Are regret minimization and gradient equilibrium equivalent problem frameworks?

Brian W. Lee   Nika Haghtalab   Michael I. Jordan   Ryan J. Tibshirani

As an overview, we will present oracle reductions between gradient equilibrium (GEQ) and a restricted form of Blackwell approachability with nonempty closed, convex target cones, which show that an approachability problem can always be solved using a GEQ algorithm, and vice versa. Furthermore, when a GEQ algorithm with error rate $\phi(T)$ is used alongside our reduction, the underlying approachability problem is solved at the same rate $\phi(T)$; and vice versa. Combining our reductions with known oracle reductions between regret minimization (RM) and this restricted form of Blackwell approachability (Abernethy et al. , 2011) provides an affirmative answer to the question under consideration.

1 Preliminaries

1.1 Regret minimization

A *regret minimization* (RM) problem is a sequential game between a learner and an adversary parameterized by $(\mathbb{R}^d, \mathcal{Z}, \mathcal{L})$, defined as follows. Let $B, L \geq 0$ be constants, and let:

- $\mathcal{Z} \subseteq \mathbb{R}^d$ be a nonempty closed convex decision set which is B-bounded in the Euclidean norm.
- $\mathcal{L} \subseteq \{\ell : \mathcal{Z} \to \mathbb{R}\}$ be a set of convex, L-Lipschitz, and differentiable loss functions.
- $\Phi_R : \cup_{t=1}^{\infty} H_t \to \mathbb{R}$ be the regret error metric, where each $H_t = \mathcal{L}^t \times \mathcal{Z}^t$. That is, for $T \geq 1$,

$$\Phi_R(\ell_{1:T}, z_{1:T}) = \sum_{t=1}^{T} \ell_t(z_t) - \min_{z \in \mathcal{Z}} \sum_{t=1}^{T} \ell_t(z),$$

  where we abbreviate $\ell_{1:T} = (\ell_1, \dots, \ell_T)$ and $z_{1:T} = (z_1, \dots, z_T)$.

At each integer round $t \geq 1$,

1. the learner selects $z_t \in \mathcal{Z}$;
2. the adversary selects $\ell_t \in \mathcal{L}$;
3. the learner observes $g_t(z_t) = \nabla \ell_t(z_t)$.

An algorithm for the learner in the problem $(\mathbb{R}^d, \mathcal{Z}, \mathcal{L})$ is a mapping $\mathcal{A}_R : \{\emptyset\} \cup (\cup_{t \geq 1} (\mathbb{R}^d)^t) \to \mathcal{Z}$ that has knowledge of $\mathbb{R}^d$ and $\mathcal{Z}$ but not $\mathcal{L}$. At each round $t \geq 1$, $\mathcal{A}_R$ selects a decision $z_t \in \mathcal{Z}$ according to

$$z_t = \mathcal{A}_R(g_1(z_1), \dots, g_{t-1}(z_{t-1}))$$

with input $\emptyset$ when $t = 1$. $\mathcal{A}_R$ guarantees that the resulting sequence $z_{1:T}$ incurs error $\Phi_R(\ell_{1:T}, z_{1:T}) = O(\phi_R(T))$, where ϕ_R is a positive and increasing error rate function which satisfies $\phi_R(T) = o(T)$.

1.2 Gradient equilibrium

A *gradient equilibrium* (GEQ) problem is a sequential game between a learner and an adversary parameterized by $(\mathbb{R}^d, \mathcal{G})$, defined as follows. Let $L, h \geq 0$ be constants, and let:

- $\mathcal{G} \subseteq \{g : \mathbb{R}^d \to \mathbb{R}^d\}$ be a set of vector fields which are L-bounded in the Euclidean norm, and satisfy the restorative condition: for each $g \in \mathcal{G}$ and $\theta \in \mathbb{R}^d$,

$$\|\theta\|_2 \geq h \implies \langle g(\theta), \theta \rangle \geq 0.$$

  These vector fields can, but need not, be gradient fields.

- $\Phi_G : \cup_{t=1}^{\infty} H'_t \to \mathbb{R}$ be the GEQ error metric, where each $H'_t = \mathcal{G}^t \times (\mathbb{R}^d)^t$. That is, for $T \geq 1$,

$$\Phi_G(g_{1:T}, \theta_{1:T}) = \left\| \sum_{t=1}^{T} g_t(\theta_t) \right\|_2,$$

  where we abbreviate $g_{1:T} = (g_1, \ldots, g_T)$ and $\theta_{1:T} = (\theta_1, \ldots, \theta_T)$.

At each integer round $t \geq 1$,

1. the learner selects $\theta_t \in \mathbb{R}^d$;
2. the adversary selects $g_t \in \mathcal{G}$;
3. the learner observes $g_t(\theta_t)$.

