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

alternative algorithms for stochastic convex optimization

Let $d \geqslant 1$, $\mathscr{X}$ a nonempty closed convex subset of $\mathbb{R}^d$, H a mirror map compatible with $\mathscr{X}$, $h = \mathrm{H} + \mathrm{I}_{\mathscr{X}}$, $(u_t)_{t\geqslant 1}$ a sequence in $\mathbb{R}^d$ and $(\eta_t)_{t\geqslant 1}$ a positive nonincreasing sequence.

Let $(h_t)_{t\in 1+\frac{1}{2}\mathbb{N}}$ be defined as $h_1 = \eta_1^{-1}(h - \min h)$ and

$$h_{t+1/2} = h_{t+1} = \frac{h - \min h}{\eta_{t+1}}, \qquad t \geqslant 1.$$

Let $((x_t, y_t))_{t\geqslant 1}$ be a UMD sequence associated with regularizers $(h_t)_{t\in 1+\frac{1}{2}\mathbb{N}}$ and sequence $(u_t)_{t\geqslant 1}$ such that

$$\forall t \geqslant 1,\ \forall x \in \operatorname{dom} h, \qquad \langle y_{t+1} - y_t - u_t, x - x_{t+1}\rangle \geqslant 0.$$

1) Explicitely give two different sequences $(x_t)_{t\geqslant 1}$ that satisfy the above.

2) Derive regret bounds.

Now consider the convex finite-sum optimization approach described in Example 6.4.1 from the lecture notes.

3) Apply the above sequences to this problem and establish guarantees with assumptions similar to Proposition 6.4.3 from the lecture notes.

4) Consider the case where H is the Euclidean mirror map. Perform numerical experiments to compare the performance of the two algorithms from question 1) and two different algorithms that belong to the class described in Proposition 6.4.3 from the lecture notes. *A possible setting is an SVM with no regularization, constrained in e.g. a Euclidean ball, but this is only a suggestion.*