Machine Learning Theory 2025
Lecture 14

Wouter M. Koolen

Download these slides now from elo.mastermath.nl!

- Two-player zero-sum games
- Nesterov Acceleration from game dynamics
  Acceleration through Optimistic No-Regret Dynamics.
  Wang and Abernethy.
  Neural Information Processing Systems (2018).

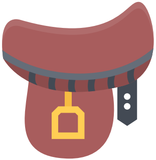

Two-Player Zero-Sum Games

Subject

Two players. One trying to maximise, one trying to minimise.

What happens when they **both behave optimally**?

Example two-player objective functions

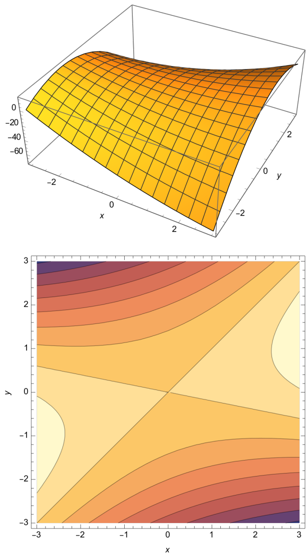

indefinite quadratic

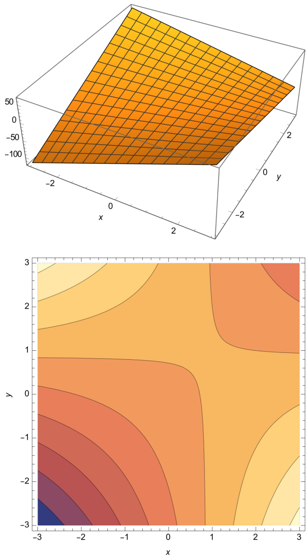

bi-linear

Games

Objective function

$$g(x, y)$$

convex in x, concave in y.

Games

Objective function

$$g(x, y)$$

convex in x, concave in y.

The game **value** is

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

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

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

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?

Motivation for Saddle Point Computation

- Analysing actual two-player situations
  - Economics
  - Security
  - ...
- Robust learning (Generative Adversarial Networks, ...)
- Applications in offline optimisation
  - Acceleration
  - Constraints (Lagrange multipliers, "primal-dual", ...)

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.

$$\begin{aligned} \sum_{t=1}^T +g(x_t, y_t) - \inf_x \sum_{t=1}^T +g(x, y_t) &\leq R_T^x \\ \sum_{t=1}^T -g(x_t, y_t) - \inf_y \sum_{t=1}^T -g(x_t, y) &\leq R_T^y \end{aligned}$$

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

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

Nesterov Acceleration

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 :

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

Approximate Saddle point

Moreover, an approximate saddle point gives an approximate minimiser. Recall that

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

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

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

Extra Assumption: Smoothness

Proposition

f is smooth $\Leftrightarrow$ f^* is strongly convex.

We see that the Fenchel game

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

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

Idea: exploit strong convexity in Fenchel game.

Elements of the Approach

The approach combines 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 scaled by α_t

$$x \;\mapsto\; \alpha_t g(x, y_t) \qquad \text{and} \qquad y \;\mapsto\; -\alpha_t g(x_t, y).$$

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

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

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

$$x_0 = 0 \qquad \text{and} \qquad x_t = x_{t-1} - \gamma\alpha_t \nabla_x g(x, y_t).$$

Expanding the Fenchel Game, this is

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

NB: Iterate x_t for round t defined in terms of outer player move y_t

Theorem

Let $\|x_*\| \leq D$. Then OGD satisfies

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

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

Combination

In total, we find

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

We now tune α_t, γ to ensure $\frac{\alpha_t^2}{A_t}L \le \frac{1}{2\gamma}$. We pick

$$\alpha_t \;=\; t \qquad \text{and} \qquad \gamma \;=\; \frac{1}{4L},$$

for then

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

We conclude

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

Final Algorithm: Nesterov Acceleration

Initialise $x_0 = 0$.
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

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

Conclusion of the Lecture

We saw

- Method to learn saddle point in two-player games
- Reduction of offline smooth convex optimisation to saddle point problem

We obtained a hierarchy for offline optimisation

- Convex: $T^{-1/2}$.
- Strongly convex: T^{-1}.
- Convex and smooth: T^{-2}.

Conclusion of the Course

We saw

- Stochastic and game-theoretic frameworks for learning
- Ways to characterise the complexity of learning problems
- Algorithms and their analysis

Advanced topics that may interest you

- Reinforcement Learning
- Learning in (strategic) multi-agent problems
- Fairness, Accountability, Transparency
- Beyond convexity (NNs, tensor dec.)

Conclusion

This concludes the lectures.

- It has been a pleasure
- Good luck for the exam
- If you have an idea that you want to work on ...