An algorithm for the learner in the problem $(\mathbb{R}^d, \mathcal{G})$ is a mapping $\mathcal{A}_G : \{\emptyset\} \cup (\cup_{t=1}^{\infty} (\mathbb{R}^d)^t) \to \mathbb{R}^d$ that has knowledge of $\mathbb{R}^d$ but not $\mathcal{G}$. At each round $t \geq 1$, $\mathcal{A}_G$ selects a decision $\theta_t \in \mathbb{R}^d$ according to

$$\theta_t = \mathcal{A}_G(g_1(\theta_1), \ldots, g_{t-1}(\theta_{t-1}))$$

with input $\emptyset$ when $t = 1$. $\mathcal{A}_G$ guarantees that the resulting sequence $\theta_{1:T}$ incurs error $\Phi_G(g_{1:T}, \theta_{1:T}) = O(\phi_G(T))$, where ϕ_G is a positive and increasing error rate function which satisfies $\phi_G(T) = o(T)$.

1.3 Blackwell approachability

A conic *Blackwell approachability* (BA) problem is a sequential game between a learner and an adversary parameterized by $(\mathcal{X}, \mathcal{Y}, f, S)$, defined as follows. Let $L \geq 0$ be a constant, and let:

- $\mathcal{X}$ be the action set of the learner, and $\mathcal{Y}$ the action set of the adversary;
- $f : \mathcal{X} \times \mathcal{Y} \to \mathbb{R}^d$ be a vector-valued payoff function which is L-bounded in the Euclidean norm;
- $S \subseteq \mathbb{R}^d$ be a nonempty closed convex target cone;
- $(\mathcal{X}, \mathcal{Y}, f, S)$ satisfy Blackwell's condition: for all $u \in \mathbb{R}^d$, there exists $x \in \mathcal{X}$ such that for all $y \in \mathcal{Y}$,

$$\langle u - \Pi_S(u), f(x, y) - \Pi_S(u) \rangle \leq 0,$$

  where $\Pi_S(u) = \arg\min_{s \in S} \|u - s\|_2$;

- $\Phi_B : \cup_{t=1}^{\infty} H''_t \to \mathbb{R}$ be the BA error metric, where each $H''_t = \mathcal{X}^t \times \mathcal{Y}^t$. That is, for $T \geq 1$,

$$\Phi_B(x_{1:T}, y_{1:T}) = \inf_{s \in S} \left\| \frac{1}{T} \sum_{t=1}^{T} f(x_t, y_t) - s \right\|_2,$$

  where we abbreviate $x_{1:T} = (x_1, \ldots, x_T)$, and $y_{1:T} = (y_1, \ldots, y_T)$.

At each integer round $t \geq 1$,

1. the learner selects $x_t \in \mathcal{X}$;
2. the adversary selects $y_t \in \mathcal{Y}$;
3. the learner observes $f(x_t, y_t)$.

An algorithm for the learner in the problem $(\mathcal{X}, \mathcal{Y}, f, S)$ is a mapping $\mathcal{A}_B : \{\emptyset\} \cup (\cup_{t\geq 1}(\mathbb{R}^d)^t) \to \mathcal{X}$ that has knowledge of $\mathcal{X}$, $\mathcal{Y}$, f, and S but not T. At each round $t \geq 1$, $\mathcal{A}_B$ selects a decision $x_t \in \mathcal{X}$ according to

$$x_t = \mathcal{A}_B(f(x_1, y_1), \ldots, f(x_{t-1}, y_{t-1}))$$

with input $\emptyset$ when $t = 1$. $\mathcal{A}_B$ guarantees that the resulting sequence $x_{1:T}$ incurs error $\Phi_B(x_{1:T}, y_{1:T}) = O(\phi_B(T)/T)$, where ϕ_B is a positive and increasing error rate function which satisfies $\phi_B(T) = o(T)$.

2 Reductions

We will consider BA algorithms of the form $\mathcal{A}_B = O \circ \mathcal{A}'_B$, where $\mathcal{A}'_B : \{\emptyset\} \cup (\cup_{t\geq 1}(\mathbb{R}^d)^t) \to S^\circ$ maps the history of payoffs $f(x_1, y_1), \ldots, f(x_{t-1}, y_{t-1})$ to an approach direction vector $u_t \in S^\circ$ for all $1 \leq t \leq T$, and O maps the approach direction vector $u_t \in S^\circ$ to the action $x_t \in \mathcal{X}$ that satisfies Blackwell's condition. When we reduce a GEQ problem to a BA problem, we assume we have access to $\mathcal{A}'_B$ and implement O using elementary operations. When we reduce a BA problem to a GEQ problem, we use calls to $\mathcal{A}_G$ and elementary operations to construct $\mathcal{A}'_B$. We assume query access to O, which is guaranteed to exist by the assumption that $(\mathcal{X}, \mathcal{Y}, f, S)$ satisfies Blackwell's condition.

