Online learning & multi-objective optimization

Shie Mannor & *Vianney Perchet* & Gilles Stoltz

Electrical Engineering
Techion

Lab. Proba. Modèles Aléa.
U. Paris-Diderot

GREGHEC
HEC

Conference On Learning Theory
June, 15 2014

Introduction

- Online linear optimization / Regret / Bandits

  "Same" guarantees with adversarial data than i.i.d. (no regret)

- "Our" question (shared with [Azar, Feige, Feldman, Tennenholtz])

  Same statement for multi-objective optimization ?

- **Spoiler Alert.**

  No

Online linear optimization in a nutshell

- Finite Action set $[A]$, (unknown) outcome set $\mathcal{U} \subset \mathbb{R}^A$
- At stage $t = 1, ..., T$:
  - $\rightarrow$ DM chooses (at random) $a_t \in [A] \sim x_t \in \Delta([A])$
  - $\rightarrow$ Nature reveals $U_t \in \mathcal{U}$
  - $\rightrightarrows$ Loss $U_t^{a_t}$ or $x_t.U_t = \mathbb{E}_{x_t} U_t^a$
- Best in hindsight $\min_{a\in[A]} \overline{U}_T^a = \min_{x\in\Delta([A])} x.\overline{U}_T = x.\frac{1}{T}\sum_{t=1}^T U_t$

Best in hindsight is "achievable" - Regret minimization

$$\frac{1}{T}\sum_{t=1}^{T} x_t.U_t - \min_{x\in\Delta([A])} x.\overline{U}_T \leq 4\sqrt{\frac{\log(A)}{T}} \max_{U\in\mathcal{U}} \|U\|_\infty$$

- "Online Optimization" of the linear function $x \mapsto x.\overline{U}_T$

Multi-objective optimization

– d different losses to control simultaneously:

"Optimization" of a vector valued mapping $F : \mathcal{X} \to \mathbb{R}^d$

– Increasing filtration of sets $\{\mathcal{C}[\alpha]\}_{\alpha \in \overline{\mathbb{R}}_+}$, $\mathcal{C}[\infty] = \mathbb{R}^d$

Optimize: Find the minimal (α^*, x^*), i.e.

i) $F(x^*) \in \mathcal{C}[\alpha^*]$

ii) if $F(x) \in \mathcal{C}[\alpha]$, then $\alpha \geq \alpha^*$

– $\mathcal{C}[\alpha]$: α-expansion of $\mathcal{C}_0$ or level sets of $f : \mathbb{R}^d \to \mathbb{R}$ when optimizing $x \mapsto f(F(x))$

Multi-objective optimization

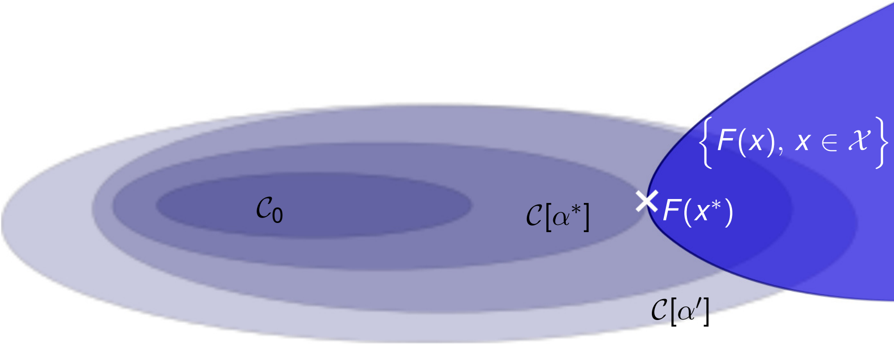

Online multi-objective optimization

- Same framework as online linear optimization except

Outcomes are vector of **vectors of dimension d**

- Nature reveals $V_t \in \mathcal{V} \subset (\mathbb{R}^d)^A$
- Vectorial loss $V_t^{a_t} \in \mathbb{R}^d$ or $x_t \odot V_t = \mathbb{E}_{x_t} V_t^a \in \mathbb{R}^d$
- Best in hindsight w.r.t. filtration $\mathcal{C}[\alpha]$:

$$\alpha^*(\overline{V}_T) = \inf \left\{ \alpha \in \mathbb{R}_+ \text{ s.t. } \exists x \in \Delta([A]),\ x \odot \overline{V}_T \in \mathcal{C}[\alpha] \right\}$$

Best in hindsight “achievable” ?

$$\frac{1}{T} \sum_{t=1}^{T} x_t \odot V_t \to \mathcal{C}[\alpha^*(\overline{V}_T)] \ \ \ldots \ \ \text{rate?}$$

