Today

Other algorithms.

For linear programming.

Online.

Perceptron Guarantees.

Separable set of points.

Perceptron.
Prove a performance bound.

Margin and Perceptron

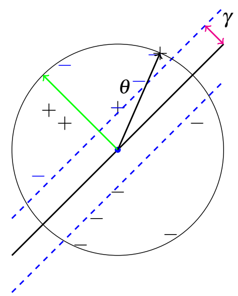

Labelled points with $x_1,\ldots,x_n$.

Hyperplane separator.

Margins.

Inside unit ball.
Margin γ
Hyperplane:
$w \cdot x \geq \gamma$ for $+$ points.
$w \cdot x \leq -\gamma$ for $-$ points.
Put points on unit ball.
$w \cdot x = cos\theta$
Will assume positive labels!
negate the negative:
$(x,-1) \to (-x,1)$

Perceptron Algorithm

An aside: a hyperplane is a perceptron.
(single layer neural network, do you see? Linear programming!)

Alg: Given $x_1,\ldots,x_n$.

Let $w_1 = x_1$.
For each x_i where $w_t \cdot x_i$ has wrong sign (negative) Mistake
$w_{t+1} = w_t + x_i$
$t = t+1$

Theorem: Algorithm only makes $\frac{1}{\gamma^2}$ mistakes.
Idea: Mistake on positive x_i:
$w_{t+1} \cdot x_i = (w_t + x_i) \cdot x_i = w_t x_i + 1$.
A step in the right direction! $w_{t+1} \cdot x_i$ is bigger.

Claim 1: $w_{t+1} \cdot w \geq w_t \cdot w + \gamma$.
A γ *in the right direction!* w_{t+1} is more like w.

Mistake on positive x_i;
$w_{t+1} \cdot w = (w_t + x_i) \cdot w = w_t \cdot w + x_i \cdot w$
$\geq w_t \cdot w + \gamma$.

□

Proof:continued.

Alg: Given $x_1,\ldots,x_n$.

Let $w_1 = x_1$.
For each x_i where $w_t \cdot x_i$ has wrong sign (negative) Mistake
$w_{t+1} = w_t + x_i$
$t = t+1$

Claim 2: $|w_{t+1}|^2 \leq |w_t|^2 + 1$

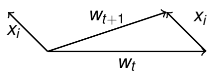

$w_{t+1} = w_t + x_i$
Less than a right angle!
$\to |w_{t+1}|^2 \leq |w_t|^2 + |x_i|^2 \leq |w_t|^2 + 1$.
Algebraically.
Positive x_i, $w_t \cdot x_i \leq 0$.
$(w_t + x_i)^2 = |w_t|^2 + 2w_t \cdot x_i + |x_i|^2$.
$\leq |w_t|^2 + |x_i|^2 = |w_t|^2 + 1$.

□

Claim 2 holds even if no separating hyperplane!

Putting it together...

Claim 1: $w_{t+1} \cdot w \geq w_t \cdot w + \gamma$. $\implies$ $w_t \cdot w \geq t\gamma$

Claim 2: $|w_{t+1}|^2 \leq |w_t|^2 + 1$. $\implies$ $|w_t|^2 \leq t$

M-number of mistakes in algorithm.
Let $t = M$.

$\gamma M \leq w_M \cdot w$
$\leq ||w_M|| \leq \sqrt{M}$.

$\to M \leq \frac{1}{\gamma^2}$

Finding fat separator.

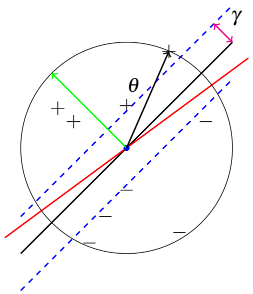

There is a γ separating hyperplane.

We might find the one.

May have bad margin.

Does perceptron find big margin separator.

Approximately Maximizing Margin Algorithm

There is a γ separating hyperplane.

Any point within $\gamma/2$ is still a mistake.

Let $w_1 = x_1$,

For each $x_2, \ldots x_n$,
if $w_t \cdot x_i < \gamma/2$, $w_{t+1} = w_t + x_i$, $t = t+1$

Claim 1: $w_{t+1} \cdot w \geq w_t \cdot w + \gamma$.

