Online learning and the Metrical Task System problem

(contd)

(+ portfolio selection)

Picking the right hypothesis

Say I am trying to figure out when the Penguins will win.

I have n "hypotheses". One is correct but I don't know which.

Goal: make predictions on the games so as not to make too many mistakes.

Solution: Take majority vote. Cross off incorrect. At most $\lg n$ mistakes.

What if there is no perfect rule?

Can we do nearly as well as the best?

One option: Just start over when every hypothesis has been discarded. Makes $\leq (m+1)\lg n$ mistakes. (m is # mistakes of best hyp).

Better: **Weighted Majority**

- Start with all hypotheses having weight 1.
- Predict based on weighted majority vote.
- Penalize mistakes by cutting weight in half.

"hypotheses" $\leftrightarrow$ "experts"

Analysis: do nearly as well as best expert

- M = # mistakes we've made so far.
- m = # mistakes best expert has made so far.
- W = total weight (starts at n).

- After each mistake, $W \leftarrow cW$ for $c \leq 3/4$.

So, after M mistakes, $W \leq n(3/4)^M$.

- Weight of best expert is $(1/2)^m$. So,

$$(1/2)^m \leq n(3/4)^M$$

$$M \leq 2.4(m + \lg n)$$

A better bound (RWM or EW)

- Instead of taking majority vote, use weights as probabilities.
  (i.e., select expert i with prob $w_i / \sum w_j$)
- Also, generalize $1/2$ to $1 - \epsilon$.

Solves to:

$$
\begin{aligned}
M &\leq \frac{-\ln(1-\epsilon)}{\epsilon} m + \frac{1}{\epsilon}\ln n \\
&\approx (1+\epsilon/2)m + \frac{1}{\epsilon}\ln n
\end{aligned}
$$

(e.g., $1.15m + 4\ln n$) $\leftarrow \epsilon = 1/4$

Analysis

- Say at time t we have fraction F_t of weight on experts that made mistake.
- So, our expected # mistakes $M = \sum_t F_t$.
- $W_{final} = n(1-\epsilon F_1)(1-\epsilon F_2)\ldots$.

$$\begin{aligned}
\ln(W_{final}) &= \ln(n) + \sum_t [\ln(1-\epsilon F_t)] \\
&\leq \ln(n) - \epsilon \sum_t F_t \\
&\quad \text{(using } \ln(1-x) < -x\text{)} \\
&= \ln(n) - \epsilon M
\end{aligned}$$

- If best expert makes m mistakes, then

$$\ln(W_{final}) \geq \ln((1-\epsilon)^m) = m\ln(1-\epsilon).$$

- Now solve....

$$M \leq \frac{-\ln(1-\epsilon)}{\epsilon} m + \frac{1}{\epsilon}\ln n$$

Game-theoretic version

- No predictions, just losses.
- At each time step, alg chooses prob dist $\mathbf{p}$, gets loss vector ℓ, then pays expected cost $\mathbf{p} \cdot \ell$.

RWM/EW algorithm and analysis still applies.

Extends to case of losses in $[0, 1]$ (rather than just $\{0, 1\}$).

Connection to MTS

- 0-step vs 1-step lookahead.
- Movement costs
- What are we comparing to?

Movement costs not so bad in oblivious adversary model.

	$\mathbf{p}$	1/4	1/4	1/4	1/4
E.g.,	ℓ	1	1	0	0
	$\mathbf{p}'$	1/8	1/8	1/4	1/2

In experts setting, we get:

$$\mathbf{E}[\text{Alg loss}] \lesssim (1+\varepsilon/2)[\text{best expert}] + \frac{1}{\epsilon}\log n.$$

In uniform MTS setting (distance $= d$),

$$\mathbf{E}[\text{Alg cost}] \lesssim (1+3\varepsilon/2)[\text{best state}] + \frac{d}{\epsilon}\log n.$$

Neat application: Combining online algorithms online.

Combining online algorithms online

Say we have n paging algorithms.

View online algs as "experts" and do nearly as well as best in hindsight?

