JMLR: Workshop and Conference Proceedings vol 30 (2013) 1–16

Approachability, fast and slow

Shie Mannor SHIE@EE.TECHNION.AC.IL
Israel Institute of Technology (Technion), Haifa, Israel

Vianney Perchet * VIANNEY.PERCHET@NORMALESUP.ORG
Université Paris Diderot, Paris, France

Abstract

Approachability has become a central tool in the analysis of repeated games and online learning. A player plays a repeated vector-valued game against Nature and her objective is to have her long-term average reward inside some target set. The celebrated results of Blackwell provide a $1/\sqrt{n}$ convergence rate of the expected point-to-set distance if this is achievable, i.e., if the set is approachable. In this paper we provide a characterization for the convergence rates of approachability and show that in some cases a set can be approached with a $1/n$ rate. Our characterization is solely based on a combination of geometric properties of the set with properties of the repeated game, and not on additional restrictive assumptions on Nature's behavior.

Introduction

Approachability goes back to the seminal paper of Blackwell (1956a) who considered a repeated game between a player and Nature. The stage outcome is a vector-valued reward and Blackwell provided sufficient conditions – that happened to be necessary in some cases – under which the player can guarantee that, asymptotically, her average reward vector belongs to some fixed (convex) target set. Such a set is then called approachable, and Blackwell also exhibited a strategy ensuring that the rate of convergence is independent of the dimension of the space (although it depends on the maximal norm of possible rewards) as, no matter the sequence of moves of Nature, the distance at stage n of the average payoff to an approachable set is smaller than $O(n^{-1/2})$.

Approachability theory is now a standard tool in the analysis of repeated games. For example, Kohlberg (1975) proved that it can be used to construct optimal strategies in a class of games with incomplete information; see also Aumann and Maschler (1995); Mertens et al. (1994). It is also widely studied in machine learning, as Blackwell (1956b) himself noticed that regret minimization, introduced by Hannan (1957) earlier, can be easily described as a special instance of approachability. The number of modifications, generalizations and improvements of his original arguments has increased dramatically during the last years; see Hart and Mas-Colell (2001); Cesa-Bianchi and Lugosi (2006); Abernethy et al. (2011); Perchet (2013) and references therein. One of the advantages of approachability theory is that it allows to treat and solve not only usual regret minimization, but also with extra assumptions (cost constraints, variable stage duration, etc.; see Mannor and Shimkin (2008); Mannor et al. (2009); Perchet (2013)) as well as other online learning problems such as

* We thank Jacob Abernethy for asking the original question during COLT '11. Vianney Perchet acknowledges funding from the ANR, under grant ANR-10-BLAN-0112. This research was partially supported by the Israel Science Foundation under grant no. 920/12.

calibration (Dawid, 1982; Foster and Vohra, 1998; Foster, 1999; Mannor and Stoltz, 2010; Rakhlin et al., 2011a).

Regret minimization has become a standard problem in the machine learning community. Much research has been devoted to min-max rates of convergence; see Cesa-Bianchi et al. (2006); Piccolboni and Schindelhauer (2001); Lugosi et al. (2008); Bartók et al. (2011); Foster and Rakhlin (2012); Perchet (2011b). These papers focus on the *finite case*, when the player and Nature have a finite number of pure actions (although they might choose to play at random at some stages) either in full monitoring (when rewards are observed) or in partial monitoring (when rewards are not observed, but only some signals related to them are received by the player). These papers provide convergence rate for the regret based solely on the geometry of both the payoff and monitoring mappings, and not on extra assumptions imposed on Nature's behavior, such as imposing her to change moves *slowly* as in Rakhlin et al. (2011b).

In this paper, we focus only on the full monitoring case when both player and Nature have finite action spaces. For regret minimization a player can only guarantee that the regret is always 0 (when there is a dominating action) or of the order of $n^{-1/2}$ (in other cases). Since approachability theory is more general than regret minimization and is already used to treat other online problems, we are looking to identify the possible convergence rates to an approachable target set. Naively, one could think that as for regret minimization there should be two possible rates: either the rate is 0 (when the set is approached with a single action) or that the rate is of the order of $n^{-1/2}$, recovering the result for regret minimization. It turns out that this conjecture is trivially incorrect as some straightforward examples we provide below show convergence in the intermediate speed of n^{-1}. Notice that in more general regret minimization frameworks, for instance with exp-concave losses, see Cesa-Bianchi and Lugosi (2006), this specific faster ?ate can also be the optimal one.

To sum up, Blackwell proved that the distance to any approachable sets decreases faster than $n^{-1/2}$, in some cases this is tight (as in regret minimization), yet sometimes convergence is proved to occur at n^{-1}. We thus call the later sets are *fast approachable*, in contrast to *slow approachable* for the former sets, as a parallel to fast learning rates in other frameworks Steinwart and Scovel (2005); Audibert and Tsybakov (2007).

Contributions. As a first step to classify sets to fast and slow approachable, we provide two easy to describe and verify conditions under which convex sets are fast approachable. We then proceed to the general results. We exhibit a general sufficient, geometric condition ensuring fast approachability. We also investigate slow approachability by providing a sufficient condition (which is unfortunately not the precise contraposition of the former one) under which Nature has a strategy that guarantee that the convergence rate to the target set is not faster than $n^{-1/2}$. We note that, however, not being fast approachable might not be equivalent to being slow approachable (a property related to the fact that a minmax is not always equal to the associated maxmin). For pedagogic purposes, we present in this version simpler proofs of main results when the underlying space is of dimension two and we point out how to generalize them to higher dimensions.

1. Model and definitions

Consider a vector-valued game between two players, a decision maker (or player) and Nature, with respective finite action sets $\mathcal{I}$ and $\mathcal{J}$, whose cardinalities are referred to as I and J. We denote by d the dimension of the reward vector and equip $\mathbb{R}^d$ with the norm $\|\cdot\|_2$. The payoff function of the player is a mapping $r : \mathcal{I} \times \mathcal{J} \to \mathbb{R}^d$, which is multi-linearly extended to $\Delta(\mathcal{I}) \times \Delta(\mathcal{J})$, the set of product-distributions over $\mathcal{I} \times \mathcal{J}$; the support of a probability distribution is denoted by $\mathrm{supp}(\cdot)$.

At each round, player and Nature simultaneously choose their actions $i_n \in \mathcal{I}$ and $j_n \in \mathcal{J}$ (according to probability distributions denoted by $x_n \in \Delta(\mathcal{I})$ and $y_n \in \Delta(\mathcal{J})$). At the end of a round, the player observe j_n (or equivalently in this framework $r_n := r(i_n, j_n)$).

The following notation will be used. We denote by R the uniform ℓ_2 bound on r and, for every $n \in \mathbb{N}$ and sequence $\{a_m\}_{m\in\mathbb{N}}$, $\bar{a}_n = (1/n)\sum_{m=1}^{n} a_m$. A behavioral strategy σ of the player is a mapping from the set of his finite histories $\cup_{n\in\mathbb{N}} (\mathcal{I} \times \mathcal{J})^n$ into $\Delta(\mathcal{I})$; similarly, a strategy τ of Nature is a mapping from $\cup_{n\in\mathbb{N}} (\mathcal{I} \times \mathcal{J})^n$ into $\Delta(\mathcal{J})$. As usual, we denote by $\mathbb{P}_{\sigma,\tau}$ the probability induced by the pair (σ, τ) onto $(\mathcal{I} \times \mathcal{J})^\infty$.