Same (ish) as before.

Margin Approximation: Claim 2

Claim 2(?): $|w_{t+1}|^2 \leq |w_t|^2 + 1$??

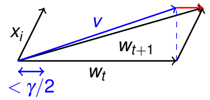

Adding x_i to w_t even if in correct direction.

Obtuse triangle.

$|v|^2 \leq |w_t|^2 + 1$

$\rightarrow |v| \leq |w_t| + \frac{1}{2|w_t|}$

(square right hand side.)

Red bit is at most $\gamma/2$.

Together: $|w_{t+1}| \leq |w_t| + \frac{1}{2|w_t|} + \frac{\gamma}{2}$

If $|w_t| \geq \frac{2}{\gamma}$, then $|w_{t+1}| \leq |w_t| + \frac{3}{4}\gamma$.

M updates $|w_M| \leq \frac{2}{\gamma} + \frac{3}{4}\gamma M$.

Claim 1: Implies $|w_M| \geq \gamma M$.

$\gamma M \leq \frac{2}{\gamma} + \frac{3}{4}\gamma M \rightarrow M \leq \frac{8}{\gamma^2}$

The multiplicative weights framework.

Experts framework.

n experts.

Every day, each offers a prediction.

"Rain" or "Shine."

	Day 1	Day 2	Day 3	$\cdots$	Day T
Expert 1	Shine	Rain	Shine	$\cdots$	
Expert 2	Shine	Shine	Shine	$\cdots$	
Expert 3	Rain	Rain	Rain	$\cdots$	
$\vdots$	$\vdots$	$\vdots$	Shine	$\cdots$	

Rained! Shined! Shined! $\cdots$

Whose advice do you follow?

"The one who is correct most often."

Sort of.

How well do you do?

Infallible expert.

One of the experts is infallible!

Your strategy?

Choose any expert that has not made a mistake!

How long to find perfect expert?

Maybe..never! Never see a mistake.

Better model?

How many mistakes could you make? Mistake Bound.

(A) 1
(B) 2
(C) $\log n$
(D) $n-1$

Adversary designs setup to watch who you choose, and make that expert make a mistake.

$n-1$!

Concept Alert.

Note.

Adversary:
makes you want to look bad.
"You could have done so well"...
but you didn't! ha..ha!

Analysis of Algorithms: do as well as possible!

Back to mistake bound.

Infallible Experts.

Alg: Choose one of the perfect experts.

Mistake Bound: $n-1$
Lower bound: adversary argument.
Upper bound: every mistake finds fallible expert.

Better Algorithm?

Making decision, not trying to find expert!

Algorithm: Go with the majority of previously correct experts.

What you would do anyway!

Alg 2: find majority of the perfect

How many mistakes could you make?

(A) 1
(B) 2
(C) $\log n$
(D) $n-1$

At most $\log n$!

When alg makes a *mistake*,
|"perfect" experts| drops by a factor of two.

Initially n perfect experts
mistake $\rightarrow$ $\leq n/2$ perfect experts
mistake $\rightarrow$ $\leq n/4$ perfect experts
$\vdots$
mistake $\rightarrow$ ≤ 1 perfect expert

≥ 1 perfect expert $\rightarrow$ at most $\log n$ mistakes!

Imperfect Experts

Goal?

Do as well as the best expert!

Algorithm. Suggestions?

Go with majority?

Penalize inaccurate experts?

Best expert is penalized the least.

1. Initially: $w_i = 1$.
2. Predict with weighted majority of experts.
3. $w_i \rightarrow w_i/2$ if wrong.

Analysis: weighted majority

1. Initially: $w_i = 1$.
2. Predict with weighted majority of experts.
3. $w_i \rightarrow w_i/2$ if wrong.

Goal: Best expert makes m mistakes.

Potential function: $\sum_i w_i$. Initially n.

For best expert, b, $w_b \geq \frac{1}{2^m}$.

Each mistake:
total weight of incorrect experts reduced by
-1? -2? factor of $\frac{1}{2}$?
each incorrect expert weight multiplied by $\frac{1}{2}$!
total weight decreases by
factor of $\frac{1}{2}$? factor of $\frac{3}{4}$?
mistake $\rightarrow$ $\geq$ half weight with incorrect experts
($\geq \frac{1}{2}$ total.

