tau^k)$	ε_1-PSI-k

with $W(t) := \{i \in S(\tau_{\varepsilon_1})^c : \nexists j \neq i : \mathrm{M}^+(i,j,\tau_{\varepsilon_1}) < 0\}$

Experiments

We benchmarked our algorithms against the state-of-the art on real-world and synthetic datasets

Real-world scenario (COV-BOOST trial (Munro et al. 2021)):

- Arms: 20 covid vaccines
- Measures: 3 immunogenicity indicators (2 indicators of antibody and 1 of cellular response)

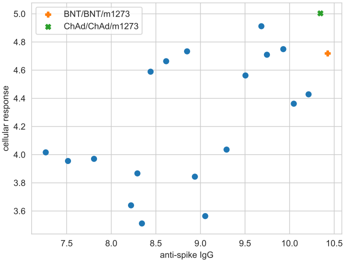

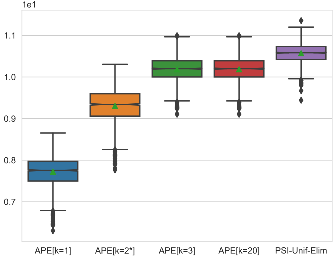

k-relaxation reduces the sample complexity

Conclusion and Future Work

- We proposed APE, an adaptive sampling rule that can be coupled with different stopping rules
- We proved the reductions in sample complexity brought by the relaxations
- We showcased the good performance of our algorithms compared to the state-of-the-art

Future working directions include

- Identify the Pareto set given a small budget
- Use component-wise slack $\boldsymbol{\varepsilon} := (\varepsilon^1, \ldots, \varepsilon^D)$