Evaluation
Online learning
links with optimization and games
Université Paris–Saclay

a hybrid of mirror descent and dual averaging

This project studies an algorithm which shares similarities with mirror descent with step-sizes and dual averaging with parameters.

Let $d \geqslant 0$, $\mathscr{X}$ a nonempty closed convex subset of $\mathbb{R}^d$, H a mirror map compatible with $\mathscr{X}$, $(u_t)_{t\geqslant 0}$ be a sequence in $\mathbb{R}^d$, $(\eta_t)_{t\geqslant 0}$ a positive sequence, and $x_0 \in \mathscr{X} \cap \operatorname{dom} \mathrm{H}$. Then define

$$x_{t+1} = \operatorname*{arg\,max}_{x\in\mathscr{X}} \left\{ \langle \nabla \mathrm{H}(x_t) + \eta_t u_t, x\rangle - \frac{\eta_t}{\eta_{t+1}} \mathrm{H}(x) \right\}, \qquad t \geqslant 0.$$

One special case of interest is when $(\eta_t)_{t\geqslant 0}$ is nonincreasing.

1) Prove that the above can be seen as UMD iterates.

2) Is there cases where the above are MD with step-sizes? MD with parameters?

3) Derive regret bounds.

4) Derive guarantees for Lipschitz convex optimization.

5) For a Lipschitz convex optimization problem, perform numerical experiments to compare the performance of the above algorithm (in the case where H is the Euclidean mirror map) with corresponding MD with nonincreasing step-sizes and DA with nonincreasing parameters. *A possible setting is an SVM with no regularization, constrained in e.g. a Euclidean ball, but this is only a suggestion.*

6) BONUS. — Derive guarantees and perform numerical experiments for (constrained) smooth convex optimization. *Possible settings are least-squares linear regression and logistic regression*.