Today

Multiplicative Weight Update Method (part III)

- Approximately solving zero-sum games & LPs

CSS.205.1
Toolkit in TCS
- Lecture #17
(19 Apr '21)
Instructor: Prahladh Harsha

Application of MWUM

- Solving zero-sum games
- Solving LPs.

n- experts

$M^{(t)} = (m_1^{(t)}, \ldots, m_n^{(t)})$

- loss vector.

$p^{(t)} = (p_1^{(t)}, \ldots, p_n^{(t)})$ - prob proportional to weights.

$\ell(t) = \langle M^{(t)}, p^{(t)} \rangle$.

Update Rule:

$$w_i^{(t+1)} \leftarrow w_i^{(t)} (1 - \varepsilon m_i^{(t)})$$

(Assumptions: $m_i^{(t)} \in [-1, 1]$
$\varepsilon \in (0, 1/2]$)

Thm [MWUM].
For any expert $i \in [n]$, the expected loss of the MWUM algorithm

$$L(T) = \sum_{t=1}^{T} \ell(t) = \sum_{t=1}^{T} \langle M^{(t)}, p^{(t)} \rangle$$
$$\leq \sum_{t=1}^{T} m_i^{(t)} + \varepsilon \sum_{t=1}^{T} |m_i^{(t)}| + \frac{\ln n}{\varepsilon}$$

Cor: p - any distribution on experts

$$L(T) \leq \sum_{t=1}^{T} \langle M^{(t)}, p \rangle + \varepsilon \sum_{t=1}^{T} \langle |M^{(t)}|, p \rangle + \frac{\ln n}{\varepsilon}$$

Application I: Approximately Solve Zero-Sum Games.

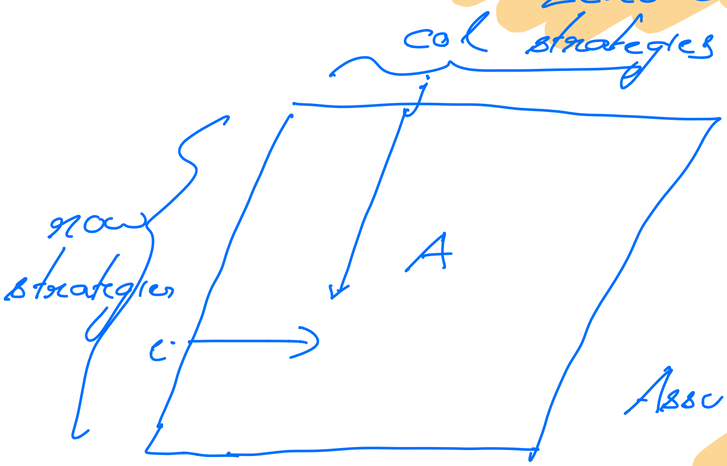

$A(i,j)$ = payoff to the column player

Assume: $A(i,j) \in [0,1]$

Row player's mixed strategy

$p \sim$ dist over pure row strategies

Col plays j

$$A(p,j) \triangleq \mathbb{E}_{i\sim p}[A(i,j)]$$

Col player find j that maximizes $A(p,j)$

$$C(p) = \max_j A(p,j)$$

<u>Similarly</u>

q – Col player's mixed strategy

i – Row player strategy

$$A(i,q) \triangleq \mathbb{E}_{j\sim q}[A(i,j)]$$

$$R(q) = \min_i A(i,q)$$

Weak duality: $R(q) \le C(p)$, $\forall p,q.$

Strong Duality: $\exists\ \lambda^*\ s\ p^*, q^*$

$$R(q) \le R(q^*) = \lambda^* = C(p^*) \le C(p)$$

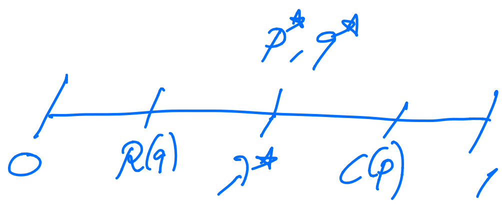

Von Neumann's MiniMax Thm

Last time: Non-constructive proof of existence of λ^*, p^* & q^*

<u>Today</u>: Algorithm using MWUM to compute λ^*, p^*, q^* approximately

Given $\delta \in (0,1)$

Will find $\tilde{p}$, $\tilde{q}$ - Row mixed strategy
Col mixed strategy

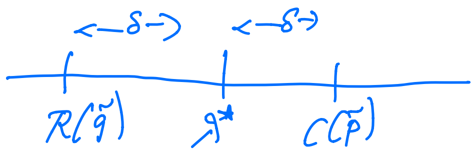

ie, $\lambda^* - \delta \le R(\tilde{q}) \le \lambda^*$

$\lambda^* \le C(\tilde{p}) \le \lambda^* + \delta$ } Additive δ-approximation

<u>Experts</u>: - row pure strategies.

row strategies i

$R(\tilde{q}) = A(i, \tilde{q}) \le \lambda^* - \delta$

At each step of MWUM,
alg o/p a prob dist over n experts:
- p - mixed row strategy.

