Today

Experts/Zero-Sum Games Equilibrium.

Boosting and Experts.

Routing and Experts.

Best Response

Column goes first:
Find y, where best row is not too low..

$$R = \max_y \min_x (x^t A y).$$

Note: x can be $(0,0,\ldots,1,\ldots 0)$.

Example: Roshambo. Value of R?

Row goes first:
Find x, where best column is not high.

$$C = \min_x \max_y (x^t A y).$$

Agin: y of form $(0,0,\ldots,1,\ldots 0)$.

Example: Roshambo. Value of C?

Two person zero sum games.

$m \times n$ payoff matrix A.

Row mixed strategy: $x = (x_1,\ldots,x_m)$.
Column mixed strategy: $y = (y_1,\ldots,y_n)$.

Payoff for strategy pair (x,y):

$$p(x,y) = x^t A y$$

That is,

$$\sum_i x_i \left(\sum_j a_{i,j} y_j\right) = \sum_j \left(\sum_i x_i a_{i,j}\right) y_j.$$

Recall row minimizes, column maximizes.

Equilibrium pair: (x^*,y^*)?

$$(x^*)^t A y^* = \max_y (x^*)^t A y = \min_x x^t A y^*.$$

(No better column strategy, no better row strategy.)

Duality.

$$R = \max_y \min_x (x^t A y).$$
$$C = \min_x \max_y (x^t A y).$$

Weak Duality: $R \leq C$.
Proof: Better to go second. □

At Equilibrium (x^*,y^*), payoff v:
row payoffs (Ay^*) all $\geq v \implies R \geq v$.
column payoffs $((x^*)^t A)$ all $\leq v \implies v \geq C$.
$\implies R \geq C$

Equilibrium $\implies R = C$!

Strong Duality: There is an equilibrium point! and $R = C$!

Doesn't matter who plays first!

Equilibrium.

Equilibrium pair: (x^*,y^*)?

$$p(x,y) = (x^*)^t A y^* = \max_y (x^*)^t A y = \min_x x^t A y^*.$$

(No better column strategy, no better row strategy.)

No row is better:
$\min_i A^{(i)} \cdot y = (x^*)^t A y^*$. [1]

No column is better:
$\max_j (A^t)^{(j)} \cdot x = (x^*)^t A y^*$.

[1] $A^{(i)}$ is ith row.

Proof of Equilibrium.

Later. Still later...

Aproximate equilibrium ...

$C(x) = \max_y x^t A y$

$R(y) = \min_x x^t A y$

Always: $R(y) \leq C(x)$

Strategy pair: (x,y)

Equilibrium: (x,y)

$R(y) = C(x) \to C(x) - R(y) = 0$.

Approximate Equilibrium: $C(x) - R(y) \leq \varepsilon$.

With $R(y) \leq C(x)$
→ "Response y to x is within ε of best response"
→ "Response x to y is within ε of best response"

Proof of approximate equilibrium.

How?

(A) Using geometry.
(B) Using a fixed point theorem.
(C) Using multiplicative weights.
(D) By the skin of my teeth.

(C) ..and (D).
Not hard. Even easy. Still, head scratching happens.

Games and experts

Again: find $(x*, y*)$, such that

$(\max_y x^* Ay) - (\min_x x^* Ay^*) \leq \varepsilon$

$C(x^*) \quad - \quad R(y^*) \leq \varepsilon$

Experts Framework: n Experts, T days, L^* -total loss.

Multiplicative Weights Method yields loss L where

$L \leq (1+\varepsilon)L^* + \frac{\log n}{\varepsilon}$

Games and Experts.

Assume: A has payoffs in $[0,1]$.

For $T = \frac{\log n}{\varepsilon^2}$ days:

1) m pure row strategies are experts.
Use multiplicative weights, produce row distribution.
Let x_t be distribution (row strategy) x_t on day t.

2) Each day, adversary plays best column response to x_t.
Choose column of A that maximizes row's expected loss.
Let y_t be indicator vector for this column.

Let $y^* = \frac{1}{T}\sum_t y_t$ and $x^* = \text{argmin}_{x_t} x_t A y_t$.

Claim: $(x^*, y)^*$ are 2ε-optimal for matrix A.

Proof Idea:
x_t minimizes the best column response is chosen. Clearly good for row.
column best response is at least what it is against x_t. Total loss, L is at least column payoff. Best row payoff, L^* is roughly less than L due to MW anlysis.
Combine bounds. Done!