Definition and some properties of approachability

A closed and convex $\mathcal{C} \subseteq \mathbb{R}^d$ is approachable if there exist a strategy σ of the player such that, for every $\varepsilon > 0$, there exists some $N \in \mathbb{N}$ satisfying, for all strategies τ of Nature:

$$\mathbb{P}_{\sigma,\tau}\left\{\exists n \geq N\, ;\, d_{\mathcal{C}}\left(\bar{r}_n\right) \geq \varepsilon\right\} \leq \varepsilon\ ,$$

where $d_{\mathcal{C}}$ is the point-to-set distance to $\mathcal{C}$. We call σ an approaching strategy of $\mathcal{C}$. Conversely, a convex set $\mathcal{C}$ is excludable if, for some $\delta > 0$, the complement of the δ-neighborhood of $\mathcal{C}$ is approachable by Nature.

Blackwell (1956a) (see also Mertens et al. (1994)) provided the following characterization of approachable convex sets; satisfying it referred as *being a B-set*:

> A set $\mathcal{C} \subset \mathbb{R}^d$ is a B-set if and only if any containing half-space $\mathcal{H} \subset \mathbb{R}^d$ (i.e., such that $\mathcal{C} \subset \mathcal{H}$) is one-shot approachable, meaning that there exists $x[\mathcal{H}] \in \Delta(\mathcal{I})$ such that $r(x, y) \in \mathcal{H}$ no matter $y \in \Delta(\mathcal{J})$.

This naturally leads to the so-called *Blackwell's approaching strategy*. At stage n, let π_n denote the projection of $\bar{r}_n$ onto $\mathcal{C}$ and consider the containing half-space $\mathcal{H}_{n+1}$ whose defining hyperplane is tangent to $\mathcal{C}$ at π_n (in other words, this hyperplane is perpendicular to $\bar{r}_n - \pi_n$ at π_n), see Figure 1. Blackwell's strategy consists of playing accordingly to $x_{n+1} = x[\mathcal{H}_{n+1}] \in \Delta(\mathcal{I})$.

Proposition 1 *Blackwell's approaching strategy of a B-set ensures that*

$$\mathbb{E}[d_{\mathcal{C}}(\bar{r}_n)] \leq \sqrt{\mathbb{E}[d^2_{\mathcal{C}}(\bar{r}_n)]} \leq \frac{2R}{\sqrt{n}} \quad \textit{and} \ \ \mathbb{P}_{\sigma,\tau}\left\{\exists n \geq N\, ;\, d_{\mathcal{C}}\left(\bar{r}_n\right) \geq \varepsilon\right\} \leq \frac{4R}{\varepsilon^2 N}\ .$$

Proof. The proof follows from the following property

$$d^2_{\mathcal{C}}(\bar{r}_{n+1}) \leq \|\bar{r}_{n+1}-\pi_n\|^2 = \frac{n^2}{(n+1)^2}\|\bar{r}_n-\pi_n\|^2+\frac{1}{(n+1)^2}\|r_{n+1}-\pi_n\|^2+2\langle\bar{r}_n-\pi_n, r_{n+1}-\pi_n\rangle.$$

Taking conditional expectation with respect to $\bar{r}_n$ gives that

$$\mathbb{E}[d^2_{\mathcal{C}}(\bar{r}_{n+1})|\bar{r}_n] \leq \frac{n^2}{(n+1)^2}d^2_{\mathcal{C}}(\bar{r}_{n+1}) + \frac{4R^2}{(n+1)^2}.$$

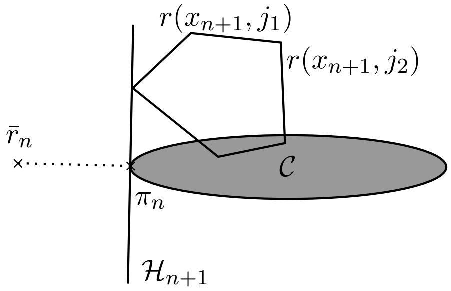

Figure 1: An illustration of Blackwell's approaching strategy. At stage $n+1$, the expected payoff $\mathbb{E}[r_{n+1}] = r(x_{n+1}, j)$ will be on the other side of $\mathcal{H}_{n+1}$ from $\bar{r}_n$.

The first part of the result follows from the convexity of $d_{\mathcal{C}}(\cdot)$. For the second part, an additional martingale (or concentration inequality) argument is required; see Mertens et al. (1994). ■

It turns out that this result holds also for non-convex sets. Yet, in the specific case of closed convex sets, the primal characterization can be transformed into the following *dual characterization* using von Neumann min-max theorem:

$$\text{A convex set } \mathcal{C} \subseteq \mathbb{R}^d \text{ is approachable } \iff \forall\, y \in \Delta(\mathcal{Y}),\ \exists\, x \in \Delta(\mathcal{I}),\ r(x, y) \in \mathcal{C}\ .$$

Note that this might be simpler to formulate and to check, but it does not provide an explicit approachability strategy and it is of no use for the scope of this paper.

Rates of convergence in approachability are space-independent, yet $n^{-1/2}$ is not be optimal in some instances, as illustrated by the following two examples where the convergence rate is $O(1/n)$. The first one consists in a simple control problem.

Example 1 Convergence of empirical distribution to probability distribution. *Consider a one player game (Nature has no role) where $r(i) = (0, \ldots, 1, \ldots, 0) \in \mathbb{R}^I$ is the i-th unit vector and $\mathcal{C} = \{x^*\}$ where $x^* = (x_1^*, \ldots, x_I^*)\} \in \Delta(\mathcal{I})$ is a fixed probability distribution.*

Consider the strategy consisting in choosing action $i_{n+1} \in \arg\max_{i \in \mathcal{I}} x_i - (\bar{r}_n)_i$ which is non-negative since both x^* and $\bar{r}_n$ belongs to $\Delta(\mathcal{I})$. As a consequence, any coordinates $i \in \mathcal{I}$ in excess (i.e., such that $x_i < (\bar{r}_n)_i$) is not played, thus for those coordinates, $(\bar{r}_n)_i - x_i$ decreases in $1/n$. The sums over all coordinates of excesses and in debts (i.e., the quantity $x_i - (\bar{r}_n)_i$ it is non-negative) has to be equal to 0, therefore the sums of debts also decreases at a rate of $1/n$. This implies that $d_{\mathcal{C}}(\bar{r}_n) \leq \frac{2I}{n}$.

This result can be immediately generalized to the framework where $r(i) \in \mathbb{R}^d$ are arbitrary vectors and $\mathcal{C}$ is any subset of the convex hull of $\{r(i), i \in \mathcal{I}\}$. ■

The second example is related to calibration, see Dawid (1982); Foster and Vohra (1998).