2.1 Reducing GEQ to conic BA

A GEQ problem parameterized by $(\mathbb{R}^d, \mathcal{G})$ can be equivalently formulated as a conic BA problem parameterized by $(\mathcal{X}, \mathcal{Y}, f, S)$, where

$$\mathcal{X} = \mathbb{R}^d, \quad \mathcal{Y} = \mathcal{G}, \quad f(\theta, g) = -g(\theta), \quad S = \{0\}. \tag{1}$$

Lemma 1. *Suppose $(\mathbb{R}^d, \mathcal{G})$ is a valid GEQ problem. Then $(\mathcal{X}, \mathcal{Y}, f, S)$, as defined in (1), is a valid conic BA problem.*

Proof. It suffices to show that f is L-bounded in the Euclidean norm, S is a nonempty closed convex cone, and Blackwell's condition holds. First, since each $g \in \mathcal{G}$ is L-bounded in the Euclidean norm, $f(\theta, g) = -g(\theta)$ is L-bounded in the Euclidean norm. Next, $S = \{0\}$ is trivially nonempty, closed, convex, and a cone. Finally, to see that Blackwell's condition holds, consider three cases. If $u = 0$, the condition holds for any choice of θ; so we can set $\theta = u$. If $h = 0$, we can set $\theta = u$ and the condition holds. If $h > 0$ and $u \in \mathbb{R}^d \setminus \{0\}$, we can set $\theta = h\frac{u}{\|u\|_2} \in \mathbb{R}^d$ and this ensures that for all $g \in \mathcal{G}$,

$$\langle u, -g(\theta)\rangle = \frac{\|u\|_2}{h} \cdot \langle \theta, -g(\theta)\rangle \leq 0,$$

where the inequality holds by the nonnegativity of $\frac{\|u\|_2}{h}$ and the restorative property of g. □

From this point, the only barrier to constructing an algorithm is a mismatch in the learner's a priori knowledge between frameworks. A learner (and hence their algorithm) in a GEQ problem has no knowledge of restorativity radius h, while a learner (and hence their algorithm) in the BA problem has full knowledge of

Algorithm 1 GEQ from BA

Require: BA algorithm $\mathcal{A}'_B$

1: Initialize $s = 0$, $\alpha_1 = 1$
2: **for** $t = 1, 2, \ldots$ **do**
3:   **if** $s = t - 1$ **then**
4:     Set $u_t = \mathcal{A}'_B(\emptyset) \in \mathbb{R}^d$
5:   **else**
6:     Set $u_t = \mathcal{A}'_B(-g_{s+1}(\theta_{s+1}), \ldots, -g_{t-1}(\theta_{t-1})) \in \mathbb{R}^d$
7:   **end if**
8:   Play $\theta_t = \alpha_t u_t / \|u_t\|_2 \in \mathbb{R}^d$ // Play $\theta_t = 0$ if $u_t = 0$
9:   Observe $g_t(\theta_t) \in \mathbb{R}^d$
10:   **if** $\langle g_t(\theta_t), \theta_t \rangle \geq 0$ **then**
11:     Set $\alpha_{t+1} = \alpha_t$
12:   **else**
13:     Set $\alpha_{t+1} = 2\alpha_t$ and $s = t$
14:   **end if**
15: **end for**

$\mathcal{Y} = \mathcal{G}$, including the radius h. To get around this discrepancy, Algorithm 1 iteratively guesses and checks the value of h while making black-box queries to an algorithm $\mathcal{A}'_B$ for BA.

Lemma 2. *Suppose $(\mathbb{R}^d, \mathcal{G})$ is a valid GEQ problem and $\mathcal{A}'_B$ is a BA algorithm with error rate function ϕ_B. Fix $T \geq 1$. Suppose that Line 10 in Algorithm 1 is active at rounds $1 \leq t_1 < t_2 < \cdots < t_m \leq T$. Let $t_0 = 0$, $t_{m+1} = T + 1$, and $T_j = t_j - t_{j-1} - 1$. Selecting a sequence of decisions $\theta_1, \ldots, \theta_T \in \mathbb{R}^d$ using Algorithm 1 ensures that for any sequence of vector fields $g_1, \ldots, g_T \in \mathcal{G}$,*

$$\Phi_G(g_{1:T}, \theta_{1:T}) \leq \sum_{j=1}^{m+1} T_j \cdot \Phi_B(\theta_{t_{j-1}+1:t_j-1}, g_{t_{j-1}+1:t_j-1}) + Lm = O(\phi_B(T)).$$

Proof. The s variable in Algorithm 1 records the last round at which the approachability algorithm $\mathcal{A}'_B$ was reset. Concretely, at each iteration $1 \leq t \leq T$, Algorithm 1 queries $\mathcal{A}'_B$ with $-g_{s+1}, \ldots, -g_{t-1}$, the history of realized vectors since the last time the algorithm was restarted. The algorithm $\mathcal{A}'_B$ returns an approach direction u_t, which Algorithm 1 translates into an action with the rule $\theta_t = \alpha_t u_t / \|u_t\|_2$. After observing the realized vector field g_t, Algorithm 1 tests whether the restorativity condition held. That is, if $\langle g_t(\theta_t), \theta_t \rangle \geq 0$, the current step size is maintained in the next round, $\alpha_{t+1} = \alpha_t$. Otherwise, $\|\theta_t\|_2 = \alpha_t < h$ must hold, so the step size is doubled for the next round, $\alpha_{t+1} = 2\alpha_t$, and $\mathcal{A}'_B$ is restarted by updating $s = t$. Clearly, this test fails on at most $m \leq \lceil \log_2(\max\{1, h\}) \rceil$ rounds. Hence

