Sparse Regret Minimization

Joon Kwon
Université Pierre-et-Marie-Curie – Paris 6
Paris, France

Vianney Perchet
INRIA & Université Paris-Diderot – Paris 7
Paris, France

GEL Workshop

January 7th, 2016

Framework

- $d \geqslant 1$ integer.
- Set of pure actions for the player: $[d] := \{1, \ldots, d\}$.
- Set of mixed actions Δ_d.
- At stage $t = 1, \ldots, T$,
  - Player chooses action $x_t \in \Delta_d$.
  - Nature reveals gain vector $g_t \in [0, 1]^d$.
  - Player gets $\langle g_t | x_t \rangle$.
- A strategy/algorithm $\sigma = (\sigma_t)_{1 \leqslant t \leqslant T}$

$$x_t = \sigma_t(x_1, i_1, g_1, \ldots, x_{t-1}, i_{t-1}, g_{t-1}).$$

$$\text{Maximize:} \quad \sum_{t=1}^{T} \langle g_t | x_t \rangle$$

The Regret

$$R_T\left\{\sigma,(g_t)_t\right\} = R_T := \max_{i\in[d]} \sum_{t=1}^{T} g_t^{(i)} - \sum_{t=1}^{T} \langle g_t | x_t \rangle$$

A strategy σ *guarantees* $B(d,T)$ if:

$$\forall (g_t)_t, \quad R_T\left\{\sigma,(g_t)_t\right\} \leqslant B(d,T).$$

- **Introduced**: Hannan (1957)
- **Surveys**: Cesa-Bianchi–Lugosi (2006), Rakhlin–Tewari (2008), Shalev-Shwartz (2011), Hazan (2012), Bubeck–Cesa-Bianchi (2012),...

The Minimax Regret

- T: number of stages
- d: number of actions

$$\min_{\sigma} \max_{(g_t)_t} R_T \{\sigma, (g_t)_t\} \quad \text{is of order} \quad \sqrt{T \log d}$$

- **Upper bound**: Cesa-Bianchi (1997)
- **Lower bound**: Cesa-Bianchi, Freund, Haussler, Helmbold, Schapire, Warmuth (1997)

Mirror Descent Strategies

- $h : \mathbb{R}^d \longrightarrow \mathbb{R} \cup \{+\infty\}$ strictly convex, lsc, dom $h = \Delta_d$
- $\delta_h := \max_{\Delta_d} h - \min_{\Delta_d} h$.
- $\eta > 0$ a parameter.

$$x_t = \nabla h^* \left(\eta \sum_{s=1}^{t-1} g_s \right)$$

Theorem (Shalev-Shwartz (2007), Bubeck (2011), etc.)

If

- *h is K-strongly-convex wrt $\|\cdot\|$*
- *$\|g\|_* \leqslant M$ for all possible gain vector g*

$$R_T \leqslant \frac{\delta_h}{\eta} + \eta \frac{TM^2}{K} \quad \underset{\eta=\sqrt{\delta_h/(TM^2)}}{\rightsquigarrow} \quad \sqrt{T\delta_h/K} \cdot M$$

The Exponential Weight Algorithm achieves $\sqrt{T \log d}$

$$R_T \leqslant \sqrt{T\delta_h/K} \cdot M$$

$$h(x) = \begin{cases} \sum_{i=1}^{d} x^{(i)} \log x^{(i)} & \text{if } x \in \Delta_d \\ +\infty & \text{otherwise} \end{cases}$$

Exponential Weights Algorithm

$$x_t^{(i)} = \frac{\exp\left(\eta \sum_{s=1}^{t-1} g_s^{(i)}\right)}{\sum_{j=1}^{d} \exp\left(\eta \sum_{s=1}^{t-1} g_s^{(j)}\right)}$$

- $\delta_h = \log d$.
- h is 1-strongly convex wrt $\|\cdot\|_1$.
- $g \in [0,1]^d \quad \implies \quad \|g\|_\infty \leqslant 1$.

$$R_T \leqslant \sqrt{T \log d}.$$

Lower bound: a probabilistic argument

$$\min_{\sigma} \max_{(g_t)_t} \; R_T \left\{ \sigma, (g_t)_t \right\} \quad \gtrsim \quad \sqrt{T \log d} \text{ ?}$$