Example 2 Easy calibration. *Assume that $\mathcal{J} = \mathcal{I} = \{0, 1\}$, $r(i, j) = i - j$ and $\mathcal{C} = \{0\}$. This represent the framework where Nature chooses at each stage whether it rains $j_n = 1$ or not and the player predicts it. The overall objective is that the average prediction matches the empirical frequency of rain. The simple strategy that selects $i_n = j_{n-1}$ ensures that $|\bar{r}_n| = d_{\mathcal{C}}(\bar{r}_n) = 1/n$.* ■

This motivates the description of geometric conditions under which some set $\mathcal{C}$ has fast rates of approachability. Notice that we seek a characterization that relies *uniquely* on the shape of $\mathcal{C}$ and not, for instance, on additional assumption on the behavior of Nature as, say, that she can only moves slowly. We are concerned with the worst case scenario, also called the *adversarial framework*.

Definition 2 *A closed and convex set is* fast approachable*, if the player has a strategy σ such that, for some constant $\kappa > 0$ and against any strategy τ of Nature,*

$$\mathbb{E}_{\sigma,\tau}\Big[d_{\mathcal{C}}(\bar{r}_n)\Big] \leq \frac{\kappa}{n}, \quad \forall n \in \mathbb{N}.$$

Of course, there exist approachable sets that are not fast approachable; one can think of regret minimization Hannan (1957); Cesa-Bianchi and Lugosi (2006) or the following toy example. Consider $\mathcal{I} = \mathcal{J} = \{0,1\}$, $r(i,j) = \mathbb{1}\{i=j\}$ and $\mathcal{C} = \{1/2\}$. If Nature chooses i.i.d. $j_n = 1$ with probability $1/2$ then, no matter the strategy of the player,

$$\mathbb{E}\Big[d_{\mathcal{C}}(\bar{r}_n)\Big] = \frac{1}{2^n}\binom{n-1}{n/2} \simeq \sqrt{\frac{2}{\pi}}\frac{1}{\sqrt{n}} \text{ for even } n,$$

which gives a $1/\sqrt{n}$ rate. Notice, on the other hand, that in this example $d_{\mathcal{C}}(\mathbb{E}[\bar{r}_n]) = 0$, thus fast approachability is a concept much stronger than fast convergence of expected payoffs.

Similarly to fast approachability, we introduce now the concept of *slow approachability*. It is a stronger requirement than not being fast approachable, as we ask for the existence of a strategy of Nature that guarantees slow rates.

Definition 3 *A closed, convex and approachable set is* slow approachable*, if Nature has a strategy τ such that, for some constant $\kappa' \geq 0$ and against any strategy σ of the player,*

$$\mathbb{E}_{\sigma,\tau}\Big[d_{\mathcal{C}}(\bar{r}_n)\Big] \geq \frac{\kappa'}{\sqrt{n}}, \quad \forall n \in \mathbb{N}.$$

Brief reminder of convex geometry.

As our concepts and proofs rely on geometric properties, we discuss in this section some well known facts from convex geometry that will be used later, see e.g., Rockafellar (1997); Rockafellar and Wets (1998); Ziegler (1995). We recall that $\mathrm{NC}_{\mathcal{C}}(c^*)$, the normal cone to $\mathcal{C}$ at $c^* \in \mathcal{C}$, is an upper-hemi continuous correspondence from $\mathcal{C}$ to $\mathbb{R}^d$ that can be defined by

$$\mathrm{NC}_{\mathcal{C}}(c^*) = \Big\{q \in \mathbb{R}^d \text{ s.t. } \langle q, c^*\rangle = \max_{c\in\mathcal{C}}\langle q, c\rangle\Big\}.$$

Let us denote by $\mathcal{S}^d$ the unit sphere of $\mathbb{R}^d$. Since $\mathrm{NC}_{\mathcal{C}}(c^*)$ is a cone, it can be written as $\mathrm{cone}\{\mathcal{Q}(c^*)\}$ where $\mathrm{cone}(.)$ stands for the positive conic hull and $\mathcal{Q}(c^*) \subset \mathbb{S}^d$ is the intersection between extreme points of $\mathrm{NC}_{\mathcal{C}}(c^*)$ and $\mathcal{S}^d$.

If $\mathcal{C}$ is a polytope with finite set of vertices $\mathcal{V}$, then the family of normal cones $\{\mathrm{NC}_{\mathcal{C}}(v); v \in \mathcal{V}\}$ forms a polyhedral decomposition of $\mathbb{R}^d$ (known as the normal fan) and every set $\mathcal{Q}(c^*)$ is finite. Moreover, for any $z \in c^* + \mathrm{NC}_{\mathcal{C}}(c^*)$ (that is projected onto $\mathcal{C}$ at c^*) there exists a constant $\gamma_{c^*} > 0$, see e.g., Perchet (2011b), property 3 in Appendix A.1, such that

$$d_{\mathcal{C}}(z) = \big\|z - c^*\big\| \leq \gamma_{c^*} \max_{q\in\mathcal{Q}(c^*)}\langle q, z - v^*\rangle.$$

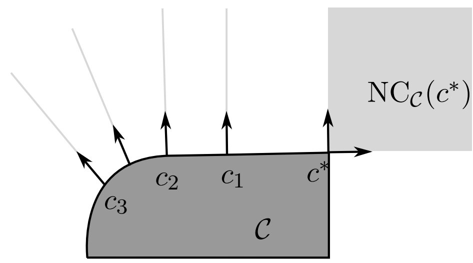

Figure 2: Normal cones are lines on a curved part of $\mathcal{C}$ (from c_3 to c_2), constant on a face (from c_2 to c^* not included) and might have non-empty interior on kinks as at c^*.

Normal cones are constant in the relative interior of a face F of $\mathcal{C}$, so we can define normal cones to a face as the normal cones of any of its interior points, i.e., $\mathrm{NC}_{\mathcal{C}}(F) = \mathrm{NC}_{\mathcal{C}}(v)$ for any v in the relative interior of F. By definition, $F + \mathrm{NC}_{\mathcal{C}}(F)$ consists in all the points in $\mathbb{R}^d$ that project onto F. Therefore $\mathbb{R}^d$ is covered by the union, over the faces of $\mathcal{C}$, of the sets $F + \mathrm{NC}_{\mathcal{C}}(F)$ that forms a polytopial complex (every intersection between two different set has empty interior), as in Figure 3.

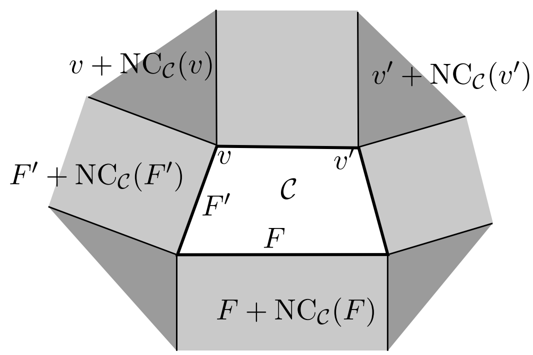

Figure 3: The space $\mathbb{R}^2$ is covered by the unions of sets $v+\mathrm{NC}_{\mathcal{C}}(v)$ (in dark gray) and $F+\mathrm{NC}_{\mathcal{C}}(F)$ (in light gray). The whole polytope $\mathcal{C}$ is a face itself, so it is actually equal to $\mathcal{C}+\mathrm{NC}_{\mathcal{C}}(\mathcal{C})$.