Approximate Equilibrium!

Experts: x_t is strategy on day t, y_t is best column against x_t.

Let $y^* = \frac{1}{T}\sum_t y_t$ and $x^* = \text{argmin}_{x_t} x_t A y_t$.

Claim: $(x^*, y)^*$ are 2ε-optimal for matrix A.

Column payoff: $C(x^*) = \max_y x^* Ay$.
Loss on day t, $x_t A y_t \geq C(x^*)$ by the choice of x .
Thus, algorithm loss, L, is $\geq TC(x^*)$.

Best expert: L^*- best row against all the columns played.

best row against $\sum_t Ay_t$ and $Ty^* = \sum_t y_t$
$\rightarrow$ best row against TAy^*.
$\rightarrow L^* \leq TR(y^*)$.

Multiplicative Weights: $L \leq (1+\varepsilon)L^* + \frac{\ln n}{\varepsilon}$

$TC(x^*) \leq (1+\varepsilon)TR(y^*) + \frac{\ln n}{\varepsilon} \rightarrow C(x^*) \leq (1+\varepsilon)R(y^*) + \frac{\ln n}{\varepsilon T}$
$\rightarrow C(x^*) - R(y^*) \leq \varepsilon R(y^*) + \frac{\ln n}{\varepsilon T}$.

$T = \frac{\ln n}{\varepsilon^2}$, $R(y^*) \leq 1$
$\rightarrow C(x^*) - R(y^*) \leq 2\varepsilon$.

Approximate Equilibrium: notes!

Experts: x_t is strategy on day t, y_t is best column against x_t.

Let $x^* = \frac{1}{T}\sum_t x_t$ and $y^* = \frac{1}{T}\sum_t y_t$.

Claim: $(x^*, y)^*$ are 2ε-optimal for matrix A.

Column payoff: $C(x^*) = \max_y x^* Ay$.
Let y_r be best response to $C(x^*)$.
Day t,y_t best response to $x_t \rightarrow x_t A y_t \geq x_t A y_r$.
Algorithm loss: $\sum_t x_t A y_t \geq \sum_t x_t A y_r$
$L \geq TC(x^*)$.

Best expert: L^*- best row against all the columns played.

best row against $\sum_t Ay_t$ and $Ty^* = \sum_t y_t$
$\rightarrow$ best row against TAy^*.
$\rightarrow L^* \leq TR(y^*)$.

Multiplicative Weights: $L \leq (1+\varepsilon)L^* + \frac{\ln n}{\varepsilon}$

$TC(x^*) \leq (1+\varepsilon)TR(y^*) + \frac{\ln n}{\varepsilon} \rightarrow C(x^*) \leq (1+\varepsilon)R(y^*) + \frac{\ln n}{\varepsilon T}$
$\rightarrow C(x^*) - R(y^*) \leq \varepsilon R(y^*) + \frac{\ln n}{\varepsilon T}$.

$T = \frac{\ln n}{\varepsilon^2}$, $R(y^*) \leq 1 \rightarrow C(x^*) - R(y^*) \leq 2\varepsilon$.

Comments

For any ε, there exists an ε-Approximate Equilibrium.

Does an equilibrium exist? Yes.

Something about math here? Fixed point theorem.

Later: will use geometry, linear programming.

Complexity?
$T = \frac{\ln n}{\varepsilon^2} \rightarrow O(nm\frac{\log n}{\varepsilon^2})$. Basically linear!

Versus Linear Programming: $O(n^3 m)$ Basically quadratic.
(Faster linear programming: $O(\sqrt{n+m})$ linear solution solves.)
Still much slower ... and more complicated.

Dynamics: best response, update weight, best response.

Also works with both using multiplicative weights.

"In practice."

Learning.

Learning just a bit.

Example: set of labelled points, find hyperplane that separates.

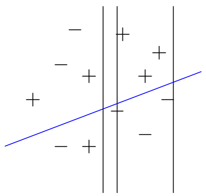

Looks hard.

Get $1/2$ on correct side? Easy.
Arbitrary line. And Scan.

Useless. A bit more than $1/2$

Weak Learner: Classify $\geq \frac{1}{2}+\varepsilon$ points correctly.

Not really important but ...

Weak Learner/Strong Learner

Input: n labelled points.

