Université de Montréal
Mila

Optimistic and Adaptive Lagrangian Hedging

Ryan D'Orazio[1] Ruitong Huang[2]

[1]Université de Montréal, MILA [2]Borealis AI

Contact: ryan.dorazio@mila.quebec

BOREALIS AI

Problem

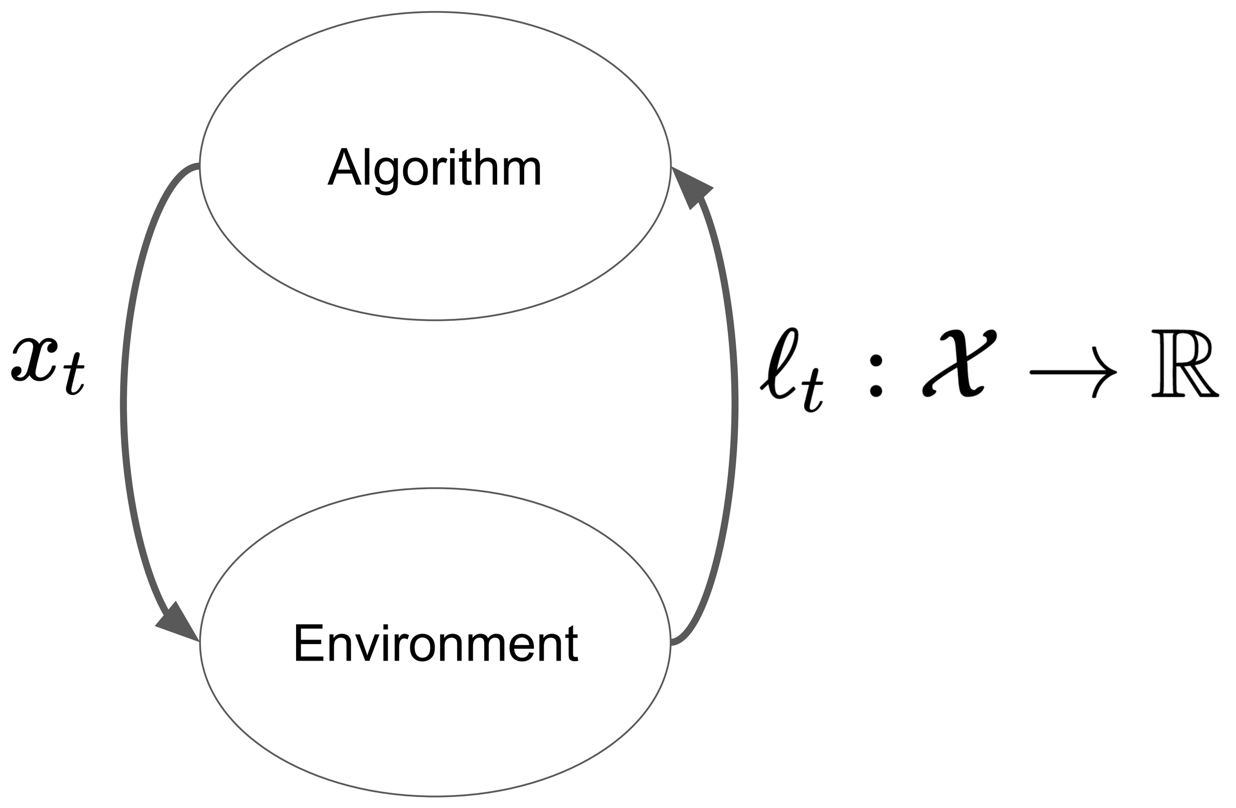

Performance after T Rounds

$$R_{\mathcal{X}}^T = \sum_{t=1}^T \ell_t(x_t) - \min_{x\in\mathcal{X}} \sum_{t=1}^T \ell_t(x)$$

Assumptions

1. ℓ_t are linear $\ell_t(\cdot) = \langle \ell_t, \cdot \rangle$[a]
2. $\mathcal{X}$ convex and compact

Objective

Vanishing average regret

$$\frac{R_{\mathcal{X}}^T}{T} \to 0 \text{ as } T \to \infty$$

Lower Bound

$$R_{\mathcal{X}}^T \in \Omega\left(\sqrt{T}\right)$$

[a]Linear losses are enough for convex losses.