<u>Oracle</u>: Given any mixed row strategy p, find the best column response.

$$p \longmapsto j \text{ s.t. } C(p) = A(p, j)$$

<u>Cost-vector</u>: Column corresponding to the j^{th} column in the payoff matrix.

$p^{(1)} \leftarrow$ uniform

At any step $t \leftarrow 1..T$

Row player plays acc to prob dist $p^{(t)}$.

$j^{(t)} \triangleq$ Best column response.

$M^{(t)} = j^{(t)}$-column of payoff matrix.

$p^{(t+1)} \leftarrow$ updated appropriately using $p^{(t)}, \varepsilon, M^{(t)}$.

MWUM Theorem.

$$\sum_{t=1}^{T} \langle M^{(t)}, p^{(t)} \rangle \leq (1+\varepsilon) \sum_{t=1}^{T} \langle M^{(t)}, p \rangle + \frac{\ln n}{\varepsilon}$$

for any mixed row strategy p.

Writing this in terms of payoff matrix.

$$\sum_{t=1}^{T} A(p^{(t)}, j^{(t)}) \leq (1+\varepsilon) \sum_{t=1}^{T} A(p, j^{(t)}) + \frac{\ln n}{\varepsilon}$$

$j^{(t)}$ is the best col response for $p^{(t)}$

Hence $A(p^{(t)}, j^{(t)}) \geq \lambda^*$

Hence, (since $A(i,j) \leq 1 \ \forall i,j$)

$$\lambda^* \leq \frac{1}{T} \sum_{t=1}^{T} A(p^{(t)}, j^{(t)}) \leq \frac{1}{T} \sum_{t=1}^{T} A(p, j^{(t)}) + \varepsilon + \frac{\ln n}{\varepsilon T}$$

for any mixed row strategy p.

In particular, it is true for $p = p^*$.

$$\lambda^* \leq \frac{1}{T} \sum A(p^{(t)}, j^{(t)}) \leq \lambda^* + \varepsilon + \frac{\ln n}{\varepsilon T}$$

Set $\varepsilon = \delta/2$; $T = \left\lceil \frac{4 \ln n}{\delta^2} \right\rceil$

$$\lambda^* \leq \frac{1}{T} \sum_{t=1}^{T} A(p^{(t)}, j^{(t)}) \leq \lambda^* + \delta.$$

(1) Approximate Value of Game $\tilde{\lambda}$

$$\tilde{\lambda} \triangleq \frac{1}{T}\sum_{t=1}^{T} A(p^{(t)}, j^{(t)})$$

$$\lambda^* \leq \tilde{\lambda} \leq \lambda^* + \delta.$$

(2) Approximate mixed row strategy $\tilde{p}$:

$\tilde{p} = \frac{1}{T}\sum p^{(t)}$, $\tilde{j}$ - best response to $\tilde{p}$

$$C(\tilde{p}) = A(\tilde{p}, \tilde{j}) = \frac{1}{T}\sum_{t=1}^{T} A(p^{(t)}, \tilde{j})$$

$$\leq \frac{1}{T}\sum_{t=1}^{T} A(p^{(t)}, j^{(t)})$$

($j^{(t)}$ - best response to $p^{(t)}$)

$$\leq \lambda^* + \delta.$$

(3) Approximate mixed Col. strategy $\tilde{q}$:

$$\tilde{q}(j) = \frac{\{t \mid j^{(t)} = j\}}{T}$$

We had proved the following for any p.

$$\lambda^* \leq \frac{1}{T}\sum_{t=1}^{T} A(p^{(t)}, j^{(t)}) \leq \frac{1}{T}\sum_{t=1}^{T} A(p, j^{(t)}) + \varepsilon + \frac{\ln n}{\varepsilon T}$$

Applying it to p - pure strategy i:

$$\lambda^* \le \frac{1}{T}\sum_{t=1}^{T} A(i, j^{(t)}) + \delta$$

$$= A(i, \tilde{q}) + \delta$$

$$A(i, \tilde{q}) \ge \lambda^* - \delta \ , \ \forall i$$

Hence, $R(\tilde{q}) \ge \lambda^* - \delta$.

