Blackwell's Approachability

Facheng Yu

October 27, 2022

Outline

1. Blackwell's approachablity
   - Setup
   - Blackwell's approachability theorem
   - Blackwell's algorithm
2. Applications of Blackwell's approachability
   - Some approachable cases
   - When the target set is known
   - When the target set is unknow
3. Accelerate Blackwell's algorithm by Fenchel duality
   - OCO-based approachability algorithm
4. Summary

Setup

Consider the following two-player vector-valued game:

1. The two players, or Player and Nature, choose strategies from compact convex sets X and Y respectively.
2. There is a bilinear vector-valued payoff function $r(x, y) \in \mathbb{R}^d$.
3. There is a compact convex target set $C \in \mathbb{R}^d$.
4. We assume that $r(x, y), C$ are bounded, for example, by $B(0, 1) = \left\{v \in \mathbb{R}^d : \|v\|_2 \leq 1\right\}$.

Goal for Player: to have the average payoff vector $\bar{r}_T = \frac{1}{T} \sum_{t=1}^{T} r(x_t, y_t)$ approach C, regardless of the strategies of Nature.

Approachable

Definition 1

A target set C is approachable if there exists an algorithm for picking x_t based on $x_1, \ldots, x_{t-1}, y_1, \ldots, y_{t-1}$ such that $d(\bar{r}_t, C) \to 0$ $(t \to \infty)$.

Note:

- $d(\bar{r}_t, C) = \inf_{z \in C} \|z - \bar{r}_t\|_2$.
- A trivial approachable case is that $\exists x' \in X$ s.t. $r(x', y) \in C, \forall y \in Y$.

Halfspace

Definition 2

A halfspace H with respect to a vector $n \in \mathbb{R}^d$ and a scalar b is defined as $H = \{z \in \mathbb{R}^d : \langle n, z \rangle \leq b\}$.

A halfspace $H = \left\{v \in \mathbb{R}^d : \langle v, n \rangle \leq b\right\}$ is approachable
$\iff \exists x' \in X$ s.t. $r(x', y) \in H, \forall y \in Y$.

Note: Consider $d(\langle \bar{r}_t, n \rangle, (-\infty, b]) \to 0$.

Blackwell's Approachability Theorem

Theorem 3

For a convex compact set C, the following statements are equivalent:

1. *C is approachable.*
2. *For each unit vector $n \in \mathbb{R}^d$, there exists $x' \in X$ such that*
$$\langle n, r(x', y)\rangle \leq \sup_{z \in C} \langle n, z\rangle, \ \forall y \in Y.$$
3. *For each $y \in Y$, there exists $x' \in X$ such that $r(x', y) \in C$.*

- Theorem 2(2) defines a halfspace H containing C, where $H = \left\{ v \in \mathbb{R}^d : \langle v, n\rangle \leq \sup_{z \in C} \langle n, z\rangle \right\}$.
- Theorem 2(3) is called the Blackwell's dual condition.

Blackwell's Algorithm

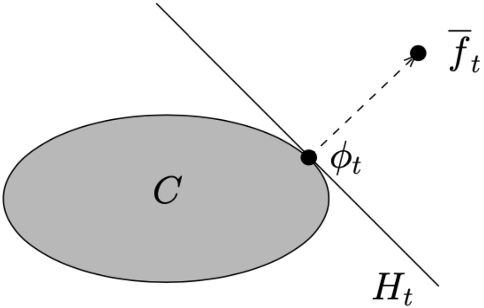

Figure: The halfspace helping select x_t, where $\phi_t = \text{Proj}_C(\bar{r}_t)$.

Blackwell's Algorithm

Assuming that a target set C satisfies Theorem 2(2), Blackwell's algorithm is established as follow:

At each timestep t, do:

1. If $\bar{r}_{t-1} \in C$, play any x_t.
2. Else, consider the unit vector $n_t \propto \bar{r}_{t-1} - \text{Proj}_C(\bar{r}_{t-1})$. Play x_t s.t.

$$\langle n_t, r(x_t, y) \rangle \leq \sup_{z \in C} \langle n_t, z \rangle, \ \forall y \in Y.$$

Note:

- This construction is the key to show C is approachable.
- It guarantees a rate of $O(\frac{1}{\sqrt{T}})$ to approach C.
- Drawback: Computing projection is expensive and sometimes even impossible.

Applications of Blackwell's Approachability

There could be two ways to apply Blackwell's approachablity:

1. When the target set is known, we just need to validate the approachability and apply Blackwell's algorithm.
2. When the target set is unknown, we might need some other techniques to help, like Fenchel duality.

Note:

- There are some occasions where the target set is unknown. For example, in problems of maximizing some objectives, we can't precisely describe the target sets.