Mistake $\rightarrow$ potential function decreased by $\frac{3}{4}$.
$\implies$ for M is number of mistakes that:

$$\frac{1}{2^m} \leq \sum_i w_i \leq \left(\frac{3}{4}\right)^M n.$$

Analysis: continued.

$\frac{1}{2^m} \leq \sum_i w_i \leq \left(\frac{3}{4}\right)^M n.$

m - best expert mistakes M algorithm mistakes.

$\frac{1}{2^m} \leq \left(\frac{3}{4}\right)^M n.$

Take log of both sides.

$-m \leq -M\log(4/3) + \log n.$

Solve for M.
$M \leq (m + \log n)/\log(4/3) \leq 2.4(m + \log n)$

Multiple by $1-\varepsilon$ for incorrect experts...

$(1-\varepsilon)^m \leq \left(1-\frac{\varepsilon}{2}\right)^M n.$

Massage...

$M \leq 2(1+\varepsilon)m + \frac{2\ln n}{\varepsilon}$

Approaches a factor of two of best expert performance!

Best Analysis?

Consider two experts: A,B

Bad example?

Which is worse?

(A) A correct even days, B correct odd days

(B) A correct first half of days, B correct second

Best expert peformance: $T/2$ mistakes.

Pattern (A): $T-1$ mistakes.

Factor of (almost) two worse!

Randomization!!!!

Better approach?

Use?

Randomization!

That is, choose expert i with prob $\propto w_i$

Bad example: A,B,A,B,A...

After a bit, A and B make nearly the same number of mistakes.

Choose each with approximately the same probabilty.

Make a mistake around $1/2$ of the time.

Best expert makes $T/2$ mistakes.

Roughly optimal!

Randomized analysis.

Some formulas:

For $\varepsilon \le 1, x \in [0,1]$,

$(1+\varepsilon)^x \le (1+\varepsilon x)$
$(1-\varepsilon)^x \le (1-\varepsilon x)$

For $\varepsilon \in [0, \frac{1}{2}]$,

$-\varepsilon - \varepsilon^2 \le \ln(1-\varepsilon) \le -\varepsilon$
$\varepsilon - \varepsilon^2 \le \ln(1+\varepsilon) \le \varepsilon$

Proof Idea: $\ln(1+x) = x - \frac{x^2}{2} + \frac{x^3}{3} - \cdots$

Randomized algorithm

Expert i loses $\ell_i^t \in [0,1]$ in round t.

1. Initially $w_i = 1$ for expert i.
2. Choose expert i with prob $\frac{w_i}{W}$, $W = \sum_i w_i$.
3. $w_i \leftarrow w_i(1-\varepsilon)^{\ell_i^t}$

$W(t)$ sum of w_i at time t. $W(0) = n$

Best expert, b, loses L^* total. $\rightarrow W(T) \ge w_b \ge (1-\varepsilon)^{L^*}$.

$L_t = \sum_i \frac{w_i \ell_i^t}{W}$ expected loss of alg. in time t.

Claim: $W(t+1) \le W(t)(1-\varepsilon L_t)$ Loss $\rightarrow$ weight loss.

Proof:

$$\begin{aligned} W(t+1) = \sum_i (1-\varepsilon)^{\ell_i^t} w_i \le \sum_i (1-\varepsilon \ell_i^t) w_i &= \sum_i w_i - \varepsilon \sum_i w_i \ell_i^t \\ &= \sum_i w_i \left(1 - \varepsilon \frac{\sum_i w_i \ell_i^t}{\sum_i w_i}\right) \\ &= W(t)(1-\varepsilon L_t) \end{aligned}$$

Analysis

$(1-\varepsilon)^{L^*} \le W(T) \le n \ \prod_t (1-\varepsilon L_t)$

Take logs

$(L^*)\ln(1-\varepsilon) \le \ln n + \sum \ln(1-\varepsilon L_t)$

Use $-\varepsilon - \varepsilon^2 \le \ln(1-\varepsilon) \le -\varepsilon$

$-(L^*)(\varepsilon + \varepsilon^2) \le \ln n - \varepsilon \sum L_t$

And

$\sum_t L_t \le (1+\varepsilon)L^* + \frac{\ln n}{\varepsilon}$.

