18.657: Mathematics of Machine Learning

Lecturer: Philippe Rigollet — Lecture 22
Scribe: Aden Forrow — Nov. 30, 2015

7. BLACKWELL'S APPROACHABILITY

7.1 Vector Losses

David Blackwell introduced approachability in 1956 as a generalization of zero sum game theory to vector payoffs. Born in 1919, Blackwell was the first black tenured professor at UC Berkeley and the seventh black PhD in math in the US.

Recall our setup for online linear optimization. At time t, we choose an action $a_t \in \Delta_K$ and the adversary chooses $z_t \in B_\infty(1)$. We then get a loss $\ell(a_t, z_t) = \langle a_t, z_t \rangle$. In the full information case, where we observe z_t and not just $\ell(a_t, z_t)$, this is the same as prediction with expert advice. Exponential weights leads to a regret bound

$$R_n \leq \sqrt{\frac{n}{2}\log(K)}.$$

The setup of a zero sum game is nearly identical:

- Player 1 plays a mixed strategy $p \in \Delta_n$.
- Player 2 plays $q \in \Delta_m$.
- Player 1's payoff is $p^\top M q$.

Here M is the game's payoff matrix.

Theorem: Von Neumann Minimax Theorem

$$\max_{p\in\Delta_n}\min_{q\in\Delta_m} p^\top M q = \min_{q\in\Delta_m}\max_{p\in\Delta_n} p^\top M q.$$

The minimax is called the value of the game. Each player can prevent the other from doing any better than this. The minimax theorem implies that if there is a good response p_q to any individual q, then there is a silver bullet strategy p that works for any q.

Corollary: If $\forall q \in \Delta_n$, $\exists p$ such that $p^\top M q \geq c$, then $\exists p$ such that $\forall q$, $p^\top M q \geq c$.

Von Neumann's minimax theorem can be extended to more general sets. The following theorem is due to Sion (1958).

Theorem: Sion's Minimax Theorem Let A and Z be convex, compact spaces, and $f : A \times Z \to \mathbb{R}$. If $f(a, \cdot)$ is upper semicontinuous and quasiconcave on Z $\forall a \in A$ and

$f(\cdot, z)$ is lower semicontinuous and quasiconvex on A $\forall z \in Z$, then

$$\inf_{a\in A} \sup_{z\in Z} f(a,z) = \sup_{z\in Z} \inf_{a\in A} f(a,z).$$

(Note - this wasn't given explicitly in lecture, but we do use it later.) Quasiconvex and quasiconcave are weaker conditions than convex and concave respectively.

Blackwell looked at the case with vector losses. We have the following setup:

- Player 1 plays $a \in A$
- Player 2 plays $z \in Z$
- Player 1's payoff is $\ell(a, z) \in \mathbb{R}^d$

We suppose A and Z are both compact and convex, that $\ell(a, z)$ is bilinear, and that $\|\ell(a,z)\| \leq R$ $\forall a \in A$, $z \in Z$. All norms in this section are Euclidean norms. Can we translate the minimax theorem directly to this new setting? That is, if we fix a set $S \subset \mathbb{R}^d$, and if $\forall z$ $\exists a$ such that $\ell(a, z) \in S$, does there exist an a such that $\forall z$ $\ell(a, z) \in S$?

No. We'll construct a counterexample. Let $A = Z = [0, 1]$, $\ell(a, z) = (a, z)$, and $S = \{(a, z) \in [0, 1]^2 : a = z\}$. Clearly, for any $z \in Z$ there is an $a \in A$ such that $a = z$ and $\ell(a, z) \in S$, but there is no $a \in A$ such that $\forall z$, $a = z$.