Finally, given two closed and convex sets $\mathcal{C}$ and $\mathcal{C}'$, we say that $\mathcal{C}'$ is a δ-shrinkage of $\mathcal{C}$ if there exists some $\delta > 0$ such that $\mathcal{C}$ contains a δ-neighborhood of $\mathcal{C}'$.

Refinements of the concepts of B-set

First, let us define for every $c \in \mathcal{C}$ and $q \in \mathrm{NC}_{\mathcal{C}}(c)$ (or more generally for any $q \in \mathbb{R}^d$) the projected zero-sum game $\mathcal{G}(c, q)$ between the player (the minimizer) and Nature (the maximizer) with payoff defined by

$$\forall x \in \Delta(\mathcal{I}), \forall y \in \Delta(\mathcal{J}), \quad r^{c,q}(x, y) = \langle r(x, y) - c, q \rangle.$$

Notice that if $q \in \mathrm{NC}_{\mathcal{C}}(c)$, then for every $c' \in \mathcal{C}$, one has by definition of normal cone

$$r^{c',q}(x, y) = r^{c,q}(x, y) + \langle c - c', q \rangle \geq r^{c,q}(x, y). \tag{1}$$

Let $\mathrm{Val}(c, q)$ denote the value of $\mathcal{G}(c, q)$ and $\mathrm{Val}^{\mathrm{pure}}(c, q)$ the min-max in pure actions, i.e.,

$$\mathrm{Val}(c, q) = \min_{x\in\Delta(\mathcal{I})} \max_{y\in\Delta(\mathcal{J})} r^{c,q}(x, y) \text{ and } \mathrm{Val}^{\mathrm{pure}}(c, q) = \min_{i\in\mathcal{I}} \max_{j\in\mathcal{J}} r^{c,q}(i, j).$$

For notational purpose and given $c \in \mathcal{C}$ and $q \in \mathrm{NC}_{\mathcal{C}}(c)$, we also denote by $\mathcal{I}^*(c, q) \subset \mathcal{I}$ the union of the support of every optimal mixed actions of the player in $\mathcal{G}^{c,q}$; the subset $\mathcal{J}^*(c, q) \subset \mathcal{J}$ is defined similarly for Nature. The value of $\mathcal{G}(c, q)$ when both players are restricted to $\mathcal{I}^*(c, q)$ and $\mathcal{J}^*(c, q)$ is denoted by $\mathrm{Val}^{\mathrm{pure}*}(c, q) \leq \mathrm{Val}^{\mathrm{pure}}(c, q)$. Notice that the equality holds for instance as soon as there exists a completely mixed Nash equilibrium in $\mathcal{G}(c, q)$.

The B-set condition can be rephrased in terms of normal cones, yielding a geometric flavor:

$$\forall c \in \mathcal{C},\ \forall q \in \mathrm{NC}_{\mathcal{C}}(c),\ \ \mathrm{Val}(c, q) \leq 0. \tag{2}$$

We now turn to the following geometric categorization of approachable B-sets.

Definition 4 *A closed convex approachable set $\mathcal{C}$ is said to be a*

– PURELY APPROACHABLE B-SET (PAB) *if it satisfies the B-set condition in pure actions, i.e.,*

$$\forall c \in \mathcal{C},\ \forall q \in \mathrm{NC}_{\mathcal{C}}(c), \mathrm{Val}^{\mathrm{pure}}(c, q) \leq 0.$$

In particular, there exists a deterministic approaching strategy of $\mathcal{C}$.

– FINITELY APPROACHABLE B-SET (FAB) *if there exists $\delta > 0$ such that*

$$\forall c \in \mathcal{C},\ \forall q \in \mathrm{NC}_{\mathcal{C}}(c),\ \textit{if } \mathrm{Val}^{\mathrm{pure}}(c, q) > 0 \textit{ then } \mathrm{Val}(c, q) < -\delta.$$

– MIXED APPROACHABLE B-SET (MAB) *if it satisfies*

$$\exists c \in \mathcal{C},\ \exists q \in \mathrm{NC}_{\mathcal{C}}(c),\ \mathrm{Val}^{\mathrm{pure}}(c, q) \geq \mathrm{Val}^{\mathrm{pure}*}(c, q) > 0 \textit{ and } \mathrm{Val}(c, q) = 0.$$

We use the word *finite* in the second condition as it yields, by upper-hemi continuity of normal cones, the existence of a finite set of mixed actions $\{x_k \in \Delta(\mathcal{I}), k \in \mathcal{K}\}$ such that, if the B-set condition is not satisfied in pure action, then $\max_{y\in\Delta(\mathcal{J})} r^{c,q}(x_k, y) \leq -\delta/2$. It is not required that x_k is optimal in these games $\mathcal{G}(c, q)$ but an approaching strategy can only play accordingly to x_k and can therefore consists of alternating between a finite set of mixed actions.

Being (MAB) implies that it is necessary to randomize the actions and therefore obtain random rewards in order to approach $\mathcal{C}$. It is not the exact contraposition of (FAB) since the latter does not require that $\mathrm{Val}^{\mathrm{pure}*}(c, q)$ is also strictly positive. On the other hand, we will use the (MAB) condition to deduce slow approachability, which is stronger than no-fast approachability.

Main results

Our overall objective is to categorize sets depending on their speed of approachability. This categorization is provided in Theorem 6, but we state before the following simpler properties that might be more practical (and its proof is simpler and gives insights for the general case).

Proposition 5 (Simple Characterization of Fast Approachability)

i. A set that contains an approachable δ-shrinkage is fast approachable.

ii. A polytope that is a purely approachable B-set (PAB) is fast approachable.

For the sake of clarity, proofs of these intuitive result are deferred to the following section. Since the case of δ-shrinkable set is settled, we focus on *minimal* approachable sets, that is, sets that do not contain any approachable strict subset. We refer to Spinat (2002) for a proof of their existence.

Theorem 6 *Consider a minimal approachable closed and convex set $\mathcal{C}$.*

i. If $\mathcal{C}$ is a Finitely Approachable B-Set (FAB), then it is fast approachable.

ii. Reciprocally, if $\mathcal{C}$ a Mixed Approachable B-Set (MAB), then it is slow approachable.