$$\begin{aligned}
\Phi_G(g_{1:T}, \theta_{1:T}) &= \left\| \sum_{t=1}^{T} g_t(\theta_t) \right\|_2 \\
&\leq \sum_{j=1}^{m+1} \left\| \sum_{t=t_{j-1}+1}^{t_j-1} g_t(\theta_t) \right\|_2 + \sum_{j=1}^{m} \|g_{t_j}(\theta_{t_j})\|_2 \\
&\leq \sum_{j=1}^{m+1} T_j \cdot \Phi_B(\theta_{t_{j-1}+1:t_j-1}, g_{t_{j-1}+1:t_j-1}) + Lm
\end{aligned}$$

$$
\begin{aligned}
&= \sum_{j=1}^{m+1} O(\phi_B(T_j)) + Lm \\
&= O(\phi_B(T))
\end{aligned}
$$

The first inequality holds by triangle inequality, and the second inequality holds because by Lemma 1, a valid conic BA problem is solved between restarts. The third inequality holds by assumption on $\mathcal{A}'_B$, and the final line holds because m is a constant and ϕ_B is increasing. $\square$

2.2 Reducing conic BA to GEQ

A conic BA problem parameterized by $(\mathcal{X}, \mathcal{Y}, f, S)$ can be equivalently formulated as a GEQ problem parameterized by $(\mathbb{R}^d, \mathcal{G})$, where

$$
\mathcal{G} = \left\{ g_y : \mathbb{R}^d \to \mathbb{R}^d \ : \ \text{for all } \theta \in \mathbb{R}^d,\ g_y(\theta) = -f(\mathcal{O}(\Pi_{S^\circ}(\theta)), y) + n_y(\theta), \text{ for some } y \in \mathcal{Y} \right\}, \tag{2}
$$

$$
n_y(\theta) = \max\left\{0, \frac{\langle f(\mathcal{O}(\Pi_{S^\circ}(\theta)), y), \theta - \Pi_{S^\circ}(\theta)\rangle}{\|\theta - \Pi_{S^\circ}(\theta)\|_2^2}\right\} \cdot (\theta - \Pi_{S^\circ}(\theta)), \tag{3}
$$

and $S^\circ = \{z \in \mathbb{R}^d \ : \ \langle z, s\rangle \le 0, \text{ for all } s \in S\}$ is the polar of S. We take $n_y(\theta)$ to be 0 if $\theta = \Pi_{S^\circ}(\theta)$.

Lemma 3. *Suppose $(\mathcal{X}, \mathcal{Y}, f, S)$ is a valid conic BA problem. Then $(\mathbb{R}^d, \mathcal{G})$, as defined in* (2), (3)*, is a valid GEQ problem.*

Proof. It suffices to show that for all $g \in \mathcal{G}$, g is $2L$-bounded in the Euclidean norm and satisfies the restorative property with $h = 0$. As f is L-bounded, if $\theta \in S^\circ$, then $\|g(\theta)\|_2 = \|f(\mathcal{O}(\Pi_{S^\circ}(\theta)), y)\|_2 \le L$. Otherwise, by the triangle inequality and Cauchy-Schwarz,

$$
\begin{aligned}
\|g(\theta)\|_2 &= \left\| - f(\mathcal{O}(\Pi_{S^\circ}(\theta)), y) + \max\left\{0, \frac{\langle f(\mathcal{O}(\Pi_{S^\circ}(\theta)), y), \theta - \Pi_{S^\circ}(\theta)\rangle}{\|\theta - \Pi_{S^\circ}(\theta)\|_2^2}\right\} \cdot (\theta - \Pi_{S^\circ}(\theta)) \right\|_2 \\
&\le \|f(\mathcal{O}(\Pi_{S^\circ}(\theta)), y)\|_2 + \frac{\|f(\mathcal{O}(\Pi_{S^\circ}(\theta)), y)\|_2 \cdot \|\theta - \Pi_{S^\circ}(\theta)\|_2}{\|\theta - \Pi_{S^\circ}(\theta)\|_2^2} \cdot \|\theta - \Pi_{S^\circ}(\theta)\|_2 \\
&= 2\|f(\mathcal{O}(\Pi_{S^\circ}(\theta)), y)\|_2 \\
&\le 2L,
\end{aligned}
$$

To see that the restorative property holds: for all $g \in \mathcal{G}$ and $\theta \in \mathbb{R}^d$,

$$
\langle g(\theta), \theta\rangle = \langle g(\theta), \theta - \Pi_{S^\circ}(\theta)\rangle + \langle g(\theta), \Pi_{S^\circ}(\theta)\rangle.
$$