Instead of looking for a single best strategy, we'll play a repeated game. At time t, player 1 plays $a_t = a_t(a_1, z_1, \ldots, a_{t-1}, z_{t-1})$ and player 2 plays $z_t = z_t(a_1, z_1, \ldots, a_{t-1}, z_{t-1})$. Player 1's average loss after n iterations is

$$\bar{\ell}_n = \frac{1}{n} \sum_{t=1}^{n} \ell(a_t, z_t)$$

Let $d(x, S)$ be the distance between a point $x \in \mathbb{R}^d$ and the set S, i.e.

$$d(x, S) = \inf_{s \in S} \|x - s\|.$$

If S is convex, the infimum is a minimum attained only at the projection of x in S.

Definition: A set S is *approachable* if there exists a strategy $a_t = a_t(a_1, z_1, \ldots, a_{t-1}, z_{t-1})$ such that $\lim_{n\to\infty} d(\bar{\ell}_n, S) = 0$.

Whether a set is approachable depends on the loss function $\ell(a, z)$. In our example, we can choose $a_0 = 0$ and $a_t = z_{t-1}$ to get

$$\lim_{n\to\infty} \bar{\ell}_n = \lim_{n\to\infty} \frac{1}{n} \sum_{t=1}^{n} (z_{t-1}, z_t) = (\bar{z}, \bar{z}) \in S.$$

So this S is approachable.

7.2 Blackwell's Theorem

We have the same conditions on A, Z, and $\ell(a, z)$ as before.

Theorem: Blackwell's Theorem Let S be a closed convex set of $\mathbb{R}^2$ with $\|x\| \leq R$ $\forall x \in S$. If $\forall z$, $\exists a$ such that $\ell(a,z) \in S$, then S is approachable.

Moreover, there exists a strategy such that

$$d(\bar{\ell}_n, S) \leq \frac{2R}{\sqrt{n}}$$

Proof. We'll prove the rate; approachability of S follows immediately. The idea here is to transform the problem to a scalar one where Sion's theorem applies by using half spaces.

Suppose we have a half space $H = \{x \in \mathbb{R}^d : \langle w, x\rangle \leq c\}$ with $S \subset H$. By assumption, $\forall z$ $\exists a$ such that $\ell(a,z) \in H$. That is, $\forall z$ $\exists a$ such that $\langle w, \ell(a,z)\rangle \leq c$, or

$$\max_{z\in Z} \min_{a\in A} \langle w, \ell(a,z)\rangle \leq c.$$

By Sion's theorem,

$$\min_{a\in A} \max_{z\in Z} \langle w, \ell(a,z)\rangle \leq c.$$

So $\exists a^*_H$ such that $\forall z$ $\ell(a,z) \in H$.

This works for any H containing S. We want to choose H_t so that $\ell(a_t, z_t)$ brings the average $\bar{\ell}_t$ closer to S than $\bar{\ell}_{t-1}$. An intuitive choice is to have the hyperplane W bounding H_t be the separating hyperplane between S and $\bar{\ell}_{t-1}$ closest to S. This is Blackwell's strategy: let W be the hyperplane through $\pi_t \in \operatorname{argmin}_{\mu\in S} \|\bar{\ell}_{t-1} - \mu\|$ with normal vector $\bar{\ell}_{t-1} - \pi_t$. Then

$$H = \{x \in \mathbb{R}^d : \langle x - \pi_t, \bar{\ell}_{t-1} - \pi_t\rangle \leq 0\}.$$

Find a^*_H and play it.

We need one more equality before proving convergence. The average loss can be expanded:

$$\begin{aligned}
\bar{\ell}_t &= \frac{t-1}{t}\bar{\ell}_{t-1} + \frac{1}{t}\ell_t \\
&= \frac{t-1}{t}(\bar{\ell}_{t-1} - \pi_t) + \frac{t-1}{t}\pi_t + \frac{1}{t}\ell_t
\end{aligned}$$

