Machine Learning Theory. Lecture 15.

Wouter M. Koolen

- Two-player zero-sum games
- Nesterov Acceleration from game dynamics

  Acceleration through Optimistic No-Regret Dynamics.

  Wang and Abernethy.

  Advances in Neural Information Processing Systems 31 (NIPS 2018).

Games

Objective function

$$g(x, y)$$

convex in x, concave in y.

The game value is

$$V^* = \inf_x \sup_y g(x, y) = \sup_y \inf_x g(x, y).$$

An ϵ-*saddle point* $(\bar{x}, \bar{y})$ satisfies

$$V^* - \epsilon \leq \inf_x g(x, \bar{y}) \leq V^* \leq \sup_y g(\bar{x}, y) \leq V^* + \epsilon.$$

Question: how to find ϵ-saddle point?

Algorithm

Idea: play regret minimisation algorithms for x and y.

- Players play y_t and x_t.
- Players see loss functions $y \mapsto -g(x_t, y)$ and $x \mapsto +g(x, y_t)$.

Output pair of average iterates: $\left(\frac{1}{T}\sum_{t=1}^{T} x_t, \frac{1}{T}\sum_{t=1}^{T} y_t\right)$.

Saddle point

Assume the players have regret (bounds) R_T^x and R_T^y, i.e.

$$\sum_{t=1}^{T} +g(x_t, y_t) - \inf_x \sum_{t=1}^{T} +g(x, y_t) \ \leq \ R_T^x$$

$$\sum_{t=1}^{T} -g(x_t, y_t) - \inf_y \sum_{t=1}^{T} -g(x_t, y) \ \leq \ R_T^y$$

Claim: $\bar{x}_T = \frac{1}{T} \sum_{t=1}^T x_t$ and $\bar{y}_T = \frac{1}{T} \sum_{t=1}^T y_t$ form an $\frac{R_T^x + R_T^y}{T}$-saddle point.

Analysis

$$
\begin{aligned}
V^* &= \inf_x \sup_y g(x, y) \\
&\le \sup_y g(\bar{x}_T, y) \\
&\le \sup_y \frac{1}{T} \sum_{t=1}^{T} g(x_t, y) \\
&\le \frac{1}{T} \sum_{t=1}^{T} g(x_t, y_t) + \frac{R_T^y}{T} \\
&\le \inf_x \frac{1}{T} \sum_{t=1}^{T} g(x, y_t) + \frac{R_T^x + R_T^y}{T} \\
&\le \inf_x g(x, \bar{y}_T) + \frac{R_T^x + R_T^y}{T} \\
&\le \inf_x \sup_y g(x, y) + \frac{R_T^x + R_T^y}{T} \\
&= V^* + \frac{R_T^x + R_T^y}{T}
\end{aligned}
$$

Offline Optimisation

Starting point: optimisation problem $\inf_x f(x)$.

Regret minimisation algorithm for $\ell_t = f$ gives $O(T^{-1/2})$ suboptimality for average iterate.

Can we do better?

Here we assume that f is L-smooth, i.e.

$$\|\nabla f(u) - \nabla f(v)\| \; \leq \; L\|u - v\|$$

(note: converse to strong convexity).

Fenchel Game

Idea: form *Fenchel game*

$$g(x, y) \ = \ \langle x, y \rangle - f^*(y)$$

where $f^*(y) = \sup_x \langle x, y \rangle - f(x)$ is the Fenchel conjugate.

Crux: saddle point for Fenchel game solves minimisation problem :

$$\inf_x \sup_y g(x, y) \ = \ \inf_x \sup_y \langle x, y \rangle - f^*(y) \ = \ \inf_x f^{**}(x) \ = \ \inf_x f(x).$$

Question

Claim: f smooth iff f^* strongly convex.

Idea: exploit strong convexity in Fenchel game.

We see that the Fenchel game

$$g(x, y) \;=\; \langle x, y\rangle - f^*(y)$$

is *strongly convex* in y and *linear* in x.

Approach 4 main ideas

1. Weighting $\alpha_1, \alpha_2, \ldots$ on rounds
2. Order the players: inner player *reacts* to outer player action.
3. Apply Optimistic Follow-The-Leader for y player
4. Apply Online Gradient Descent for x player.