Online multi-objective optimization

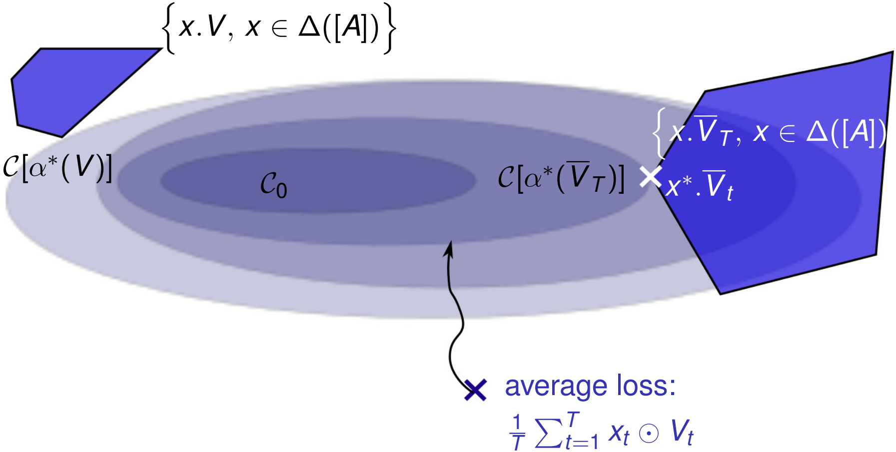

Examples

- Blackwell's approachability
  - $\rightarrow$ $\mathcal{C}[\alpha] = \mathcal{C}_0$ for all $\alpha \in \mathbb{R}_+$
- Approachability under budget limit
  - $\rightarrow$ Approachability: $\mathcal{C}[\alpha] = \mathcal{C}_0$ for all α.
  - $\rightarrow$ cost. $\frac{1}{T}\sum_{t=1}^{T} x_t.U_t \in [0, b]$
- Regret minimization under path constraints
  - $\rightarrow$ Regret: minimize $\frac{1}{T}\sum_{t=1}^{T} x_t.U_t \in \mathbb{R}$.
  - $\rightarrow$ Path/budget constraints $\frac{1}{T}\sum_{t=1}^{T} x_t \odot V_t \in \mathcal{C} \subset \mathbb{R}^d$
- Global cost minimization
  - $\rightarrow$ Surplus/Deficit minimization; min $\left|\frac{1}{T}\sum_{t=1}^{T} x_t.U_t\right|$
  - $\rightarrow$ Minimize $f\left(\frac{1}{T}\sum_{t=1}^{T} x_t \odot V_t\right) \in \mathbb{R}$

Blackwell's approachability [56]

- Constant Filtration. $\mathcal{C}[\alpha] = \mathcal{C}_0$ the "target set"
- Known vector outcome set $\mathcal{V}$. (not required for our results)
- Blackwell's necessary and sufficient condition s.t. $\frac{1}{T}\sum_{t=1}^{T} x_t \odot V_t$ can converge to $\mathcal{C}_0$ against any $\{V_t\}_{t\in\mathbb{N}}$

$$\forall V \in \mathcal{V}, \quad \exists x \in \Delta([A]), \quad x \odot V \in \mathcal{C}_0$$

Blackwell's algorithm – and many variants

At each stage $t \in \mathbb{N}$, two steps:

1) Projection on $\mathcal{C}_0$ + 2) Resolution of a LP

Main results – Outline

$\mathcal{C}[\alpha]$: α-expansion of $\mathcal{C}_0$

- Best in hindsight α^* not always achievable

  with a "simple" example

- A trivial relaxation – cav$[\alpha^*]$ – achievable

  a consequence of Blackwell's result

- We can do better

  A family of achievable mappings $\beta(\cdot) \leq$ cav$[\alpha^*](\cdot)$

- New algorithm.

  No projections, rate $T^{-1/4}$

Best in hindsight not achievable

- $[A] = \{1, 2\}$ and $\mathcal{V} = [-1, 1]^2$ so $x \odot V = x_1 V_1 + x_2 V_2 \in [-1, 1]$
- Minimization $\left| \frac{1}{T} \sum_{t=1}^{T} x_t \odot V_t \right| \in [0, 1]$ (surplus/deficit)
- $\mathcal{C}[\alpha] = [0, \alpha]$

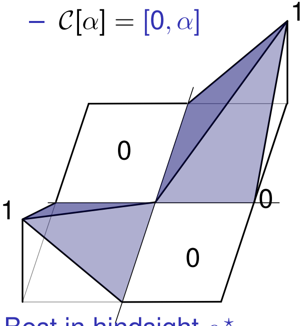

Best in hindsight $\alpha^\star$

Ideal objectives:

$$\left| \frac{1}{T} \sum_{t=1}^{T} x_t \odot V_t \right| \leq \alpha^*(\overline{V}_t) + o(1)$$

$$\left(\overline{V}_T, \left| \frac{1}{T} \sum_{t=1}^{T} x_t \odot V_t \right| \right) \to \text{hypograph}$$