Weak Learner:
produce hypothesis correctly classifies $\frac{1}{2}+\varepsilon$ fraction

Strong Learner:
produce hypothesis correctly classifies $1+\mu$ fraction
That's a really strong learner!
produce hypothesis correctly classifies $1-\mu$ fraction

Same thing?

Can one use weak learning to produce strong learner?

Boosting: use a weak learner to produce strong learner.

Poll.

Given a weak learning method (produce ok hypotheses.)
produce a great hypothesis.

Can we do this?

(A) Yes
(B) No

If yes. How?

Multiplicative Weights!

The endpoint to a line of research.

Experts Picture

Boosting/MW Framework

Experts are points. “Adversary” weak learner.

Points want to be misclassified.

Learner wants to maximize probability
of classifying random point correctly.
Strong learner algorithm will come from adversary.

Do $T = \frac{2}{\gamma^2}\log\frac{1}{\mu}$ rounds

1. Row player: multiplicative weights($1-\gamma$) on points.
2. Column: run weak learner on row distribution.
3. Hypothesis $h(x)$: majority of $h_1(x), h_2(x), \ldots, h_T(x)$.

Claim: $h(x)$ is correct on $1-\mu$ of the points ! ! !

Cool!

Really? Proof?

Some intuition

Intuition 1: Each point classified correctly independently in each round with probability $\frac{1}{2}+\varepsilon$.

After enough rounds, majority rule correct for almost all points.

Intuition 2:
Say some point classified correctly $\leq 1/2$ of time.

High probability of choosing such point in distribuiont.

In limit, whole distribution becomes such point.

This subset will be classified correctly with probability $1/2+\varepsilon$.

Adaboost proof.

Claim: $h(x)$ is correct on $1-\mu$ of the points ! !

Let S_{bad} be the set of points where $h(x)$ is incorrect.

majority of $h_t(x)$ are wrong for $x \in S_{bad}$.

$x \in S_{bad}$ is a good expert – loses less than $\frac{1}{2}$ the time.

$W(T) \geq (1-\varepsilon)^{\frac{T}{2}}|S_{bad}|$

Each day, weak learner gets $\geq \frac{1}{2}+\gamma$ payoff.

$\rightarrow L_t \geq \frac{1}{2}+\gamma$.

$\rightarrow W(T) \leq n(1-\varepsilon)^L \leq ne^{-\varepsilon L} \leq ne^{-\varepsilon(\frac{1}{2}+\gamma)T}$

Combining

$|S_{bad}|(1-\varepsilon)^{T/2} \leq W(T) \leq ne^{\varepsilon(\frac{1}{2}+\gamma T)}$

Calculation..

$|S_{bad}|(1-\varepsilon)^{T/2} \leq ne^{\varepsilon(\frac{1}{2}+\gamma)}T$

Set $\varepsilon = \gamma$, take logs.

$\ln\left(\frac{|S_{bad}|}{n}\right) + \frac{T}{2}\ln(1-\gamma) \leq -\gamma T(\frac{1}{2}+\gamma)$

Again, $-\gamma-\gamma^2 \leq \ln(1-\gamma)$,

$\ln\left(\frac{|S_{bad}|}{n}\right) + \frac{T}{2}(-\gamma-\gamma^2) \leq -\gamma T(\frac{1}{2}+\gamma) \rightarrow \ln\left(\frac{|S_{bad}|}{n}\right) \leq -\frac{\gamma^2 T}{2}$

And $T = \frac{2}{\gamma^2}\log\frac{1}{\mu}$,

$\rightarrow \ln\left(\frac{|S_{bad}|}{n}\right) \leq \log\mu \rightarrow \frac{|S_{bad}|}{n} \leq \mu$.

The misclassified set is at most μ fraction of all the points.

The hypothesis correctly classifies $1-\mu$ of the points ! ! !

Claim: Multiplicative weights: $h(x)$ is correct on $1-\mu$ of the points !

Claim: Weak learning $\rightarrow$ strong learning!

not so weak after all.

Some details...

Weak learner learns over distributions of points not points.

Make copies of points to simulate distributions.

Used often in machine learning.

Example.

Set of points on unit ball in d-space.

Learner: learns hyperplanes through origin.

Can learn if

there is a hyperplane, $\mathscr{H}$, that separates all the points.

and find $\frac{1}{2}+\varepsilon$ weighted separating plane.