Motivation

Some problems are **predictable**!

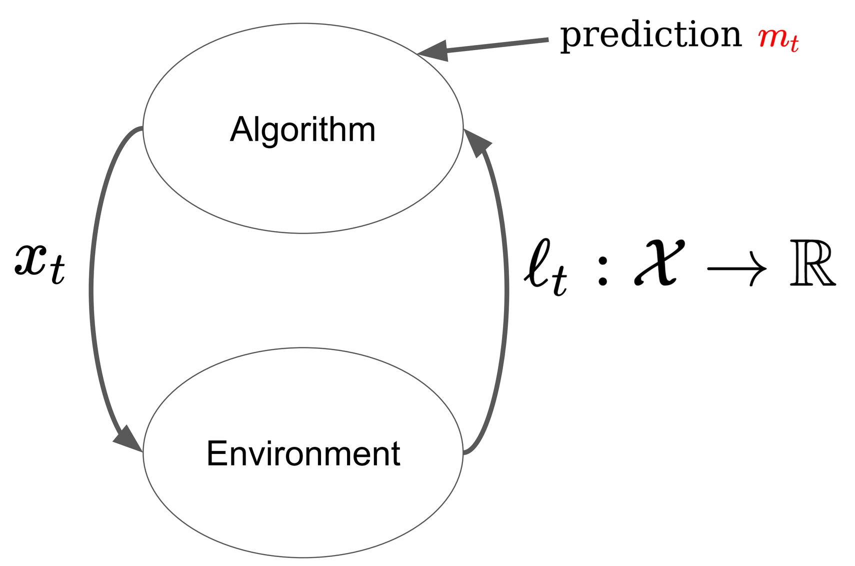

- If $m_t \approx \ell_t$ then regret should be small
- How do we:
  1. incorporate the prediction m_t to leverage predictable problems?
  2. maintain worst-case performance in online setting but perform well on predictable problems?

$$\underbrace{\textbf{Optimistic}}_{\text{Use a prediction } m_t} + \underbrace{\textbf{adaptive algorithms}}_{\text{Use adaptive stepsizes } \eta_t}$$

Examples of predictable problems:

- L smooth convex optimization $\|\nabla\ell(x) - \nabla\ell(y)\| \leq L\,\|x-y\|_*$
- Convex concave games with L smooth gradients