This induces the following classification of minimal approachable convex sets:

	$\mathbb{E}[d_{\mathcal{C}}(\bar{r}_n)] = 0$	$\Omega\left(\frac{1}{n}\right) \leq \mathbb{E}[d_{\mathcal{C}}(\bar{r}_n)] \leq O\left(\frac{1}{n}\right)$	$\mathbb{E}[d_{\mathcal{C}}(\bar{r}_n)] \geq \Omega\left(\frac{1}{\sqrt{n}}\right)$
Type of sets	One-shot purely	(FAB) sets	(MAB)-sets

where *one-shot purely* means that there exists $i \in \mathcal{I}$ such that $g(i,j) \in \mathcal{C}$, for all $j \in \mathcal{J}$. Notice that if a set is not one-shot purely approachable then, immediately, $\Omega\left(n^{-1}\right) \leq \mathbb{E}[d_{\mathcal{C}}(\bar{r}_n)]$. Unfortunately, the (MAB) condition (which corresponds to a maxmin of some problem), although quite close, is not necessarily the contraposition of the (FAB) condition (which, on the other hand, is the associated minmax). The full proof of Theorem 6, due to the length restrictions, is omitted; we only point out key steps in Section 3.

2. Shrinkage and purely approachable polytopes: Proof of Proposition 5

We decompose the proof in three major steps: first we treat the first claim, then the second claim in two dimension (in details, as it gives most of the intuitions of results) and we point out how to generalize it to higher dimension.

2.1. Proof of the first claim

We shall prove that if $\mathcal{C}$ contains an approachable δ-shrinkage, then it is fast approachable and convergence rates can be explicitly stated as there exists a strategy σ such that, against any strategy τ of Nature,

$$\mathbb{E}_{\sigma,\tau}\Big[d_{\mathcal{C}}(\bar{r}_n)\Big] \leq \frac{R^2}{\delta n}, \quad \forall n \in \mathbb{N}.$$

Let $\mathcal{C}'$ be any approachable δ-shrinkage of $\mathcal{C}$ and, for any $z \notin \mathcal{C}$, let $\pi_{\mathcal{C}'}(z)$ denotes the projection of z onto $\mathcal{C}'$. By definition of δ-shrinkage, $\pi_{\mathcal{C}'}(z) + \delta(z - \pi_{\mathcal{C}'}(z))/\|z - \pi_{\mathcal{C}'}(z)\|$ belongs to $\mathcal{C}$, thus

$$d_{\mathcal{C}}(z) \leq \left\| z - \left(\pi_{\mathcal{C}'}(z) + \delta \frac{z - \pi_{\mathcal{C}'}(z)}{\|z - \pi_{\mathcal{C}'}(z)\|}\right)\right\| \leq \frac{d^2_{\mathcal{C}'}(z)}{4\delta}.$$

As a consequence, Blackwell's approachability strategy of $\mathcal{C}'$ ensures that

$$\mathbb{E}_{\sigma,\tau}\Big[d_{\mathcal{C}}(\bar{r}_n)\Big] \leq \mathbb{E}_{\sigma,\tau}\left[\frac{d^2_{\mathcal{C}'}(\bar{r}_n)}{4\delta}\right] \leq \frac{R^2}{\delta n},$$

hence $\mathcal{C}$ is fast approachable. ■

2.2. Proof of the second claim in two dimensions

The proof relies on two main arguments (a part of the proof is dedicated to each of them). The first one relies on geometrical properties of $\mathcal{C}$: we prove that the whole space $\mathbb{R}^d$ can be decomposed into some polyhedra in which the same pure action can be played as an approaching strategy. And we show how the distance to $\mathcal{C}$ is controlled in each of these polyhedra. In the second part, we define some sequence $\theta_n \in \mathbb{R}$ that decreases in $1/n$ and such that $d_{\mathcal{C}}(\bar{r}_n) \leq \theta_n$, which will give the result.

FIRST PART. Recall that the union over faces of $\mathcal{C}$ of sets of types $F + \mathrm{NC}_{\mathcal{C}}(F)$ and $v + \mathrm{NC}_{\mathcal{C}}(v)$ covers $\mathbb{R}^d$ (where F ranges over the set of 1-faces of $\mathcal{C}$ and v over its sets of vertices $\mathcal{V}$). We are going to prove that each of these sets can be further subdivided into polyhedra on which the B-set condition is satisfied by the same pure action.

First we consider sets of types $F + \mathrm{NC}_{\mathcal{C}}(F)$. Since, on such a $F + \mathrm{NC}_{\mathcal{C}}(F)$, the normal cone is constant and generated by a unique $q_F \in \mathbb{S}^2$, there exists $i_F \in \mathcal{I}$ such that:

$$\max_{j \in \mathcal{J}} \langle q_F, r(i_F, j) - f \rangle = \max_{j \in \mathcal{J}} r^{f,q_F}(i_F, j) \leq 0, \quad \forall f \in F,$$

hence the B-condition is satisfied by i_F for every point in the relative interior of F.

Moreover, for any point $z \in F + \mathrm{NC}_{\mathcal{C}}(F)$, the distance to $\mathcal{C}$ is controlled as $d_{\mathcal{C}}(z) = \langle z, q_f \rangle$.

We now consider sets of type $v^* + \mathrm{NC}_{\mathcal{C}}(v^*)$ for some vertex $V^* \in \mathcal{V}$. Given an action $i \in \mathcal{I}$, we define the convex subset $\mathrm{NC}_{\mathcal{C}}(v^*)^i \subset \mathrm{NC}_{\mathcal{C}}(v^*)$ by

$$\mathrm{NC}_{\mathcal{C}}(v^*)^i = \left\{ q \in \mathrm{NC}_{\mathcal{C}}(c), \max_{j \in \mathcal{J}} \ r^{c,q}(i,j) = \max_{j \in \mathcal{J}} \langle q, r(i,j) - c \rangle \leq 0 \right\}.$$

The convex set $\mathcal{C}$ is purely approachable (PAB), so the family $\{\mathrm{NC}^i_{\mathcal{C}}(v^*), i \in \mathcal{I}\}$ covers $\mathrm{NC}_{\mathcal{C}}(v^*)$. All sets $\mathrm{NC}^i_{\mathcal{C}}(v^*)$ are cones so, up to considering intersection, that they can be assumed to form a polyhedral complex of $\mathrm{NC}^i_{\mathcal{C}}(F)$, i.e., any intersection between two different $\mathrm{NC}^i_{\mathcal{C}}(v^*)$ has an empty interior.

