

PS0 solutions on EdStem

PS1 due next week

Last time - finished analysis of Hedge
- application to minimax theorem
for zero-sum games


Application to online convex optimization

Say $f: \mathbb{R}^N \rightarrow \mathbb{R}$ is convex if

$\forall x_0 \in \mathbb{R}^N \quad \exists \lambda \in \mathbb{R}^N$ such that

$$(*) \quad f(x) \geq f(x_0) + \langle x - x_0, \lambda \rangle \quad \forall x \in \mathbb{R}^N$$

The set of $\lambda \in \mathbb{R}^N$ satisfying $(*)$ is denoted by $\partial f(x_0)$ "subgradients of f at x_0 "



$$a(x) = f(x_0) + \langle x - x_0, \lambda \rangle$$

If f is also differentiable, then $\lambda = \nabla f(x_0)$

Examples

- Affine/linear functions
- PSD quadratic functions

$$f(x) = \frac{1}{2} \langle x, Ax \rangle + \langle x, b \rangle + c$$

A is $N \times N$ psd.

- Compositions of convex fn w/ affine fn

$$f(x) = \exp(\langle x, w \rangle)$$

Online Convex Optimization over Δ^{N-1}

For $t=1, 2, \dots, T$:

- Learner chooses $p_t \in \Delta^{N-1}$
- Learner "gets" convex loss function $f_t: \mathbb{R}^N \rightarrow \mathbb{R}$
- Learner incur loss $f_t(p_t) \in \mathbb{R}$

All notions of loss, regret carry over replacing $\langle q, \ell_t \rangle$ with $f_t(q)$ for any $q \in \Delta^{N-1}$

Think about regret of learner to alloc. $q \in \Delta^{N-1}$

$$\sum_{t=1}^T f_t(p_t) - f_t(q)$$

Round t contributes

$$f_t(p_t) - f_t(q)$$

Affine approx of f_t at p_t

$$f_t(q) \geq f_t(p_t) + \langle q - p_t, \lambda_t \rangle$$

say learner gets this λ_t

for $\lambda_t \in \partial f_t(p_t)$

$$f_t(p_t) - f_t(q) \leq \langle p_t - q, \lambda_t \rangle$$

like the contribution of round t to learner's regret to q in O.A. scenario when $\ell_t = \lambda_t$

Punchline: If learner for O.A. (eg. Hedge) that uses λ_t as loss vector in round t has regret bound (for O.A.)

$$\sum_{t=1}^T \langle p_t, \lambda_t \rangle - \sum_{t=1}^T \langle q, \lambda_t \rangle \leq \text{Blah}$$

then it also guarantees

$$\sum_{t=1}^T f_t(p_t) - \sum_{t=1}^T f_t(q) \leq \text{Blah}$$

Lower bound for online allocation aka why Hedge is near-optimal

$R(A; l_1, \dots, l_T) = \text{regret of alg. } A \text{ on}$
loss vectors $l_1, \dots, l_T \in [0,1]^N$.

Thm $\exists C > 0$
For all large enough N, T
all A

$\exists l_1, \dots, l_T \in \{0,1\}^N$ s.t.

$$R(A; l_1, \dots, l_T) \geq C \sqrt{T \log N}$$

Idea ① Define some prob. dist. over
the space of possible T -tuples of loss vectors
and show

$$E[R(A; l_1, \dots, l_T)] \geq C \sqrt{T \log N}$$

② Conclude $\exists \omega \in \text{sample space}$
s.t.

$$R(A; l_1(\omega), \dots, l_T(\omega)) \geq C \sqrt{T \log N}.$$

these are the vectors
we needed to exhibit.

The prob. dist. that we treat treats all $T \times N$ entries
as independent Bernoulli($\frac{1}{2}$) r.v.'s.

$$E[R(A; l_1, \dots, l_T)] = \frac{1}{2} E \left[\max_{i \in [N]} \sum_{t=1}^T \sigma_{t,i} \right]$$

where $\{\sigma_{t,i}\}_{i \in [N], t \in [T]}$ are iid r.v.'s
w/ $P(\sigma_{t,i} = 1) = P(\sigma_{t,i} = -1) = \frac{1}{2}$.

Turns out

$$E \left[\max_{i \in [N]} \sum_{t=1}^T \sigma_{t,i} \right] \geq \sqrt{T \log N}.$$

ONLINE CONCEPT LEARNING

(Closer to "standard" M.L.)

Experts $\longrightarrow$ Concepts

Functions that map "context info" to predictions

$$f: \mathcal{X} \rightarrow \{0,1\}$$

$\uparrow$
input space

$\nwarrow$ binary outcomes

Concept class $\mathcal{C}$: set of possible concepts that you think might be producing the "true" outcomes

We'll think of there being a "ground truth" concept $c^* \in \mathcal{C}$

and on day t , we have context info $x_t \in \mathcal{X}$

we'll think of the true outcome y_t as being the "label" $c^*(x_t)$

Online Concept Learning (aka Mistake Bounded Learning)

For $t=1, 2, \dots$

- Learner sees context info $x_t \in \mathcal{X}$
- Learner predicts $a_t \in \{0,1\}$
- Learner observes true outcome $y_t \in \{0,1\}$

$\mathcal{C}$ -realizability assumption: Outcomes $y_1, \dots, y_t$

are consistent w/ some $c \in \mathcal{C}$, $\forall t$

$\uparrow$

$$c(x_1)=y_1, c(x_2)=y_2, \dots, c(x_t)=y_t$$

Expert Advice Scenario - special case

$$\mathcal{X} = \{0,1\}^N$$

$$x = (x_1, \dots, x_N) \in \mathcal{X}$$

x_i = advice from i^{th} expert

$$\mathcal{C}_{\text{experts}} = \{x \mapsto x_i \mid i \in [N]\}$$

Mistake bound via Halving: $\log_2 N$

In fact for any (finite) $\mathcal{C}$, can just pretend each $c \in \mathcal{C}$ is an "expert"

and run the Halving alg. w/ these "experts"

$\rightarrow$ Mistake bound $\log_2 |\mathcal{C}|$

Problem: runtime is $\Omega(|\mathcal{C}|)$.

We're hoping that for some "structured" concept classes, we can get efficient algs that still have good mistake bounds.

Boolean function concept classes

$$\mathcal{X} = \{0,1\}^n \quad n = \text{dimension}$$

$\uparrow$ Boolean n -cube

$$c: \mathcal{X} \rightarrow \{0,1\}$$

ALL_n = all functions from $\mathcal{X} \rightarrow \{0,1\}$

x	$c(x)$
0000	0
0001	1
0010	1
0011	0

$$|\text{ALL}_n| = 2^{2^n}$$

$$\log_2 |\text{ALL}_n| = 2^n$$

$$\text{MCONJ}_n = \{c_S \mid S \subseteq [n]\}$$

$$\text{where } c_S(x) = \bigwedge_{i \in S} x_i$$

$$|\text{MCONJ}_n| = 2^n$$

Halving gets $\log 2^n = n$ mistake bound.

$\uparrow$

too slow! We'll see an $O(n)$ -time alg.