Some Approachable Cases

Before discuss the situation where the target set is known, we first give some approachable cases.

1. A closed convex cone C is approachable $\iff \forall z \in C^o, \exists x' \in X$ s.t. $\langle r(x', y), z \rangle \leq 0, \forall y \in Y$, where C^o is the polar cone of C.
2. Let $r(x, y) = (x \cdot y - y_1, ..., x \cdot y - y_d)$, then the negative halfspace $\mathbb{R}^d_-$ is approachable.

Note:

- These can be proved by Theorem 2(2) and (3) respectively.

When the Target Set is Known

A typical example is the online linear learning. Consider an online learning setting where loss vectors $l^1, l^2, \ldots \in [0,1]^d$ are observed. We want to choose weight $w^1, w^2, \ldots \in \Delta_d$ so that

$$\frac{1}{T}\sum_{t=1}^{T} l^t \cdot w^t - \min_{i \in \{1,\ldots,d\}} \frac{1}{T}\sum_{t=1}^{T} l_i^t \leq 0, \ \ T \to \infty.$$

In the corresponding approachability problem, define the payoff fuction r and target set C as follow:

$$r(w,l) = (w \cdot l - l_1, \ldots, w \cdot l - l_d),$$
$$C = \mathbb{R}_-^d := \{v \in \mathbb{R}^d : v_i \leq 0, i = 1, \ldots, d\}.$$

Note: we've already shown that C is approachable under such a payoff function.

When the Target Set is Known

Other applications when the target set is known can be found below:

1. To show the existence of the calibrated forecaster:
   Abernethy, J., Bartlett, P. L., Hazan, E. (2011, December). Blackwell approachability and no-regret learning are equivalent.

2. To show the that MaxWeight in queueing is an instance of Blackwell's policy:
   Walton, N., Xu, K. (2021). Learning and information in stochastic networks and queues.

3. To bulid an approchability algorithm for the partial monitoring problem.
   Kwon, J., Perchet, V. (2017, April). Online learning and blackwell approachability with partial monitoring: optimal convergence rates.

When the Target Set is Unknown

Let's consider a maximizing problem with a concave objective function $f(\bar{r}_t) : \mathbb{R}^d \to \mathbb{R}$. To use approachability, it would be helpful to maximize an upper bound of $f(\bar{r}_t)$ instead of maximizing $f(\bar{r}_t)$ directly.

Define the upper bound with respect to some z as follow:

$$l_f(r(x,y);z) = f(z) - \nabla f(z) \cdot (z - r(x,y)).$$

Note:

- $l_f(\bar{r}_t;z) = \bar{l}_f(r;z)$, $\ l_f(r(x,y);z) \ \geq \ f(r(x,y))$.
- $\max l_f(r(x_t,y_t);z) \iff \min -\nabla f(z) \cdot r(x_t,y_t)$.
- Compared with a benchmark algorithm generating x^* at time t, $l_f(r(x_t,y_t);z) \geq l_f(r(x^*,y^*);z) \geq f(r(x^*,y^*))$.

When the Target Set is Unknown

We now can adapt $\max l_f(\bar{r}_t; z)$ into an approachability version, since $\min -\nabla f(z) \cdot r(x_t, y_t)$ is like to find an optimal halfspace containing the unknown target set.

Idea:

1. Since the target set is unknown, we need a projection-free algorithm to generate the direction vector $n = -\nabla f(z)$ for some z.
2. Online convex optimization can be great, and Fenchel conjugate helps to build a bijection between n and $-\nabla f(z)$, plus the convex function for OCO.

Fenchel Duality

Definition 4

The Fenchel conjugate of a function $f : \mathbb{R}^d \to \mathbb{R}$ is defined as

$$f^*(\theta) = \sup_{z \in \mathbb{R}^d} \{\theta \cdot z + f(z)\}$$

Let $g_t(\theta) = f^*(\theta) - \theta \cdot r(x_t, y_t)$, we have:

1. $\theta_t = -\nabla f(z_t)$, where $z_t = \arg\max_{z \in \mathbb{R}^d}(\theta_t \cdot z + f(z))$.
2. $g_t(\theta_t) = \theta_t \cdot z_t + f(z_t) - \theta_t \cdot r(x_t, y_t) = l_f(r(x_t, y_t); z_t)$.

We can generate θ_t by doing an OCO update for the convex funciton g_t.

The Complete Procedure

The algorithm is established as follow:

Initialize θ_1. For $t = 1, 2, ..., T$, do

1. Set $x_t = \arg\min_{x_t} \max_y \langle \theta_t, r(x_t, y) \rangle$.
2. Choose θ_{t+1} by doing an OCO update for $g_t(\theta) = f^*(\theta) - \theta \cdot r(x_t, y_t)$.