- Experts have state. Loss is f(state).
- To switch from following expert X to Y need (at least lazily) to switch state too.

Can view as changing metric space, but always diameter $\leq k$. Apply EW to get:

$$\mathbf{E}[\text{cost}] \lesssim (1 + 3\varepsilon/2)[\text{best alg}] + \frac{1}{\epsilon} k \log n.$$

Which is better: Marking or EW?

Bad case for Marking:

$$\begin{array}{cccc} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ . & . & . & \end{array}$$

Bad case for EW:

$$\begin{array}{cccc} 1 & 1 & 0 & 0 \\ . & . & . & \\ 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ . & . & . & \\ 0 & 0 & 1 & 1 \end{array}$$

Weight-sharing and randomized WF algs give "best of both".

Best of both worlds

Weight-sharing version of EW: take half of weight removed from system and distribute evenly among experts.

- In "experts" setting, for all k,

  $$\mathbf{E}[\text{Alg loss}] \lesssim (1+\varepsilon)L_k + (\frac{1}{\epsilon}\log n)(k+1)$$

  where L_k is loss of best strategy in hindsight that switches k times.

- In uniform MTS setting, can view as:

  $$\mathbf{E}[\text{Alg cost}] \le O(1) \times \text{OPT}_{local} + O(\log n) \times \text{OPT}_{move}.$$

Simpler alg: Phased-RWM

- Run RWM until every expert has incurred total loss $\geq r$.
- Then, reset weights to be all equal again.

Per phase,

$$\mathbf{E}[\text{Alg loss}] \lesssim (1 + \varepsilon/2)r + \frac{1}{\epsilon}\log n.$$

If choose $r = \frac{1}{\alpha\varepsilon}\log n$, then in "experts" setting,

$$\mathbf{E}[\text{Alg loss}] \lesssim (1 + \alpha + \varepsilon/2)L_k + (\frac{1}{\alpha\epsilon}\log n)(k+1)$$

where L_k is loss of best strategy in hindsight that switches k times.

Not quite as good as Share bounds but simpler.

Portfolio selection

- We're considering n stocks.
- Our distribution of wealth is some vector $\mathbf{b}$. E.g., (1/3, 1/3, 1/3).
- Each period, get vector of "price-relatives" $\mathbf{x}$. E.g., (0.98, 1.02, 1.00).
- Our wealth gets multiplied by $\mathbf{b} \cdot \mathbf{x}$.

E.g., 2 stocks, we have 50% of our wealth in each. Stock 1 doubles, stock 2 cut in half.

New wealth $= 2 \cdot \frac{1}{2} + \frac{1}{2} \cdot \frac{1}{2} = 5/4$ times old wealth.

If each period, alg A does 99% as well as alg B, then in t steps,

$$\text{A's wealth} = (0.99)^t(\text{B's wealth}).$$

Looks nicer if take logs:

$$\ln(A) = \ln(B) - 0.01t$$

Competing against the best stock

In a sense, stock problem is "harder" than experts problem:

- In worst case, can't hope to do better than $\frac{1}{n}\times$(performance of best stock in hindsight).

On the other hand, in a sense, it's easier:

- There's a simple strategy to perform at least this well. Or, if take logs,

$$\ln(\text{online}) \geq \ln(\text{optimal stock}) - \ln(n).$$

"Split money evenly at start and let it sit."

Constant-rebalanced Portfolios

CRP($\mathbf{b}$): at end of each period, rebalance back to same distribution of wealth $\mathbf{b}$.

- Common type of strategy.
- Takes advantage of market volatility. "Buy low, sell high."

E.g., look a 2 stocks. Strategy CRP$(1/2, 1/2)$. Suppose stock #1 stays constant (cash), price of stock #2 alternately doubles and halves:

stock 1	stock 2	our wealth
1	1	1
1	2	$3/2$
1	1	$9/8$
1	2	$(3/2)(9/8)$
1	1	$(9/8)^2$

How well can we compete against best of these?