If $\theta \in S^\circ$, then the first term is zero. Otherwise, expanding the first term yields

$$
\begin{aligned}
&\langle -g(\theta), \theta - \Pi_{S^\circ}(\theta)\rangle \\
&= \langle f(\mathcal{O}(\Pi_{S^\circ}(\theta)), y), \theta - \Pi_{S^\circ}(\theta)\rangle - \max\left\{0, \frac{\langle f(\mathcal{O}(\Pi_{S^\circ}(\theta)), y), \theta - \Pi_{S^\circ}(\theta)\rangle}{\|\theta - \Pi_{S^\circ}(\theta)\|_2^2}\right\} \cdot \langle \theta - \Pi_{S^\circ}(\theta), \theta - \Pi_{S^\circ}(\theta)\rangle.
\end{aligned}
$$

If $\langle f(\mathcal{O}(\Pi_{S^\circ}(\theta)), y), \theta - \Pi_{S^\circ}(\theta)\rangle \le 0$, the above is at most zero. Otherwise, it is exactly zero. Hence

$$
\begin{aligned}
&\langle -g(\theta), \theta\rangle \\
&\le \langle -g(\theta), \Pi_{S^\circ}(\theta)\rangle \\
&\le \langle f(\mathcal{O}(\Pi_{S^\circ}(\theta)), y), \Pi_{S^\circ}(\theta)\rangle - \max\left\{0, \frac{\langle f(\mathcal{O}(\Pi_{S^\circ}(\theta)), y), \theta - \Pi_{S^\circ}(\theta)\rangle}{\|\theta - \Pi_{S^\circ}(\theta)\|_2^2}\right\} \cdot \langle \theta - \Pi_{S^\circ}(\theta), \Pi_{S^\circ}(\theta)\rangle.
\end{aligned}
$$

Algorithm 2 BA from GEQ

Require: GEQ algorithm $\mathcal{A}_G$, Blackwell condition oracle $\mathcal{O}$

for $t = 1, 2, \ldots$ **do**
  if $t = 1$ **then**
    Set $\theta_t = \mathcal{A}_G(\emptyset) \in \mathbb{R}^d$
  else
    Set $\theta_t = \mathcal{A}_G(g_1(\theta_1), \ldots, g_{t-1}(\theta_{t-1})) \in \mathbb{R}^d$
  end if
  Play $x_t = \mathcal{O}(\Pi_{S^\circ}(\theta_t)) \in \mathcal{X}$
  Observe $f(x_t, y_t) \in \mathbb{R}^d$
  Set $n_t(\theta_t) = \max\left\{0, \frac{\langle f(x_t,y_t), \theta_t - \Pi_{S^\circ}(\theta_t)\rangle}{\|\theta_t - \Pi_{S^\circ}(\theta_t)\|_2^2}\right\} \cdot (\theta_t - \Pi_{S^\circ}(\theta_t)) \in \mathbb{R}^d$ // Set $n_t(\theta_t) = 0$ if $\theta_t = \Pi_{S^\circ}(\theta_t)$
  Set $g_t(\theta_t) = -f(x_t, y_t) + n_t(\theta_t) \in \mathbb{R}^d$
end for

To bound the second term above, recall that by the first-order optimality condition for Euclidean projection, for all $z \in S^\circ$, $\langle \theta - \Pi_{S^\circ}(\theta), \Pi_{S^\circ}(\theta) - z\rangle \geq 0$. Applying this inequality to $z = 0 \in S^\circ$, we see that the second term is at most zero. Hence

$$\langle -g(\theta), \theta\rangle \leq \langle f(\mathcal{O}(\Pi_{S^\circ}(\theta)), y), \Pi_{S^\circ}(\theta)\rangle.$$

Denote $u = \Pi_{S^\circ}(\theta)$. Since u is in the polar set of S, its projection onto S is $\Pi_S(u) = 0$. Thus the above can be rewritten as

$$\langle -g(\theta), \theta\rangle \leq \langle u - \Pi_S(u), f(\mathcal{O}(u), y) - \Pi_S(u)\rangle \leq 0,$$

where the second inequality holds by Blackwell's condition. □

Lemma 4. *Suppose $(\mathcal{X}, \mathcal{Y}, f, S)$ is a valid conic BA problem and $\mathcal{A}_G$ is a GEQ algorithm with error rate function ϕ_G. Fix $T \geq 1$. Selecting a sequence of approach directions $\theta_1, \ldots, \theta_T$ and corresponding actions $x_1, \ldots, x_T \in \mathcal{X}$ using Algorithm 2 ensures that for any sequence of actions $y_1, \ldots, y_T \in \mathcal{Y}$,*

$$\Phi_B(x_{1:T}, y_{1:T}) \leq \frac{\Phi_G(g_{1:T}, \theta_{1:T})}{T} = O\left(\frac{\phi_G(T)}{T}\right).$$