Consider the benchmark algorithm generating x^* at time t, we have:

$$g_t(\theta_t) = l_f(r(x_t, y_t); z_t) \geq l_f(r(x^*, y^*); z_t) \geq f(r(x^*, y^*)),$$

$$\min_\theta \frac{1}{T} \sum_{t=1}^{T} g_t(\theta) = \min_\theta f^*(\theta) - \theta \cdot \bar{r}_T = f(\bar{r}_T),$$

which implies $f(\bar{r}_T) \geq f(\bar{r}(x^*, y^*)) - \frac{1}{T} Regret_T$.

A mixed Example of Bandits

A referential example is the bandits with global convex constraints.
Setup

- A finite set of m arms. A convex set $S \in [0,1]^d$.
- A concave objective g. An unknown value matrix $V \in [0,1]^{m\times d}$.
- Each time t, play an arm $i_t \sim p_t$ and observe a vector v_t.

Goal: $\max g(\frac{1}{T}\sum_{t=1}^{T} v_t)$ s.t. $\frac{1}{T}\sum_{t=1}^{T} v_t \in S$,

which can be transformed, with the UCB technique, to:

$$p_t = \arg\min_{p\in\Delta_m} \min_{\tilde{U}\in H_t} \theta_t \cdot \tilde{U}p$$

$$s.t. \min_{\tilde{V}\in H_t} \phi_t \cdot \tilde{V}p \leq h_S(\phi_t),$$

Agrawal, S., Devanur, N. R. (2019). Bandits with global convex constraints and objective.

Some Concrete Examples

Two examples where Fenchel Duality has been used are listed here:

1. To solve bandits with global convex constraints and objective: *Agrawal, S., Devanur, N. R. (2019). Bandits with global convex constraints and objective.*

2. To solve vector-valued two-player tabular Markov game: *Yu, T., Tian, Y., Zhang, J., Sra, S. (2021, July). Provably efficient algorithms for multi-objective competitive rl.*

Accelerate Blackwell's Algorithm by Fenchel Duality

Since the projection of the average payoff vector can be avoided using Fenchel conjugate, we can build a faster approachability algorithm.

Let $h_C(\theta) = \sup_{z \in C} \langle \theta, z \rangle$ and $g_t(\theta) = h_C(\theta) - \langle \theta, r_t \rangle$, we have

$$
\begin{aligned}
-d(v, C) &= -\inf_{z \in C} \|v - z\|_2 \\
&= -\inf_{z \in C} \sup_{\|\theta\|_2 = 1} \langle \theta, v - z \rangle \\
&= \inf_{\|\theta\|_2 = 1} \{h_C(\theta) - \langle \theta, v \rangle\} = \inf_{\|\theta\|_2 = 1} g_t(\theta),
\end{aligned}
$$

and we have that $h_C(\theta)$ is the Fenchel conjugate of $-d(v, C)$.

We now show that constructing an approachability algorithm can be restated with an OCO update.

OCO-based Approachability Algorithm

For an approachable target set C, consider following OCO update:
Initialize θ_1. For t=1,2,...,T do

1. Set x_t such that $\langle \theta_t, r(x_t, y) \rangle \leq h_C(\theta_t), \forall y \in Y$.
2. Choose θ_{t+1} by doing an OCO update for $g_t(\theta) = h_C(\theta) - \langle \theta, r_t \rangle$.

We have the following statements:

$$-d(\bar{r}_T, C) = \inf_{\|\theta\|_2 = 1} \{h_C(\theta) - \langle \theta, \bar{r}_T \rangle\} = \inf_{\|\theta\|_2 = 1} \frac{1}{T} \sum_{t=1}^{T} g_t(\theta),$$

$$0 \leq \frac{1}{T} \sum_{t=1}^{T} h_C(\theta_t) - \langle \theta_t, r_t \rangle = \frac{1}{T} \sum_{t=1}^{T} g_t(\theta_t),$$

which implies $d(\bar{r}_T, C) \leq \frac{1}{T} Regret_T$.

Shimkin, N. (2016). An online convex optimization approach to Blackwell's approachability.

Summary

To use Blackwell's approachability, one must prove that the target set is approachable or assume that Theorem 2(3) holds, which is that for each $y \in Y$, there exists $x' \in X$ such that $r(x', y) \in C$.

Advantages:

1. It's a natural method when deal with vector-valued games.
2. It provides theoretically feasible and efficient algorithms.

Disadvantages:

1. $h_C(\theta_t)$ is hard to compute, for example, when C is a simplex $\{Ax \leq b\}$.

The End

Thanks for listening.