The cone $\mathrm{NC}^i_{\mathcal{C}}(v^*)$ is an affine cone with vertex v^*, generated by two directions $q, q' \in \mathbb{S}^2$. We denote by $\phi^i(v^*)$ the angle between q and q' which is, without loss of generality (up to dividing $\mathrm{NC}^i_{\mathcal{C}}(v^*)$ in two), smaller than $\pi/2$. As a consequence, for any $z \in v^* + \mathrm{NC}_{\mathcal{C}}(v^*)$, simple trigonometric considerations show that

$$d_{\mathcal{C}}(z) \leq \max \left\{ \frac{\langle z, q \rangle}{\cos(\phi^i(v^*)/2)}, \frac{\langle z, q' \rangle}{\cos(\phi^i(v^*)/2)} \right\} \leq \sqrt{2} \max \left\{ \langle z, q \rangle, \langle z, q' \rangle \right\}.$$

So we just proved the existence of a finite polyhedral decomposition $\{P_\ell; \ell \in \mathcal{L}\}$ of $\mathbb{R}^d$ and $\gamma > 0$ such that, on every polyhedron P_ℓ, the same $i_\ell \in \mathcal{I}$ can be used to satisfy Blackwell's condition (this is illustrated in Figure 4), i.e.

$$\forall z \in P_\ell, \ \max_{q \in \mathcal{Q}_\ell} \langle q, r(i_\ell, j) - \pi_{\mathcal{C}}(z) \rangle \leq 0, \forall j \in \mathcal{J}.$$

and such that

$$d_{\mathcal{C}}(z) \leq \gamma \max_{q \in \mathcal{Q}_\ell} \langle q, z - \pi_{\mathcal{C}}(z) \rangle := \gamma \, \theta_\ell(z). \tag{3}$$

where $\mathcal{Q}_\ell \subset \mathbb{S}^2$ is the set of directions that generate P_ℓ (the cardinality of $\mathcal{Q}_\ell$ is either one or two).

We introduced this mapping θ_ℓ for the following reason. Assume for the moment that $\bar{r}_n$ always stays in the same P_ℓ, then playing at each stage the same i_ℓ ensures that $\theta_\ell(\bar{r}_{n+1}) \leq \frac{n}{n+1}\theta_\ell(\bar{r}_n)$. Hence both $\theta_\ell(\bar{r}_n)$ and $d_{\mathcal{C}}(\bar{r}_n)$ will decrease at a rate of $1/n$.

We now generalize this idea when $\bar{r}_n$ can change from one P_ℓ to another. The proof relies on the fact that the polyhedral decomposition also satisfies the following crucial property.

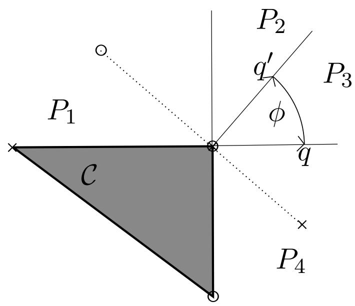

Figure 4: In this game, the player has two pure actions X and O and Nature 3 pure actions. Possible payoffs are represented, by crosses or circles. When $\bar{r}_n$ is in P_1 or P_2, Blackwell's strategy dictates to play X, while if if is in P_3 and P_4 to play O.

Define $\varepsilon = \min_{i\in I, v\in\mathcal{V}} \tan(\phi^i(v)/2)$. Consider some $z \in \mathbb{R}^d$ such that $d_{\mathcal{C}}(z) = 1$ and denote by ℓ^* the index of a polyhedron P_{ℓ^*} that contains z. Once again, simple trigonometric computations show that for every $\eta \in \mathbb{R}^d$ such that $\|\eta\| \leq \varepsilon$, then $z + \eta$ belongs to a polyhedron P_ℓ which is a neighbor of P_{ℓ^*}. More importantly $\theta_\ell(z+\eta) = \theta_{\ell^*}(z+\eta)$ or, stated otherwise, the maximum in the definition of $\theta_{\ell'}(z+z')$ is attained at the point in q that belongs to both $\mathcal{Q}_\ell$ and $\mathcal{Q}_{\ell^*}$

SECOND PART. At stage n, we denote by ℓ_n any polytope to which $\bar{r}_n$ belongs and we define $\theta_n := \theta_{\ell_n}(\bar{r}_n)$. We now claim that if $n\theta_n \geq 2\|r\|\varepsilon^{-1}$, then $(n+1)\theta_{n+1} \leq n\theta_n$. This would immediately entail the result as it implies that

$$d_{\mathcal{C}}(\bar{r}_n) \leq \gamma\theta_n \leq 2\|r\|\gamma\frac{(\varepsilon^{-1}+1)}{n}.$$

It only remains to prove the claim. If $n\theta_n \geq 2\|r\|\varepsilon^{-1}$, then $d_{\mathcal{C}}(\bar{r}_n) \geq 2\|r\|\varepsilon^{-1}/n$ and since $\|\bar{r}_{n+1} - \bar{r}_n\| \leq 2\|r\|/n$, one has, after renormalization, that $\|\bar{r}_{n+1} - \bar{r}_n\|/d_{\mathcal{C}}(\bar{r}_n) \leq \varepsilon$. As a consequence,

$$\theta_{\ell_{n+1}}(\bar{r}_{n+1}) = \theta_{\ell_n}(\bar{r}_{n+1}) = \max_{q\in\mathcal{Q}_{\ell_n}} \langle q, \bar{r}_{n+1} - \pi_{\mathcal{C}}(\bar{r}_{n+1})\rangle = \max_{q\in\mathcal{Q}_{\ell_n}} \langle q, \bar{r}_{n+1} - \pi_{\mathcal{C}}(\bar{r}_n)\rangle,$$

where the last equality comes from the fact that $\bar{r}_{n+1}$ and $\bar{r}_n$ must be in two neighbor polyhedra, thus $\langle q, \bar{r}_n - \bar{r}_{n+1}\rangle = 0$. Therefore,

$$\theta_{n+1} \leq \frac{n}{n+1}\theta_n + \frac{1}{n+1}\max_{q\in\mathcal{Q}_{\ell_n}} \langle q, r_{n+1} - \pi_{\mathcal{C}}(\bar{r}_n)\rangle \leq \frac{n}{n+1}\theta_n$$

as soon as $i_{n+1} = i_{\ell_n}$, as prescribed by Blackwell's strategy. ■

2.3. Generalization to higher dimensions

The following issue emerges in dimension bigger than two. It is no longer true, with the previous definition of ε and θ_ℓ, that $\theta_\ell(z+\eta) = \theta_{\ell^*}(z+\eta)$ (see the counter example in Figure 5); we have to adapt those definitions as follows.

Let δ be the smallest distance between a point $v+q$ (for some vertex $v \in \mathcal{V}_\ell$ and $q \in Q_\ell$) and any face of any $P_{\ell'}$ that does not contain it. Since there exists a finite number of such faces, δ is strictly positive. We define now the following subset of P_ℓ:

$$W_\ell = \left\{ \frac{\omega - \pi_\mathcal{C}(\omega)}{\|\omega - \pi_\mathcal{C}(\omega)\|} \text{ for } \omega \in \partial P_\ell \right\} \backslash B_\ell, \text{ where } B_\ell = \bigcup_{q \in Q_\ell} B(q,\delta) \backslash \bigcup_{q \in Q_\ell} \{q\},$$

where ∂P_ℓ is the boundary of P_ℓ and $B(q,\delta)$ is the open ball centered in q of radius δ. This construction is illustrated in Figure 5.