Weighted rounds

In round t we assign losses $x \mapsto \alpha_t g(x, y_t)$ and $y \mapsto -\alpha_t g(x_t, y)$.

We now analyse the weighted average iterates

$$\bar{x}_T = \frac{1}{A_T} \sum_{t=1}^{T} \alpha_t x_t \qquad\qquad \bar{y}_T = \frac{1}{A_T} \sum_{t=1}^{T} \alpha_t y_t$$

where $A_t = \sum_{s=1}^{t} \alpha_s$.

Approximate Saddle point

Let

$$V^* = \inf_x \sup_y g(x,y) = \inf_x f(x).$$

An ϵ saddle point $(\bar{x}, \bar{y})$ for the Fenchel game satisfies

$$V^* - \epsilon \le \inf_x g(x, \bar{y}) \le V^* \le \sup_y g(\bar{x}, y) \le V^* + \epsilon$$

In particular

$$f(\bar{x}) = \sup_y g(\bar{x}, y) \le \inf_x f(x) + \epsilon.$$

Result for y player

Weighted Optimistic FTL plays

$$y_t = \arg\min_y -\alpha_t g(x_{t-1}, y) + \sum_{s=1}^{t-1} -\alpha_s g(x_s, y)$$

Expanding the Fenchel game, this is

$$y_t = \nabla f(\tilde{x}_t) \quad \text{where} \quad \tilde{x}_t = \frac{\alpha_t x_{t-1} + \sum_{s=1}^{t-1} \alpha_s x_s}{A_t}$$

Theorem 1. *Optimistic FTL satisfies*

$$\sup_y \sum_{t=1}^{T} \alpha_t \left(g(x_t, y) - g(x_t, y_t) \right) \le L \sum_{t=1}^{T} \frac{\alpha_t^2}{A_t} \|x_t - x_{t-1}\|^2.$$

Result for x player

Weighted Online Gradient Descent plays $x_1 = 0$ and

$$x_t \;=\; x_{t-1} - \gamma\alpha_t \nabla_x g(x, y_t)$$

Expanding the Fenchel Game, this is

$$x_t \;=\; x_{t-1} - \gamma\alpha_t y_t$$

Theorem 2. *Let $\|x^*\| \leq D$. Then*

$$\sum_{t=1}^{T} \alpha_t \left(g(x_t, y_t) - g(x^*, y_t)\right) \;\leq\; \frac{D^2}{\gamma} - \sum_{t=1}^{T} \frac{1}{2\gamma} \|x_t - x_{t-1}\|^2.$$

The reason we get a negative regret is that x plays second, with knowledge of y_t.

Combination

In total, we find

$$f(\bar{x}_T) - \min_x f(x) \;\leq\; \frac{1}{A_T}\left(\frac{D^2}{\gamma} + \sum_{t=1}^{T}\left(\frac{\alpha_t^2}{A_t}L - \frac{1}{2\gamma}\right)\|x_t - x_{t-1}\|^2\right).$$

Setting $\frac{\alpha_t^2}{A_t}L = \frac{1}{2\gamma}$, i.e. $\alpha_t = t$ and $\gamma = \frac{1}{4L}$, we obtain

$$\frac{\alpha_t^2}{A_t}L \;=\; \frac{t^2}{t(t+1)/2}L \;\leq\; 2L \;=\; \frac{1}{2\gamma}$$

and hence we have

$$f(\bar{x}_T) - \min_x f(x) \;\leq\; \frac{8LD^2}{T^2}$$

Final Algorithm: Nesterov Acceleration

Initialise $x_1 = 0$.

For $t = 1, \ldots, T$

- $\tilde{x}_t = \frac{\alpha_t x_{t-1} + \sum_{s=1}^{t-1} \alpha_s x_s}{A_t}$
- $y_t = \nabla f(\tilde{x}_t)$
- $x_t = x_{t-1} - \gamma \alpha_t y_t$

Output average iterate

$$\frac{1}{A_T} \sum_{t=1}^{T} \alpha_t x_t$$