A simple counterexample – Horizon $2T$

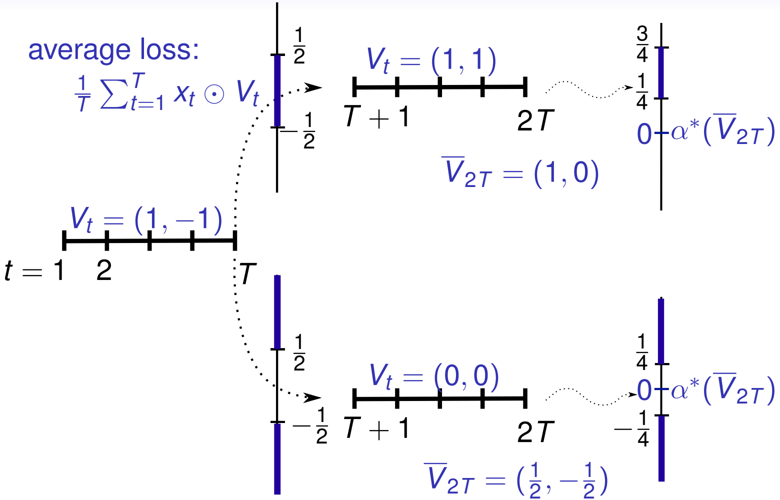

Achievable guarantees

- cav[α^*], concavification of α^* is achievable

  $\Rightarrow$ Hypo-graph cav[α^*] is convex and approachable

- Analytic formulation of cav:

$$\text{cav}[\alpha^*](V) = \sup \left\{ \sum \lambda_i \left| x^*(V_i) \odot V_i \right|, \quad \sum \lambda_i V_i = V \right\}$$

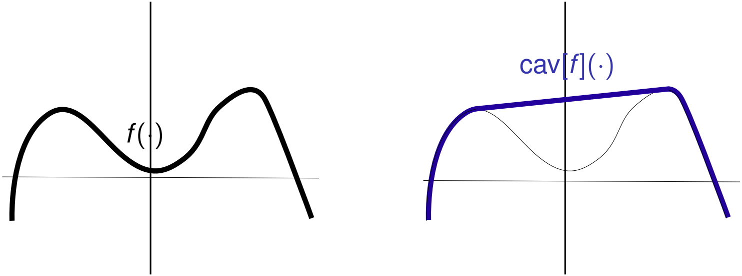

Concavification of $f : \mathbb{R} \to \mathbb{R}$

Achievable guarantees

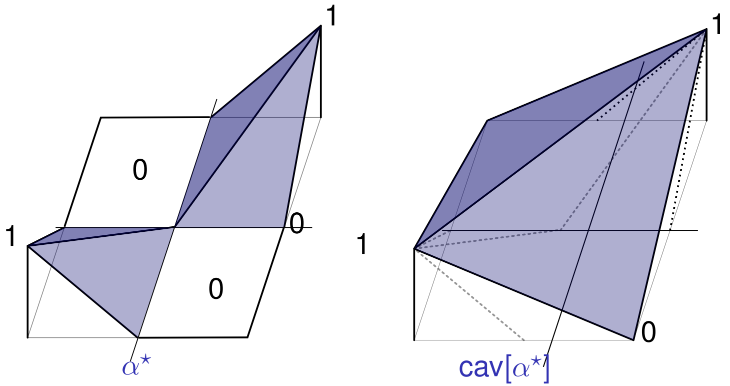

Achievable guarantees

- cav[α^*], concavification of α^* is achievable

  $\Rightarrow$ Hypo-graph cav[α^*] is convex and approachable

- Analytic formulation of cav:

$$\text{cav}[\alpha^*](V) = \sup\left\{\sum \lambda_i \left| x^*(V_i) \odot V_i \right|, \quad \sum \lambda_i V_i = V\right\}$$

Actually achieve better

$$\underline{\text{cav}}[x^\star](V) := \sup\left\{\left| \sum \lambda_i x^*(V_i) \odot V_i \right|, \quad \sum \lambda_i V_i = V\right\}$$