Now we look at the distance of the average from S, using the above equation and the definition of π_{t+1}:

$$\begin{aligned}
d(\bar{\ell}_t, S)^2 &= \|\bar{\ell}_t - \pi_{t+1}\|^2 \\
&\leq \|\bar{\ell}_t - \pi_t\|^2 \\
&= \left\| \frac{t-1}{t}(\bar{\ell}_{t-1} - \pi_t) + \frac{1}{t}\left(\ell_t - \pi_t\right) \right\|^2 \\
&= \left(\frac{t-1}{t}\right)^2 d(\bar{\ell}_{t-1}, S)^2 + \frac{\|\ell_t - \pi_t\|^2}{t^2} + 2\frac{t-1}{t^2}\langle \ell_t - \pi_t, \bar{\ell}_{t-1} - \pi_t\rangle
\end{aligned}$$

Since $\ell_t \in H$, the last term is negative; since ℓ_t and π_t are both bounded by R, the middle term is bounded by $\frac{4R^2}{t^2}$. Letting $\mu_t^2 = t^2 d(\bar{\ell}_t, S)^2$, we have a recurrence relation

$$\mu_t^2 \leq \mu_{t-1}^2 + 4R^2,$$

implying

$$\mu_n^2 \leq 4nR^2.$$

Rewriting in terms of the distance gives the desired bound,

$$d(\bar{\ell}_t, S) \leq \frac{2R}{\sqrt{n}}$$

□

Note that this proof fails for nonconvex S.

7.3 Regret Minimization via Approachability

Consider the case $A = \Delta_K$, $Z = B_\infty^K(1)$. As we showed before, exponential weights $R_n \leq c\sqrt{n\log(K)}$. We can get the same dependence on n with an approachability-based strategy. First recall that

$$\begin{aligned} \frac{1}{n}R_n &= \frac{1}{n}\sum_{t=1}^n \ell(a_t, z_t) - \min_j \frac{1}{n}\sum_{t=1}^n \ell(e_j, z_t) \\ &= \max_j \left[\frac{1}{n}\sum_{t=1}^n \ell(a_t, z_t) - \frac{1}{n}\sum_{t=1}^n \ell(e_j, z_t)\right] \end{aligned}$$

If we define a vector average loss

$$\bar{\ell}_n = \frac{1}{n}\sum_{t=1}^n (\ell(a_t, z_t) - \ell(e_1, z_t), \ldots, \ell(a_t, z_t) - \ell(e_K, z_t)) \in \mathbb{R}^K,$$

$\frac{R_n}{n} \to 0$ if and only if all components of $\bar{\ell}_n$ are nonpositive. That is, we need $d(\bar{\ell}_n, O_K^-) \to 0$, where $O_K^- = \{x \in \mathbb{R}^K : -1 \leq x_i \leq 0, \forall i\}$ is the nonpositive orthant. Using Blackwell's approachability strategy, we get

$$\frac{R_n}{n} \leq d(\bar{\ell}_n, O_K^-) \leq c\sqrt{\frac{K}{n}}.$$

The K dependence is worse than exponential weights, $\sqrt{K}$ instead of $\sqrt{\log(K)}$.

How do we find a_H^*? As a concrete example, let $K = 2$. We need a_H^* tp satisfy

$$\langle w, \ell(a_H^*, z)\rangle = \langle w, \langle a_H^*, z\rangle y - z\rangle \leq c$$

for all z. Here y is the vector of all ones. Note that $c \geq 0$ since 0 is in S and therefore in H. Rearranging,

$$\langle a_H^*, z\rangle\langle w, y\rangle \leq \langle w, z\rangle + c,$$

Choosing $a_H^* = \frac{w}{\langle w, y\rangle}$ will work; the inequality reduces to

$$\langle w, z\rangle \leq \langle w, z\rangle + c.$$

Approachability in the bandit setting with only partial feedback is still an open problem.

References

[Bla56] D. Blackwell, *An analog of the minimax theorem for vector payoffs*, Pacific J. Math. 6 (1956), no. 1, 1–8

[Sio58] M. Sion, *On general minimax theorems*. Pacific J. Math. 8 (1958), no. 1, 171–176.

18.657 Mathematics of Machine Learning
Fall 2015