<u>Thm</u>: For any $\delta \in (0,1)$, δ-additive approximation to the zero-sum game $(\tilde{\lambda}, \tilde{p}, \bar{q})$ can be obtained by making $O\left(\frac{\log n}{\delta^2}\right)$ calls to the oracle $+ O(n)$ processing time for each call.

<u>Application II</u>: Approximately Solving Linear Programs.

LP:

$$\exists? \ \overbrace{x \in \mathcal{P}}^{\text{easy constraint}}, \ \underbrace{Ax \ge b}_{\text{hard constraint}}$$

$\mathcal{P}$ - convex set

Guiding Example:

Set Cover: U - universe

$S_1 \ldots S_m \subseteq U$

Find smallest sub-collection of sets that covers U.

LP relaxation:

$$\min \sum_S x_S$$
$$\sum_{S \ni e} x_S \geq 1 \quad \forall e \in U$$
$$x_S \geq 0$$

x_S - $\mathbb{1}[S \text{ is in sub-collection}]$

$$\mathcal{P} = \{x \in \mathbb{R}^n \mid x_S \geq 0;\ \sum x_S = L\}$$

$$\exists?\ x \in \mathcal{P} \quad \forall e \in U,\ \sum_{S \ni e} x_S \geq 1$$

$$\exists?\ x \in \mathcal{P},\ Ax \geq b \quad \ldots \quad (*)$$

Goal: Approximate Solve above Problem

- Either Find $x \in \mathcal{P}$ s.t for all constraints.

$$A_i x \geq b_i - \delta$$

$$\left\{ \begin{array}{c} \underbrace{Ax \geq b} \\ A_1 x \geq b_1 \\ A_2 x \geq b_2 \\ \vdots \\ A_m x \geq b_m \end{array} \right\}$$

or

- Declare that (*) is infeasible.

MWUM - paradigm

m experts - m constraints.

$p_1 \quad A_1 x \geq b_1$

$p_2 \quad A_2 x \geq b_2$

$p_m \quad A_m x \geq b_m$

$$\sum p_i A_i x \geq \sum p_i b_i$$

$\hookrightarrow$ Single constraint

Assume the existence of an oracle that solves the single constraint problem

---

Set cover

$$\sum_{S \ni e} x_S \geq 1 \quad ; \; e \in U$$

p - prob dist over the elements of universe

$$\sum_{e \in U} p(e) \sum_{S \ni e} x_S \geq \sum_{e \in U} p_e \cdot 1$$

ie $\sum_S x_S \sum_{e \in S} p(e) \geq 1$

$$\sum_S x_S \cdot p(S) \geq 1 \quad \text{where}$$

$p_S :=$ ~~total~~ total prob of all elts in S.

Single Constraint Problem

$$\exists ?\ x \in \mathcal{P}, \quad \sum x_S\, p(S) \geq 1$$

$$\exists x, \quad x_S \geq 0; \ \sum x_S = 1 \ ; \ \sum x_S\, p(S) \geq 1.$$

Easy to solve by setting $x_S = \begin{cases} 1 & \text{if } S = \underset{p(S)}{\operatorname{argmax}} \\ 0 & \text{o.w.} \end{cases}$

<u>Oracle:</u> (l, ρ)-bounded oracle. $(0 \leq l \leq \rho \leq 1)$

There exists an oracle $\mathcal{O}$ &
$I \subseteq [m]$ (subset of constraints) s.t.

on input any p (prob dist over the m constraints), $\mathcal{O}$ outputs

a soln to

$$p^T A x \geq p^T b \quad \text{or declares}$$

it infeasible

Whenever, it outputs a soln x, x satisfie

$$-l \le A_i x - b_i \le \rho \quad , \forall i \in I$$

$$-\rho \le A_i x - b_i \le l \quad , \forall i \notin I$$

Set cover:

$I = [m]$.

$$-1 \le \sum_{S \ni e} x_S - 1 \le L - 1$$

The oracle we constructed above is a $(1, L-1)$-bounded oracle.

$$\left.\begin{aligned} A_1 x &\ge b_1 \\ A_2 x &\ge b_2 \\ &\vdots \\ A_m x &\ge b_m \end{aligned}\right\} \text{ m-experts}$$

Cost of i^{th} expert $= \dfrac{[A_i x - b_i]}{\rho}$