Proof. Recall that $x_t = \mathcal{O}(\Pi_{S^\circ}(\theta_t))$ and observe that

$$\begin{aligned}
&\Phi_B(x_{1:T}, y_{1:T}) \\
&= \inf_{s\in S}\left\|\frac{1}{T}\sum_{t=1}^T f(x_t, y_t) - s\right\|_2 \\
&\leq \left\|\frac{1}{T}\sum_{t=1}^T f(x_t, y_t) - \frac{1}{T}\sum_{t=1}^T n_t(\theta_t)\right\|_2 \\
&= \left\|\frac{1}{T}\sum_{t=1}^T\left(-f(\mathcal{O}(\Pi_{S^\circ}(\theta_t)), y_t) + \max\left\{0, \frac{\langle f(\mathcal{O}(\Pi_{S^\circ}(\theta_t)), y_t), \theta_t - \Pi_{S^\circ}(\theta_t)\rangle}{\|\theta_t - \Pi_{S^\circ}(\theta_t)\|_2^2}\right\} \cdot (\theta_t - \Pi_{S^\circ}(\theta_t))\right)\right\|_2 \\
&= \frac{\|\sum_{t=1}^T g_t(\theta_t)\|_2}{T} \\
&= \frac{\Phi_G(g_{1:T}, \theta_{1:T})}{T}.
\end{aligned}$$

Algorithm 3 RM from BA

Require: BA algorithm $\mathcal{A}'_B$

for $t = 1, 2, \ldots$ **do**
  if t=1 **then**
    Set $u_t = \mathcal{A}'_B(\emptyset) \in \mathbb{R}^{d+1}$
  else
    Set $u_t = \mathcal{A}'_B(f(z_1, \ell_1), \ldots, f(z_{t-1}, \ell_{t-1})) \in \mathbb{R}^{d+1}$
  end if
  Play $z_t = (u_t)_{2:(d+1)}/(u_t)_1 \in \mathcal{Z}$ // Play arbitrary $z_0 \in \mathcal{Z}$ if $(u_t)_1 = 0$
  Observe $g_t(z_t) \in \mathbb{R}^d$
  Set $f(z_t, \ell_t) = \langle g_t(z_t), z_t \rangle \oplus -g_t(z_t) \in \mathbb{R}^{d+1}$
end for

To explain the first inequality above, we first claim that $n_t(\theta_t) \in S$. To see this, note that $n_t(\theta_t) \in N_{S^\circ}(\Pi_{S^\circ}(\theta_t))$ by construction, where $N_{S^\circ}(z)$ denotes the normal cone of S° at z; furthermore, we have $N_{S^\circ}(\Pi_{S^\circ}(\theta_t)) \subseteq S$, as verified in Lemma 5 below. Thus $n_t(\theta_t) \in S$, together with convexity of S, verifies the first inequality above. The remaining lines follow from definitions. By Lemma 3, $g_1, \ldots, g_T \in \mathcal{G}$ are $2L$-bounded and restorative with $h = 0$. Therefore, $\mathcal{A}_G$ ensures that $\Phi_G(g_{1:T}, \theta_{1:T}) = O(\phi_G(T))$, and $\Phi_B(x_{1:T}, y_{1:T}) = O(\phi_G(T)/T)$. □

Lemma 5. *Suppose S is a closed and convex cone, and S° is its polar. Denote the normal cone of S° at $z \in S^\circ$ by $N_{S^\circ}(z) = \{s \in \mathbb{R}^d \,:\, \langle s, z - z' \rangle \geq 0, \text{ for all } z' \in S^\circ\}$. Then for all $z \in S^\circ$, it holds that $N_{S^\circ}(z) \subseteq S$.*

Proof. Fix any $z \in S^\circ$. For any $s \in N_{S^\circ}(z)$, $\langle s, z' \rangle \leq \langle s, z \rangle$ holds for all $z' \in S^\circ$. Since S° is a cone, $0 \in S^\circ$ and $2z \in S^\circ$. Instantiating the last inequality with $z' = 0$ and $z' = 2z$ shows that $\langle s, z \rangle = 0$. Hence, $\langle s, z' \rangle \leq 0$ for all $z' \in S^\circ$, which shows $s \in (S^\circ)^\circ$. Since S is closed and convex, $(S^\circ)^\circ = S$. □

2.3 Conic BA and RM are equivalent

We collect relevant results from Abernethy et al. (2011) as propositions.

First, we recall that RM can be reduced to conic BA. For this, it helps to introduce additional notation: for $x \in \mathbb{R}$ and $y \in \mathbb{R}^d$, denote $x \oplus y = (x, y^\top)^\top \in \mathbb{R}^{d+1}$. A RM problem parameterized by $(\mathbb{R}^d, \mathcal{Z}, \mathcal{L})$ can be equivalently formulated as a conic BA problem parameterized by $(\mathcal{X}, \mathcal{Y}, f, S)$, where

$$\mathcal{X} = \mathcal{Z}, \quad \mathcal{Y} = \mathcal{L}, \quad f(z, \ell) = \langle g(z), z \rangle \oplus -g(z), \quad S = \operatorname{cone}(\{1\} \times \mathcal{Z})^\circ. \tag{4}$$

Proposition 1. *Suppose $(\mathbb{R}^d, \mathcal{Z}, \mathcal{L})$ is a valid RM problem. Then $(\mathcal{X}, \mathcal{Y}, f, S)$, as defined in* (4)*, is a valid conic BA problem. Now suppose $\mathcal{A}'_B$ is a BA algorithm with error rate function ϕ_B. Fix $T \geq 1$. Selecting a sequence of decisions $z_1, \ldots, z_T \in \mathcal{Z}$ using Algorithm 3 ensures that for any sequence of losses $\ell_1, \ldots, \ell_T \in \mathcal{L}$,*

$$\Phi_R(\ell_{1:T}, z_{1:T}) \leq cT \cdot \Phi_B(z_{1:T}, \ell_{1:T}) = O(\phi_B(T)),$$

where $c > 0$ may depend on $\mathcal{Z}$ but is independent of T.