Experts output is average of hyperplanes ...a hyperplane!

$\frac{1}{2}+\varepsilon$ separating hyperplane?
Assumption: margin γ.

Random hyperplane?
Not likely to be exactly normal to $\mathscr{H}$.
Should get $\frac{1}{2}+\gamma/\sqrt{d}$

$O(\frac{d\log n}{\gamma^2})$ to find separating hyperplane.

Weak learner: random Wow. That's weak.

Better weak learner?

Hyperplane that separates weighted average of +/- points?

Change loss a bit, and get better results.

Toll/Congestion

Given: $G = (V,E)$.
Given $(s_1,t_1)\ldots(s_k,t_k)$.
Row: choose routing of all paths.
Column: choose edge.
Row pays if column chooses edge on any path.

Matrix:
row for each routing: r
column for each edge: e

$A[r,e]$ is congestion on edge e by routing r

Offense: (Best Response.)
Router: route along shortest paths.
Toll: charge most loaded edge.

Defense: Toll: maximize shortest path under tolls.
Route: minimize max congestion on any edge.

Two person game.

Row for every roting. ($A[r,e]$)

An exponential number of rows!

Two person game with experts won't be so easy to implement.

Version with row and column flipped may work.

$A[e,r]$ - congestion of edge e on routing r.

m rows. Exponential number of columns.

Multiplicative Weights only maintains m weights.

Adversary only needs to provide best column each day.

Runtime only dependent on m and T (number of days.)

Fractional versus Integer.

Did we solve path routing?
Yes? No?

No! Average of T routings.
We approximately solved fractional routing problem.

No solution to the path routing problem that is $(1+\varepsilon)$ optimal!

Homework 2. Problem 1.

Decent solution to path routing problem?

Congestion minimization and Experts.

Will use gain and $[0,\rho]$ version of experts:

$G \geq (1-\varepsilon)G^* - \frac{\rho \log n}{\varepsilon}$.

Let $T = \frac{k \log n}{\varepsilon^2}$

1. Row player runs multiplicative weights:
   $w_i = w_i(1+\varepsilon)^{g_i/k}$.
2. Route all paths along shortest paths.
3. Output the average of all routings: $\frac{1}{T}\sum_t f(t)$.

Claim: The congestion, c_{max} is at most $(1+\varepsilon)C^* + \varepsilon/(1-\varepsilon)$.

Proof:
$G \geq G^*(1-\varepsilon) - \frac{k \log n}{\varepsilon}$
$G^* = c_{\max} T$ — Best row payoff against average routing.

$G \leq C^* T$ — each day, gain is average congestion $\leq C^*$
since each day cost is toll solution which is at most C^*
$C^* T \geq c_{max} T(1-\varepsilon) - \frac{k \log n}{\varepsilon}$
For $T = \frac{k \log n}{\varepsilon^2}$
$\rightarrow C^* \frac{1}{1-\varepsilon} + \varepsilon \geq c_{\max}$ plus $\frac{1}{1-\varepsilon} \leq 1+\varepsilon \rightarrow$
$c_{max} - C^* \leq \varepsilon C + \varepsilon/(1-\varepsilon)$ □

Randomized Rounding

For each s_i, t_i, choose path p_i with probability $f(p_i)$.

Congestion $c(e)$ edge rounds to $\tilde{c}(e)$.

Edge e.
used by paths $p_1, \ldots, p_m$.
Let $X_i = 1$,
if path p_i is chosen.
otherwise, $X_i = 0$.

Rounded congestion, $\tilde{c}(e)$, is $\sum_i X_i$.

Expected Congestion: $\sum_i E(X_i)$.
$E(X_i) = 1Pr[X_i = 1] + 0Pr[X_i = 0] = f(p_i)$
$\rightarrow \sum_i E(X_i) = \sum_i f(p_i) = c(e)$.
$\rightarrow E(\tilde{c}(e)) = c(e)$.

Concentration (law of large numbers)
$c(e)$ is relatively large ($\Omega(\log n)$)
$\rightarrow \tilde{c}(e) \approx c(e)$.

Concentration results? later.

Better setup.

Runtime: $O(km)$ to route in each step.
$O(k \log n(\frac{1}{\varepsilon^2}))$ steps
$\rightarrow O(k^2 m \log n)$ to get a constant approximation.

Homework: $O(km \log n)$ algorithm.

See you on Tuesday.