Achievable guarantees

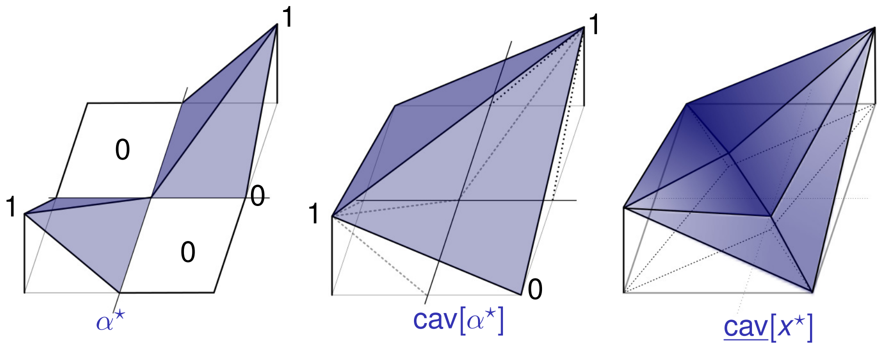

Achievable guarantees

- cav[α^*], concavification of α^* is achievable

  $\Rightarrow$ Hypo-graph cav[α^*] is convex and approachable

- Analytic formulation of cav:

$$\text{cav}[\alpha^*](V) = \sup\left\{\sum \lambda_i \left| x^*(V_i) \odot V_i \right|, \quad \sum \lambda_i V_i = V\right\}$$

Actually achieve better

$$\underline{\text{cav}}[x^\star](V) := \sup\left\{\left| \sum \lambda_i x^*(V_i) \odot V_i \right|, \quad \sum \lambda_i V_i = V\right\}$$

- And even

$$\underline{\text{cav}}[\psi](V)) := \sup\left\{\left| \sum \lambda_i \psi(V_i) \odot V_i \right|, \quad \sum \lambda_i V_i = V\right\}$$

Achievable guarantees

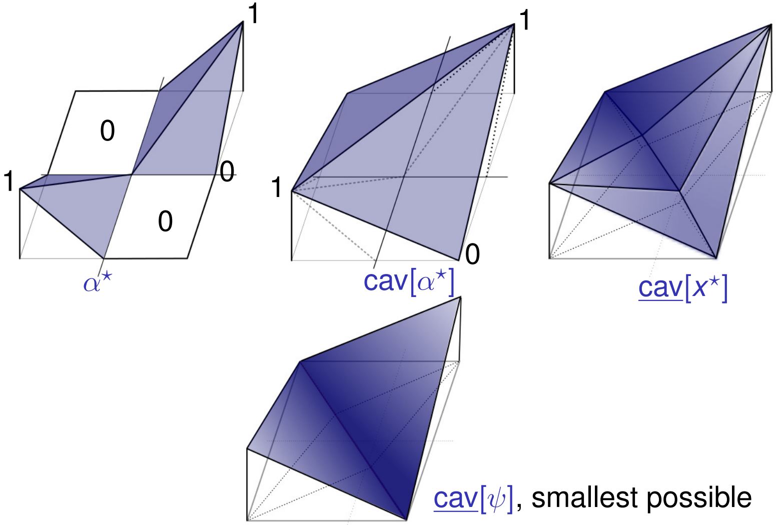

New algo based on regret minimization

1) Decompose $\{1, ..., T\}$ in $\sqrt{T}$ blocks of lengths $\sqrt{T}$.

$\mathbf{V}_k$ average outcome on k-th block

2) At the end of block K, compute direction

$$\vec{\lambda}_K = \frac{1}{T}\sum_{t=1}^{K\sqrt{T}} x_t \odot V_t - \frac{1}{T}\sqrt{T}\sum_{k=1}^{K} \psi(\mathbf{V}_k) \odot \mathbf{V}_k$$

3) Block $K+1$: minimize regret w.r.t. $U_t = \langle V_t, \vec{\lambda}_K \rangle$

Distance between $\frac{1}{T}\sum_{t=1}^{T} x_t \odot V_t$ and $\mathcal{C}[\underline{\text{cav}}[\psi](\overline{V}_T)]$

roughly $\|\vec{\lambda}_{\sqrt{T}}\|$, smaller than $16\,\sup_{V\in\mathcal{V}} \|V\| \sqrt{\log(A)}\, T^{-1/4}$

Conclusion, links

- Cannot achieve best in hindsight... can still achieve something
- In the classical approachability framework:
  - → Approach. algo. without projections [Bernstein-Shimkin] or LP
  - → Only $\sqrt{T}$ calls to $\psi(\cdot)$.
  - → cv. to the smallest expansion of $\mathcal{C}$ (without computing it / NP hard)
  - → Knowledge of $\mathcal{V}$ not required
  - → Another example of Approachability via Regret [Abernethy, Bartlett, Hazan]
- Same guarantee in anytime version
- Improve results (explicit rates – simpler strategies) on global cost [Even-Dar, Kleinberg, Mannor, Mansour], regret minimization with path constraints [Mannor, Tsitsiklis, Yu], optimistic approachability [Bernstein, Mannor, Shimkin]...
- Go beyond existing criteria [Azar, Feige, Feldman, Tennenholtz]