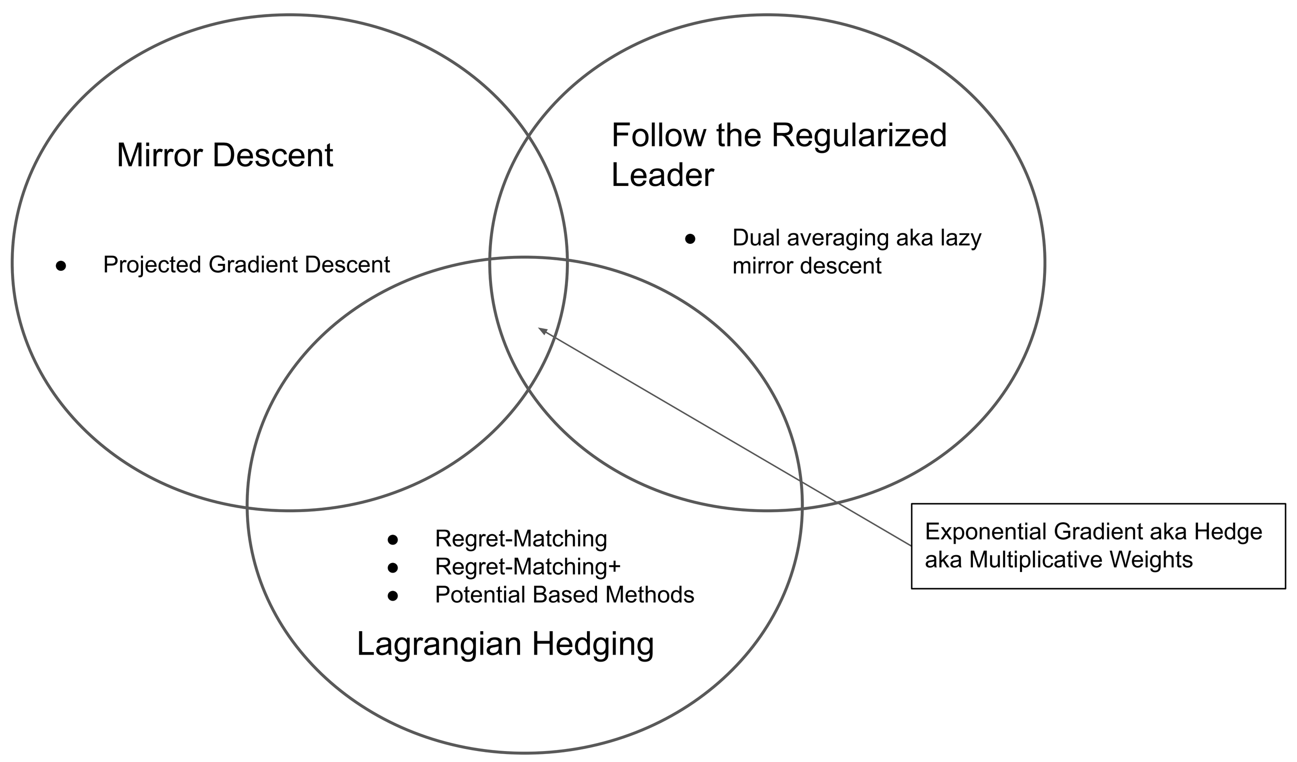

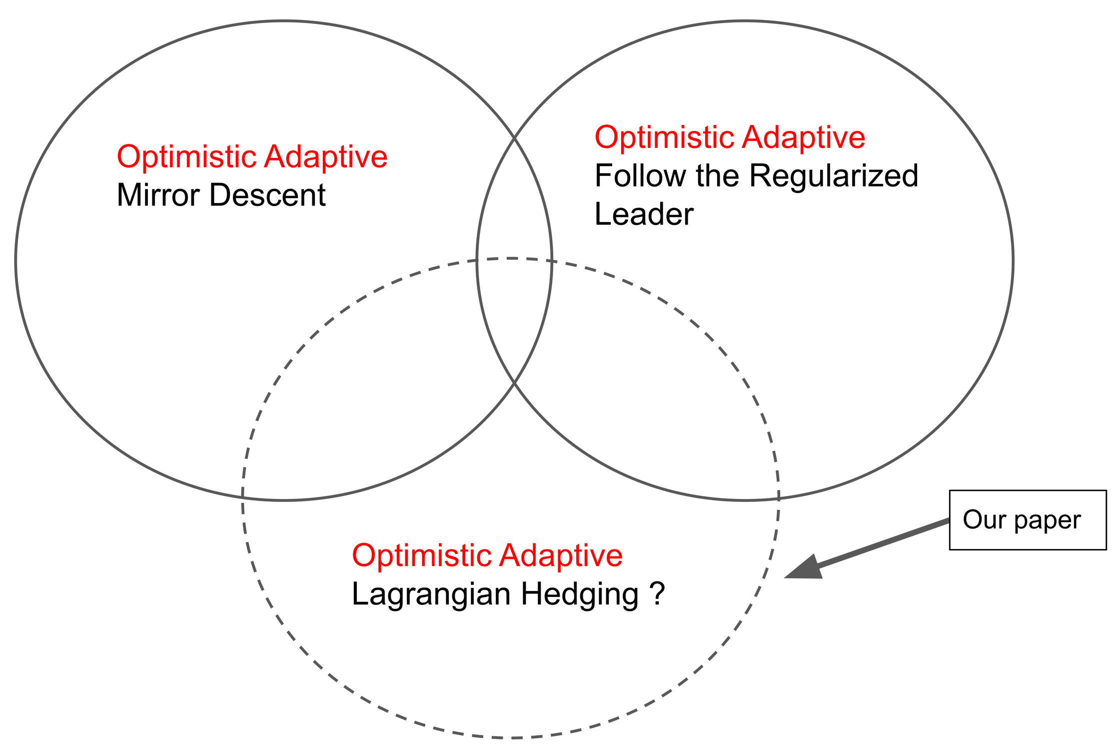

Lagrangian Hedging

Generalizes the following in one framework:

- Regret-matching and regret-matching+
- Polynomial weighted averaging and polynomial weighted averaging+
- Hedge aka multiplicative weights

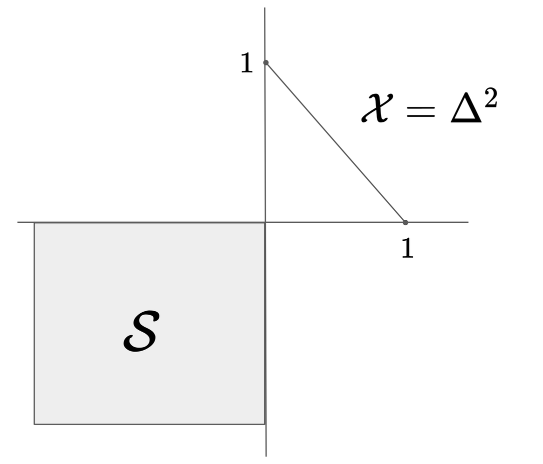

$$R_{\mathcal{X}}^T \leq \inf_{s\in\mathcal{S}} \|s_{1:T} - s\| \max_{x\in\mathcal{X}} \|x\|_*$$

Lagrangian Hedging Setup

- Find u such that $\langle u, x\rangle = 1 \quad \forall x \in \mathcal{X}$
  - Otherwise set $\mathcal{X} = \{(x,1) \,|x\in\mathcal{X}\}$ and pick $u = (0,\cdots,0,1)$
- Set $\mathcal{S}$ to be the polar cone to $\mathcal{X}$, $\mathcal{S} = \{s|\langle s, x\rangle \leq 0 \quad \forall x\in\mathcal{X}\}$
- At round t maintain regret vector

$$s_{1:t-1} = \sum_{k=1}^{t-1} \langle \ell_k, x_k\rangle u - \ell_k = \sum_{k=1}^{t-1} s_k$$

Objective

$$\frac{\inf_{s\in\mathcal{S}} \|s_{1:T} - s\|}{T} \to 0 \text{ as } T \to \infty$$

Optimistic and Adaptive Lagrangian Hedging

1. Pick a convex smooth potential function F
2. At time t play the strategy

$$x_t = \begin{cases} \frac{\nabla F(\eta_t(s_{1:t-1}+m_t))}{\langle \nabla F(\eta_t(s_{1:t-1}+m_t)), u\rangle} & \text{if } \langle \nabla F(\eta_t(s_{1:t-1}+m_t)), u\rangle > 0 \\ \text{arbitrary } x\in\mathcal{X} & o.w. \end{cases}$$