Next, recall that conic BA can be reduced to RM. A conic BA problem parameterized by $(\mathcal{X}, \mathcal{Y}, f, S)$ can be equivalently formulated as a RM problem parameterized by $(\mathbb{R}^d, \mathcal{Z}, \mathcal{L})$, where

$$\mathcal{Z} = S^\circ \cap B_2(1), \quad \mathcal{L} = \{\langle -f(O(z'), y), \cdot \rangle : z' \in \mathcal{Z},\ y \in \mathcal{Y}\}. \tag{5}$$

In the above, $B_2(1)$ denotes the Euclidean unit ball centered at the origin.

Algorithm 4 BA from RM

Require: RM algorithm $\mathcal{A}_R$, Blackwell condition oracle O

for $t = 1, 2, \ldots$ **do**
    if t=1 **then**
        Set $z_t = \mathcal{A}_R(\emptyset) \in \mathbb{R}^d$
    else
        Set $z_t = \mathcal{A}_R(g_1(z_1), \ldots, g_{t-1}(z_{t-1})) \in \mathbb{R}^d$
    end if
    Play $x_t = O(z_t) \in \mathcal{X}$
    Observe $f(x_t, y_t) \in \mathbb{R}^d$
    Set $g_t(z_t) = -f(x_t, y_t) \in \mathbb{R}^d$
end for

Algorithm 5 RM from GEQ

Require: GEQ algorithm $\mathcal{A}_G$

for $t = 1, 2, \ldots$ **do**
    if $t = 1$ **then**
        Set $\theta'_t = \mathcal{A}_G(\emptyset) \in \mathbb{R}^{d+1}$
    else
        Set $\theta'_t = \mathcal{A}_G(g'_1(\theta'_1), \ldots, g'_{t-1}(\theta'_{t-1})) \in \mathbb{R}^{d+1}$
    end if
    Set $\theta_t = \Pi_{S^\circ}(\theta'_t)$
    Play $z_t = (\theta_t)_{2:(d+1)}/(\theta_t)_1 \in \mathcal{Z}$ // Play arbitrary $z_0 \in \mathcal{Z}$ if $\theta_t = 0$
    Observe $g_t(z_t) \in \mathbb{R}^d$
    Set $f(z_t, \ell_t) = \langle g_t(z_t), z_t \rangle \oplus -g_t(z_t) \in \mathbb{R}^{d+1}$
    Set $n_t(\theta'_t) = \max\left\{0, \frac{\langle f(z_t,\ell_t), \theta'_t - \Pi_{S^\circ}(\theta'_t)\rangle}{\|\theta'_t - \Pi_{S^\circ}(\theta'_t)\|_2^2}\right\} \cdot (\theta'_t - \Pi_{S^\circ}(\theta'_t)) \in \mathbb{R}^{d+1}$ // Set $n_t(\theta'_t) = 0$ if $\theta'_t = \Pi_{S^\circ}(\theta'_t)$
    Set $g'_t(\theta'_t) = -f(z_t, \ell_t) + n_t(\theta'_t) \in \mathbb{R}^{d+1}$
end for

Proposition 2. *Suppose $(\mathcal{X}, \mathcal{Y}, f, S)$ is a valid conic BA problem. Then $(\mathbb{R}^d, \mathcal{Z}, \mathcal{L})$, as defined in (5), is a valid RM problem. Now suppose $\mathcal{A}_R$ is a RM algorithm with error rate function ϕ_R. Fix $T \geq 1$. Selecting a sequence of approach directions $z_1, \ldots, z_T \in S^\circ \cap B_2(1)$ and corresponding actions $x_1, \ldots, x_T \in \mathcal{X}$ using Algorithm 4 ensures that for any sequence of actions $y_1, \ldots, y_T \in \mathcal{Y}$,*

$$\Phi_B(x_{1:T}, y_{1:T}) \leq \frac{\Phi_R(\ell_{1:T}, z_{1:T})}{T} = O\left(\frac{\phi_R(T)}{T}\right).$$

2.4 GEQ and RM are equivalent

We state our main results. Theorems 6 and 7 prove an algorithmic equivalence between RM and GEQ.

Theorem 6. *Suppose $(\mathbb{R}^d, \mathcal{Z}, \mathcal{L})$ is a valid RM problem and $\mathcal{A}_G$ is a GEQ algorithm with error rate function ϕ_G. Fix $T \geq 1$. Selecting a sequence of decisions $z_1, \ldots, z_T \in \mathcal{Z}$ using Algorithm 5 ensures that for any sequence of losses $\ell_1, \ldots, \ell_T \in \mathcal{L}$,*

$$\Phi_R(\ell_{1:T}, z_{1:T}) \leq c\Phi_G(g'_{1:T}, \theta'_{1:T}) = O(\phi_G(T)),$$

where $c > 0$ may depend on $\mathcal{Z}$ but is independent of T.