$\sum_t L_t$ is total expected loss of algorithm.

Within $(1+\varepsilon)$ ish of the best expert!

No factor of 2 loss!

Gains.

Why so negative?

Each day, each expert gives gain in $[0,1]$.

Multiplicative weights with $(1+\varepsilon)^{g_i^t}$.

$$G \ge (1-\varepsilon)G^* - \frac{\log n}{\varepsilon}$$

where G^* is payoff of best expert.

Scaling:

Not $[0,1]$, say $[0,\rho]$.

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

Summary: multiplicative weights.

Framework: n experts, each loses different amount every day.

Perfect Expert: $\log n$ mistakes.

Imperfect Expert: best makes m mistakes.

Deterministic Strategy: $2(1+\varepsilon)m + \frac{\log n}{\varepsilon}$

Real numbered losses: Best loses L^* total.

Randomized Strategy: $(1+\varepsilon)L^* + \frac{\log n}{\varepsilon}$

Strategy:
  Choose proportional to weights
    multiply weight by $(1-\varepsilon)^{\text{loss}}$.

Multiplicative weights framework!

Applications next!

Strategic Games.

N players.

Each player has strategy set. $\{S_1, \ldots, S_N\}$.

Vector valued payoff function: $u(s_1, \ldots, s_n)$ (e.g., $\in \Re^N$).

Example:

2 players

Player 1: { **D**efect, **C**ooperate }.
Player 2: { **D**efect, **C**ooperate }.

Payoff:

	C	**D**
C	(3,3)	(0,5)
D	(5,0)	(1,1)

Famous because?

	C	**D**
C	(3,3)	(0,5)
D	(5,0)	(.1.1)

What is the best thing for the players to do?

Both cooperate. Payoff $(3,3)$.

If player 1 wants to do better, what does he do?

Defects! Payoff $(5,0)$

What does player 2 do now?

Defects! Payoff $(.1,.1)$.

Stable now!

Nash Equilibrium: neither player has incentive to change strategy.

Digression..

What situations?

Prisoner's dilemma:
Two prisoners separated by jailors and asked to betray partner.

Basis of the free market.
Companies compete, don't cooperate.
No Monopoly:
E.G., OPEC, Airlines, .
Should defect.
Why don't they?
Free market economics ...not so much?
More sophisticated models ,e.g, iterated dominance, coalitions, complexity..
Lots of interesting Game Theory!

This class(today): simpler version.

Two Person Zero Sum Games

2 players.

Each player has strategy set:
m strategies for player 1 n strategies for player 2

Payoff function: $u(i,j) = (-a, a)$ (or just a).
"Player 1 pays a to player 2."

Zero Sum: Payoff for any pair of strategies sums to 0.

Payoffs by m by n matrix: A.
Row player minimizes, column player maximizes.

Roshambo: rock,paper, scissors.

	R	P	S
R	0	1	-1
P	-1	0	1
S	1	-1	0

Any Nash Equilibrium?

(R,R)? no. (R,P)? no. (R,S)? no.

Mixed Strategies.

		R	P	S
		.33	.33	.33
R	.33	0	1	-1
P	.33	-1	0	1
S	.33	1	-1	0

How do you play?

Player 1: play each strategy with equal probability.
Player 2: play each strategy with equal probability.

Definitions.

Mixed strategies: Each player plays distribution over strategies.

Pure strategies: Each player plays single strategy.

Payoffs: Equilibrium.

		R	P	S
		.33	.33	.33
R	.33	0	1	-1
P	.33	-1	0	1
S	.33	1	-1	0

Payoffs? Can't just look it up in matrix!.

Average Payoff. Expected Payoff.

Sample space: $\Omega = \{(i,j) : i,j \in [1,..,3]\}$
Random variable X (payoff).

$$E[X] = \sum_{(i,j)} X(i,j) Pr[(i,j)].$$

Each player chooses independently:
$Pr[(i,j)] = \frac{1}{3} \times \frac{1}{3} = \frac{1}{9}$.

$$E[X] = 0.$$ [1]

[1] Remember zero sum games have one payoff.