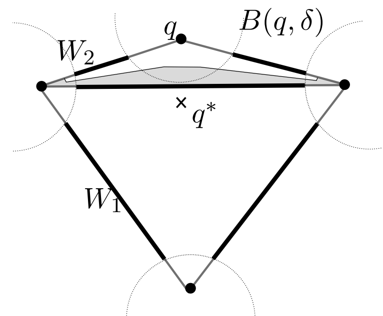

Figure 5: A cut of the normal cone at some vertex v is represented here with two polyhedra P_1 (the triangle below) and P_2 (the one at the top). W_i is the boundary of P_i, minus balls centered at vertices, plus these vertices (it is represented by the thick black lines and points). The shaded area represents points that are closer to $W_1 \cap W_2$ (the horizontal line) than to $W_2 \backslash W_1$; it is bounded away from P_1 by some positive distance. If we just consider for W_i the union of vertices, then the point q^* will be closer to W_2 than to W_1, even though it is in the interior of P_1 – notice that this property only happens when $d > 2$.

We now redefine the key mappings $\theta_\ell : \mathbb{R}^d \to \mathbb{R}$ as follows:

$$\theta_\ell(z) := \max_{\omega \in W_\ell} \langle \omega, z - \pi_\mathcal{C}(z) \rangle = \frac{1}{2}\left(1 + \|z - \pi_\mathcal{C}(z)\|^2 - \min_{\omega \in W_\ell} \|\omega - (z - \pi_\mathcal{C}(z))\|^2\right).$$

We claim that there still exists a small $\varepsilon > 0$ satisfying the crucial property that, for any pair $z, z' \in \mathbb{R}^d$ with $z \in \mathbb{S}_\ell := \{\omega \in P_\ell \text{ s.t. } d(\omega, \mathcal{C}) = 1\}$ and $z + z' \in P_{\ell'}$, if $\|z'\| \leq \varepsilon$ then $\theta_{\ell'}(z+z') = \theta_\ell(z+z')$. Stated otherwise, the claim is that the maximum in the definition of $\theta_{\ell'}(z+z')$ is attained at a point in $W_{\ell,\ell'} := W_\ell \cap W'_\ell$.

Indeed, for ε small enough, P_ℓ and $P_{\ell'}$ must be neighbors, i.e., they share a common face $F_{\ell,\ell'}$ (this is due to the fact that the distance from any $\mathbb{S}_\ell$ to P_k that are not neighbors is lower bounded by some positive quantity). By continuity of θ_ℓ, the result follows from

$$\max_{\omega \in W'_\ell \cap W_\ell} \langle \omega, z'' - \pi_\mathcal{C}(z'') \rangle > \max_{\omega' \in W'_\ell \backslash W_\ell} \langle \omega', z'' - \pi_\mathcal{C}(z'') \rangle, \quad \forall z'' \in \mathbb{S}_\ell \cap \mathbb{S}_{\ell'}.$$

This inequality is immediately true if $z'' - \pi_{\mathcal{C}}(z'') \in W_{\ell,\ell'}$. It also holds in the remaining cases, when $z'' - \pi_{\mathcal{C}}(z'') \in B(q^*, \delta/2)$ for some $q^* \in \mathcal{Q}_\ell \cap \mathcal{Q}_{\ell'}$. The most difficult configuration is when z'' belongs to the common face $F_{\ell,\ell'}$, so we can restrict ourself to this case.

Recall that if q and q' are such that $\|q - q^*\| = \|q' - q^*\| = \delta$ then every points in the segment $(q^*, q]$ is strictly closer to q than to q'. Therefore, if the maximum in the left side term is not attained at q^*, it is attained at $q^* + \delta(q^* - (z'' - \pi_{\mathcal{C}}(z''))/\|(q^* - (z'' - \pi_{\mathcal{C}}(z''))\| \in W'_\ell \cap W_\ell$. And the left side term is, in that latter case, strictly greater than the right side term.

3. Insights and examples of the main result

It is clear that the condition (PAB) is not necessary, as it essentially relies on the finite polyhedral decomposition of the space. So, if such a decomposition exists, but with respect to a finite set of mixed strategies, then the result still holds, at least in expectation. In the following example, we provide a minimal non-trivial approachable set that does not satisfy (PAB) but is fast approachable.

Example 3 *Consider the game described in Figure 3. The player has* 4 *pure actions, cross, circle, square and triangle. Nature has 5 actions, enumerated from 1 to 5. Payoffs associated to a pair of action are represented by the associated symbol and a number (or two numbers if they generate the same payoff) next to it. For instance, if the player chooses action "cross" and Nature chooses action "3", the payoff is the top left corner of* $\mathcal{C}$. *If they choose, on the contrary, action "circle" and "5", the payoff is located on the right of* $\mathcal{C}$. *We also depict by a black filled triangle the expected payoff induced by playing triangle and square with probability one half each.*

The set $\mathcal{C}$ *is minimum approachable, it requires at least one mixed action (for instance* $1/2$ *triangle+*$1/2$ *square) to approach it, but whenever it has to be used, Equation (2) is not tight. In particular, this example show that it is not correct to assume that, for a minimal convex approachable set, Equation (2) must hold with equality in every directions.* ■

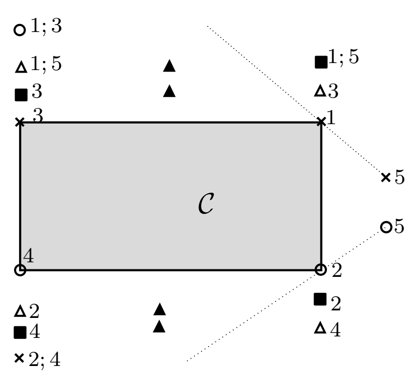

Figure 6: A minimal approachable convex set, only in mixed actions.

Due to the length restriction, the quite technical and geometric proof of Theorem 6 is omitted. Yet, we can describe the steps it follows:

1) A minimal convex set satisfying the (FAB) condition must be a polytope. This is solely due to the fact that there exists a finite number of mixed actions needed to approach it.

2) The proof of the second part of Proposition 5 can be extended with mixed actions, using the core idea behind the proof of its first part: although averages of random i.i.d. variables converge typically to their expectation at a rate of $n^{-1/2}$, the distance to any δ-neighborhood of the latter is at a rate of n^{-1}. Consider for instance $Z \sim \mathcal{N}(0, \sigma^2/n)$, then indeed

$$\mathbb{E}\Big[d_{[-\delta,\delta]}(Z)\Big] = \mathbb{E}\Big[\big(|Z| - \delta\big)\mathbb{1}\{|Z| \geq \delta\}\Big] \leq \sqrt{\frac{2\sigma^2}{\pi n}} \exp\left(-\frac{\delta^2 n}{2\sigma^2}\right)$$

and fast convergence occurs (it is actually exponentially fast in this example).

3) To prove the last part of Theorem 6, we consider a specific projected game $\mathcal{G}(c,q)$ such that $\mathrm{Val}(c,q) = 0$ but $\mathrm{Val}^{\mathrm{pure}*}(c,q)$. In this game Nature has an optimal mixed action and the variance when this action is played ensures that convergence cannot be faster than $n^{-1/2}$.

We conclude this section by An intriguing and counter intuitive example.