Algorithm 6 GEQ from RM

Require: RM algorithm $\mathcal{A}_R$

1: Initialize $s = 0$, $\alpha_1 = 1$
2: **for** $t = 1, 2, \ldots$ **do**
3:   **if** $s = t - 1$ **then**
4:     Set $z_t = \mathcal{A}_R(\emptyset) \in \mathbb{B}_2(1)$
5:   **else**
6:     Set $z_t = \mathcal{A}_R(g_{s+1}(\theta_{s+1}), \ldots, g_{t-1}(\theta_{t-1})) \in \mathbb{B}_2(1)$
7:   **end if**
8:   Play $\theta_t = \alpha_t z_t / \|z_t\|_2 \in \mathbb{R}^d$ // Play $\theta_t = 0$ if $z_t = 0$
9:   Observe $g_t(\theta_t) \in \mathbb{R}^d$
10:   **if** $\langle g_t(\theta_t), \theta_t \rangle \geq 0$ **then**
11:     Set $\alpha_{t+1} = \alpha_t$
12:   **else**
13:     Set $\alpha_{t+1} = 2\alpha_t$ and $s = t$
14:   **end if**
15: **end for**

Proof. A RM problem $(\mathbb{R}^d, \mathcal{Z}, \mathcal{L})$ can be formulated as a conic BA problem $(\mathcal{X}, \mathcal{Y}, f, S)$, as defined in (4). Algorithm 5 specifies how Algorithm 2 can be used to solve this choice of $(\mathcal{X}, \mathcal{Y}, f, S)$. Hence,

$$\Phi_R(\ell_{1:T}, z_{1:T}) \leq cT \cdot \Phi_B(z_{1:T}, \ell_{1:T}) \leq c\Phi_G(g'_{1:T}, \theta'_{1:T}) = O(\phi_G(T)).$$

The first inequality holds by Proposition 1, the second inequality by Lemma 4, and the final equality by the assumption on $\mathcal{A}_G$. □

Theorem 7. *Suppose $(\mathbb{R}^d, \mathcal{G})$ is a valid GEQ problem and $\mathcal{A}_R$ is a RM algorithm with error rate function ϕ_R. Fix $T \geq 1$. Suppose that Line 10 in Algorithm 6 is active at rounds $1 \leq t_1 < t_2 < \cdots < t_m \leq T$. Let $t_0 = 0$, $t_{m+1} = T + 1$, and $T_j = t_j - t_{j-1} - 1$. Selecting a sequence of decisions $\theta_1, \ldots, \theta_T \in \mathbb{R}^d$ using Algorithm 6 ensures that for any sequence of vector fields $g_1, \ldots, g_T \in \mathcal{G}$,*

$$\Phi_G(g_{1:T}, \theta_{1:T}) \leq \sum_{j=1}^{m+1} \Phi_R(\ell_{t_{j-1}+1:t_j-1}, z_{t_{j-1}+1:t_j-1}) + Lm = O(\phi_R(T)).$$

Proof. A GEQ problem $(\mathbb{R}^d, \mathcal{G})$ can be formulated as a conic BA problem $(\mathcal{X}, \mathcal{Y}, f, S)$, as defined in (1). Algorithm 6 is a specialization of Algorithm 1 where $\mathcal{A}'_B$ is instantiated with $\mathcal{A}_R$, identically to how $\mathcal{A}'_B$ is instantiated with $\mathcal{A}_R$ in Algorithm 4. By Lemma 2, we have that

$$\Phi_G(g_{1:T}, \theta_{1:T}) \leq \sum_{j=1}^{m+1} T_j \cdot \Phi_B(\theta_{t_{j-1}+1:t_j-1}, g_{t_{j-1}+1:t_j-1}) + Lm.$$

Between each restart, Algorithm 6 is equivalent to Algorithm 4, where Line 8 of Algorithm 6 corresponds to Line 7 of Algorithm 4. Hence, by Proposition 2, for each $1 \leq j \leq m + 1$,

$$T_j \cdot \Phi_B(\theta_{t_{j-1}+1:t_j-1}, g_{t_{j-1}+1:t_j-1}) \leq \Phi_R(\ell_{t_{j-1}+1:t_j-1}, z_{t_{j-1}+1:t_j-1}).$$

So we conclude that

$$\Phi_G(g_{1:T}, \theta_{1:T}) \leq \sum_{j=1}^{m+1} O(\phi_R(T_j)) + Lm = O(\phi_R(T)),$$

where the final equality holds because m is constant and ϕ_R is increasing in T. □

References

Jacob Abernethy, Peter L. Bartlett, and Elad Hazan. *Blackwell approachability and no-regret learning are equivalent*. In Proceedings of the 24th Annual Conference on Learning Theory, pages 27–46, 2011.