Equilibrium

		R	P	S
		.33	.33	.33
R	.33	0	1	-1
P	.33	-1	0	1
S	.33	1	-1	0

Will Player 1 change strategy? Mixed strategies uncountable!

Expected payoffs for pure strategies for player 1.

Expected payoff of Rock? $\frac{1}{3} \times 0 + \frac{1}{3} \times 1 + \frac{1}{3} \times -1 = 0$.
Expected payoff of Paper? $\frac{1}{3} \times -1 + \frac{1}{3} \times 0 + \frac{1}{3} \times 1 = 0$.
Expected payoff of Scissors? $\frac{1}{3} \times 1 + \frac{1}{3} \times -1 + \frac{1}{3} \times 0 = 0$.

No better pure strategy. $\implies$ No better mixed strategy!

Mixed strat. payoff is weighted av. of payoffs of pure strats.

$E[X] = \sum_{(i,j)} (Pr[i] \times Pr[j]) X(i,j) = \sum_i Pr[i] (\sum_j Pr[j] \times X(i,j))$

Mixed strategy can't be better than the best pure strategy.

Player 1 has no incentive to change! Same for player 2.

Equilibrium!

Another example plus notation.

Rock, Paper, Scissors, prEempt.
PreEmpt ties preEmpt, beats everything else.
Payoffs.

	R	P	S	E
R	0	1	-1	1
P	-1	0	1	1
S	1	-1	0	1
E	-1	-1	-1	0

Equilibrium? **(E,E)**. Pure strategy equilibrium.
Notation: Rock is 1, Paper is 2, Scissors is 3, prEmpt is 4.
Payoff Matrix.

$$A = \begin{bmatrix} 0 & 1 & -1 & 1 \\ -1 & 0 & 1 & 1 \\ 1 & -1 & 0 & 1 \\ -1 & -1 & -1 & 0 \end{bmatrix}$$

Playing the boss...

Row has extra strategy:Cheat.
Ties with Rock, Paper, beats scissors.
Payoff matrix:
Rock is strategy 1, Paper is 2, Scissors is 3, and Cheat is 4 (for row.)

$$A = \begin{bmatrix} 0 & 1 & -1 \\ -1 & 0 & 1 \\ 1 & -1 & 0 \\ 0 & 0 & -1 \end{bmatrix}$$

Note: column knows row cheats.
Why play?
Row is column's advisor.
... boss.

Equilibrium: play the boss...

$$A = \begin{bmatrix} 0 & 1 & -1 \\ -1 & 0 & 1 \\ 1 & -1 & 0 \\ 0 & 0 & -1 \end{bmatrix}$$

Equilibrium:
Row: $(0, \frac{1}{3}, \frac{1}{6}, \frac{1}{2})$. Column: $(\frac{1}{3}, \frac{1}{2}, \frac{1}{6})$.

Payoff? Remember: weighted average of pure strategies.

Row Player.

Strategy 1: $\frac{1}{3} \times 0 + \frac{1}{2} \times 1 + \frac{1}{6} \times -1 = \frac{1}{3}$
Strategy 2: $\frac{1}{3} \times -1 + \frac{1}{2} \times 0 + \frac{1}{6} \times 1 = -\frac{1}{6}$
Strategy 3: $\frac{1}{3} \times 1 + \frac{1}{2} \times -1 + \frac{1}{6} \times 0 = -\frac{1}{6}$
Strategy 4: $\frac{1}{3} \times 0 + \frac{1}{2} \times 0 + \frac{1}{6} \times -1 = -\frac{1}{6}$

Payoff is $0 \times \frac{1}{3} + \frac{1}{3} \times (-\frac{1}{6}) + \frac{1}{6} \times (-\frac{1}{6}) + \frac{1}{2} \times (-\frac{1}{6}) = -\frac{1}{6}$

Column player: every column payoff is $-\frac{1}{6}$.

Both only play optimal strategies! Complementary slackness.

Why not play just one? Change payoff for other guy!

Next time: Multiplicative weights and games.