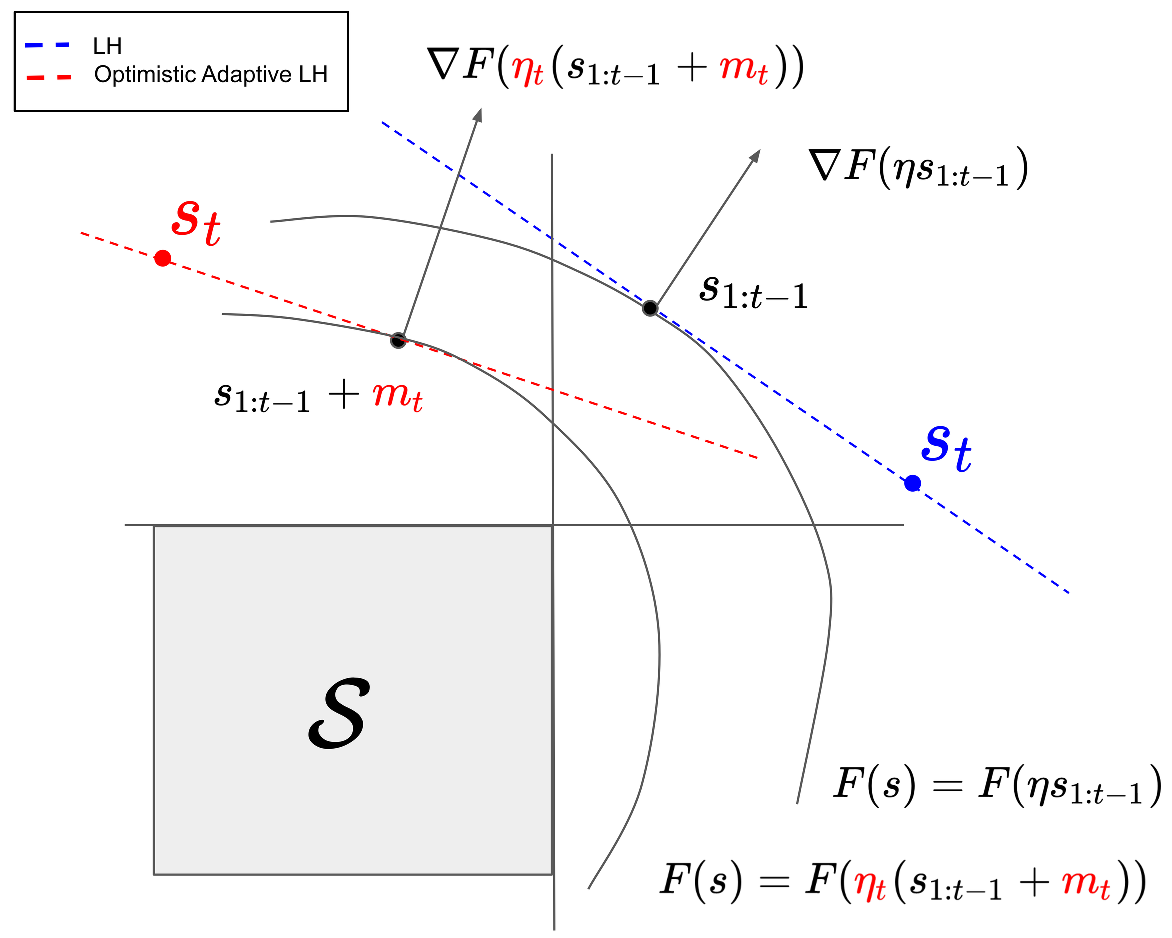

Results

1. Step-size free algorithms (e.g. regret-matching and regret-matching+) that guarantee

$$R_{\mathcal{X}}^T \in O\left(\sqrt{\sum_{t=1}^T \|s_t - m_t\|^2}\right)$$

2. When a stepsize is needed (e.g. hedge)

$$\eta_t = \frac{1}{\sqrt{\frac{1}{\eta_1^2} + \sum_{k=1}^{t-1}\|s_k - m_k\|^2}} \implies$$

$$R_{\mathcal{X}}^T \in O\left(\sqrt{\frac{1}{\eta_1^2} + \sum_{t=1}^{T-1}\|s_t - m_t\|^2}\right)$$

3. Regret with a fixed Smooth Convex Loss

$$R_{\mathcal{X}}^T \in O\left(\sqrt{\sum_{t=1}^T \|x_t - x_{t-1}\|^2}\right)$$

4. Adaptive and Optimistic Regret Bounds for Φ-Regret
   - Includes new optimistic and adaptive algorithms to minimize internal regret and more!

Related Work

Matching regret-bounds (up to a constant) with:

Gabriele Farina, Christian Kroer, Tuomas Sandholm G. Farina, C. Kroer, T. Sandholm, "Faster Game Solving via Predictive Blackwell Approachability: Connecting Regret Matching and Mirror Descent." *AAAI-21*

See our full paper on Arxiv!