Competing against optimal CRP

What if we split up $ evenly among all CRPs and let it sit?

[Cover][Cover&Ordentlich]:

- Wealth $\geq$ (best CRP)$/(t+1)^{n-1}$.
- Better split: Wealth $\geq$ (best CRP)$/\sqrt{(t+1)^{n-1}}$.

Per-day ratio $\rightarrow 1$.

[Blum & Kalai]:

- Simpler proof of 1st result.
- Extends to transaction costs.

[Helmbold,Schapire,Singer,Warmuth]:

- "experts"-based algorithm. More aggressive.
- Worse guarantees but better on historical data.

Analysis (no commissions)

Split up \$ evenly among all CRPs.

- Nearby portfolios perform similarly:

  If $\mathbf{b} = (1-\alpha)\mathbf{b}_1 + \alpha\mathbf{b}_2$ then

  $$\text{wealth}(\mathbf{b}) \geq (1-\alpha)^t \text{wealth}(\mathbf{b}_1).$$

- Set of CRPs is a simplex.
  Set of CRPs near to $\mathbf{b}^*$ is a smaller simplex.

  $$Vol_{n-1}\{(1-\alpha)\mathbf{b}^* + \alpha\mathbf{b}_2\} = \alpha^{n-1} Vol_{n-1}(\mathbf{b})$$

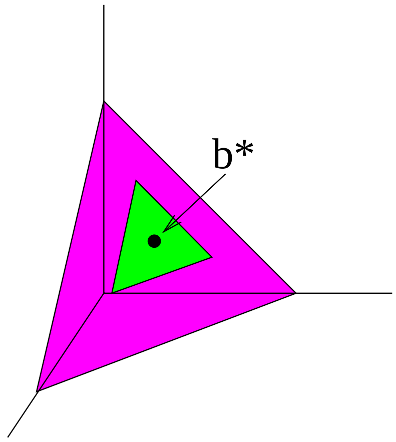

- $(1/t)^{n-1}$ fraction of initial \$ is performing within $1/e$ of optimal.

Discussion

- Performance on two sets of $n = 3$ stocks, over 22 years, rebalancing monthly.

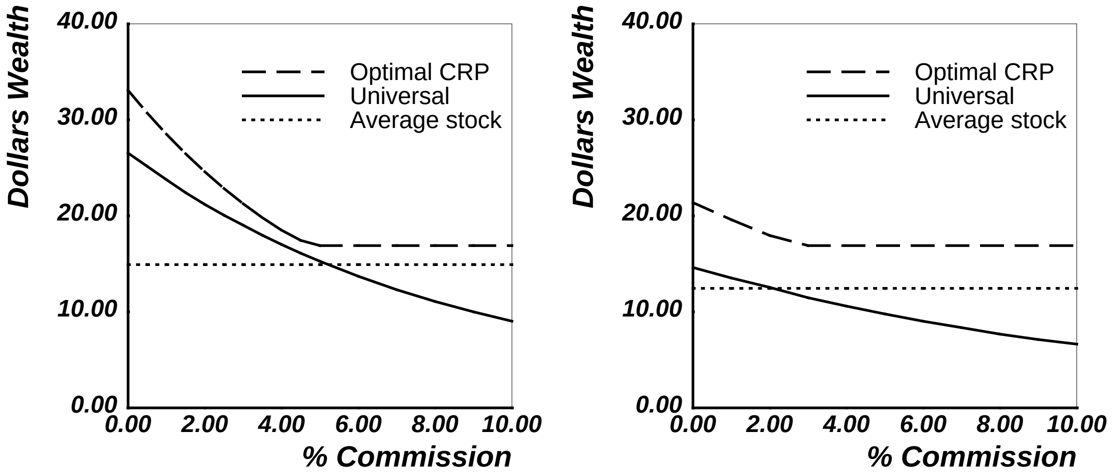

Market 1 - vary commission ***Market 2 - vary commission***

- If commissions are high, can rebalance less often.

Lots of other kinds of online trading results.

Would be good to get a nice direct translation of MTS.