Example 4 *Consider any of the two games where payoff matrices are given by*

	L	R
T	*-2*	*1*
B	*2*	*-1*

and

	L	R
C	*1*	*-2*
M	*-1*	*2*

and define the same convex target set $\mathcal{C} = \{0\}$ in both games. It is easy to show that $\mathcal{C}$ is slow approachable in any of these games. Indeed, in the left one, Nature just has to play i.i.d. action L with probability $1/3$ and R with probability $2/3$ and similarly in the right one. On the other hand, to approach $\mathcal{C}$, the player has to play at each stage T and B (or C and M) with probability $1/2$.

On the contrary, in a game where the player's action set is $\{T, B, C, M\}$ and Nature's one is $\{L, R\}$ (basically, we concatenate the two payoff matrices by putting them one on top of the other) then $\mathcal{C}$ becomes fast approachable. However, the strategy of the player is different. Indeed, playing i.i.d. T and B (or similarly C and M) with the probability $1/2$ only ensure slow approachability.

The strategy that consists in choosing T and M with probability $1/2$ if $\bar{r}_n \geq 0$ (so that the expected payoff is $1/2$) or B and C with probability $1/2$ if $\bar{r}_n < 0$ is fast approaching. Indeed, when $\bar{r}_n > 0$, it decreases at the rate of $1/n$, and similarly if $\bar{r}_n < 0$. ■

4. Conclusion

We provided a partial answer to the question when is approachability fast and when is it slow? We did that for the most natural model where one looks for approaching a target set in a vector-valued game and the reward is deterministic. There are three variations of this model that are of interest. The first is the case where the reward itself is stochastic. In this case, it is not hard to prove that a (FAB) set is still fast approachable under mild conditions on the reward distribution. On the other hand, a (PAB) set may be slow approachable. The second is the case where instead of looking at the expected distance to the set, we consider the distance of the expectation to the set. In this case, it is not hard to prove that a sufficient condition for fast approachability is if there exists a strategy that uses finitely many mixed actions. The last variation is concerned with approachability with partial monitoring (Perchet, 2011a; Mannor et al., 2011). In this last framework, determining even slow (or worst case) convergence rates is still an open problem. Complete solutions to the rate of approachability in all three variations are left for future research.

References

J. Abernethy, P. L. Bartlett, and E. Hazan. Blackwell approachability and low-regret learning are equivalent. In *Proceedings of the Twenty-Fourth Annual Conference on Learning Theory (COLT'11)*, 2011.

Jean-Yves Audibert and Alexandre B. Tsybakov. Fast learning rates for plug-in classifiers. *Ann. Statist.*, 35:608–633, 2007.

R.J. Aumann and M.B. Maschler. *Repeated Games with Incomplete Information*. MIT Press, 1995.

G. Bartók, D. Pál, and C. Szepesvári. Minimax regret of finite partial monitoring games in stochastic environments. In *Proceedings of the Twenty-Fourth Annual Conference on Learning Theory (COLT'11)*, 2011.

D. Blackwell. An analog of the minimax theorem for vector payoffs. *Pacific Journal of Mathematics*, 6:1–8, 1956a.

D. Blackwell. Controlled random walks. In *Proceedings of the International Congress of Mathematicians, 1954, Amsterdam, vol. III*, pages 336–338, 1956b.

N. Cesa-Bianchi and G. Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, 2006.

N. Cesa-Bianchi, G. Lugosi, and G. Stoltz. Regret minimization under partial monitoring. *Mathematics of Operations Research*, 31:562–580, 2006.

A.P. Dawid. The well-calibrated Bayesian. *Journal of the American Statistical Association*, 77: 605–613, 1982.

D. Foster and R. Vohra. Asymptotic calibration. *Biometrika*, 85:379–390, 1998.

D. P. Foster. A proof of calibration via Blackwell's approachabilitytheorem. *Games Econom. Behav.*, 29:73–78, 1999.

D. P. Foster and A. Rakhlin. No internal regret via neighborhood watch. In *Proceedings of the Fifteenth International Conference on Artificial Intelligence and Statistics (AISTATS'12)*, 2012.

J. Hannan. Approximation to Bayes risk in repeated play. In *Contributions to the Theory of Games*, volume 3 of *Annals of Mathematics Studies*, pages 97–139. Princeton University Press, Princeton, N. J., 1957.

S. Hart and A. Mas-Colell. A general class of adaptive strategies. *Journal of Economic Theory*, 98: 26–54, 2001.

E. Kohlberg. Optimal strategies in repeated games with incomplete information. *International Journal of Game Theory*, 4:7–243, 1975.

G. Lugosi, S. Mannor, and G. Stoltz. Strategies for prediction under imperfect monitoring. *Mathematics of Operations Research*, 33:513–528, 2008.

S. Mannor and N. Shimkin. Regret minimization in repeated matrix games with variable stage duration. *Games and Economic Behavior*, 63(1):227–258, 2008.

S. Mannor and G. Stoltz. A geometric proof of calibration. *Mathematics of Operations Research*, 35:721–727, 2010.

S. Mannor, J. Tsitsiklis, and J. Y. Yu. Online learning with sample path constraints. *Journal of Machine Learning Research*, 10:569–590, 2009.

S. Mannor, V. Perchet, and G. Stoltz. Robust approachability and regret minimization in games with partial monitoring. In *Proceedings of the Twenty-Fourth Annual Conference on Learning Theory (COLT'11)*, 2011.

J.-F. Mertens, S. Sorin, and S. Zamir. Repeated games. Technical Report no. 9420, 9421, 9422, Université de Louvain-la-Neuve, 1994.

V. Perchet. Approachability of convex sets in games with partial monitoring. *Journal of Optimization Theory and Applications*, 149:665–677, 2011a.

V. Perchet. Internal regret with partial monitoring calibration-based optimal algorithms. *Journal of Machine Learning Research*, 2011b. In press.

V. Perchet. Approachability, regret and calibration, implications and equivalences. *ArXiv.org:1301.2663*, 2013.

A. Piccolboni and C. Schindelhauer. Discrete prediction games with arbitrary feedback and loss. In *Proceedings of the Fourteenth Annual Conference on Computational Learning Theory (COLT'01)*, pages 208–223, 2001.

A. Rakhlin, K. Sridharan, and A. Tewari. Online learning: Beyond regret. In *Proceedings of the Twenty-Fifth Annual Conference on Neural Information Processing Systems (NIPS'11)*, 2011a.

A. Rakhlin, K. Sridharan, and A. Tewari. Online learning: Stochastic and constrained adversaries. In *Proceedings of the Twenty-Fourth Annual Conference on Learning Theory (COLT'11)*, 2011b.

R. T. Rockafellar. *Convex analysis*. Princeton Landmarks in Mathematics. Princeton University Press, Princeton, NJ, 1997. Reprint of the 1970 original, Princeton Paperbacks.

R. T. Rockafellar and R. J.-B. Wets. *Variational analysis*, volume 317 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 1998.

X. Spinat. A necessary and sufficient condition for approachability. *Math. Oper. Res.*, 27:31–44, 2002.

I. Steinwart and C. Scovel. Fast rates for support vector machine. In *Proceedings of the eighteenth Annual Conference on Learning Theory (COLT'05)*, 2005.

G. Ziegler. *Lectures on Polytopes*, volume 152 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1995.