Fix a strategy σ_0. Let $(\tilde{g}_t)_t$ be i.i.d $\tilde{g}_t^{(i)} = 0$ or 1 (with prob. $(\frac{1}{2}, \frac{1}{2})$).

$$\max_{(g_t)_t} \; R_T \left\{ \sigma_0, (g_t)_t \right\} \geqslant R_T \left\{ \sigma_0, (\tilde{g}_t)_t \right\} \; = \; \max_i \sum_{t=1}^{T} \tilde{g}_t^{(i)} - \sum_{t=1}^{T} \langle \tilde{g}_t | x_t \rangle$$

$$\begin{aligned}
\min_{\sigma} \; \max_{(g_t)_t} \; R_T \left\{ \sigma_0, (g_t)_t \right\} &\geqslant \mathbb{E} \left[\max_i \sum_{t=1}^{T} \left(\tilde{g}_t^{(i)} - \frac{1}{2} \right) \right] \\
&= \sqrt{T} \cdot \mathbb{E} \left[\max_i \frac{1}{\sqrt{T}} \sum_{t=1}^{T} \left(\tilde{g}_t^{(i)} - \frac{1}{2} \right) \right] \\
X \sim \mathcal{N}(0, \tfrac{1}{4} I_d) \qquad\qquad &\sim \sqrt{T} \cdot \mathbb{E} \left[\max_i X^{(i)} \right] \\
&\sim \sqrt{T \log d}.
\end{aligned}$$

Gains and Losses are Equivalent

- Nature chooses loss vectors $\ell_t \in [0,1]^d$

$$\sum_{t=1}^{T} \langle \ell_t | x_t \rangle - \min_{i\in[d]} \sum_{t=1}^{T} \ell_t^{(i)}$$

- $g_t^{(i)} := 1 - \ell_t^{(i)}$
- $\ell_t \in [0,1]^d \implies g_t \in [0,1]^d$.

$$\max_{i\in[d]} \sum_{t=1}^{T} g_t^{(i)} - \sum_{t=1}^{T} \langle g_t | x_t \rangle = \sum_{t=1}^{T} \langle \ell_t | x_t \rangle - \min_{i\in[d]} \sum_{t=1}^{T} \ell_t^{(i)}$$

A Sparsity Assumption

Let $s \geqslant 1$ be an integer.

Assumption

All gain (resp. loss) vectors are s-sparse, i.e. have at most s nonzero components.

Example

$d = 3$ and $s = 1$.

$$g_1 = \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} \qquad g_2 = \begin{pmatrix} \frac{1}{2} \\ 0 \\ 0 \end{pmatrix} \qquad g_3 = \begin{pmatrix} 0 \\ \frac{1}{3} \\ 0 \end{pmatrix}$$

$$\ell_1 := \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix} - g_1 = \begin{pmatrix} 1 \\ 1 \\ 0 \end{pmatrix} \quad \rightsquigarrow \quad \text{not 1-sparse}$$

Minimax Regret for s-sparse Gains ?

$$\begin{pmatrix} s \text{ actions} \\ \text{no sparsity} \end{pmatrix} \underset{\text{easier}}{\leqslant} \begin{pmatrix} d \text{ actions} \\ \text{sparsity } s \end{pmatrix} \underset{\text{easier}}{\leqslant} \begin{pmatrix} d \text{ actions} \\ \text{no sparsity} \end{pmatrix}$$

$$\sqrt{T \log s} \quad \leqslant \quad \text{minimax regret} \quad \leqslant \quad \sqrt{T \log d}.$$

$$\sqrt{T \log s}$$

Upper bound for s-sparse Gains

$$\begin{cases} h \text{ is } K\text{-strongly convex wrt } \|\cdot\| \\ \forall g, \quad \|g\|_* \leqslant M \end{cases} \implies \quad R_T \leqslant \sqrt{T\delta_h/K} \cdot M$$

- $h_p(x) = \frac{1}{2}\|x\|_p^2 \quad (\text{if } x \in \Delta_d)$
- $\delta_h \leqslant 1$
- h_p is $(p-1)$-strongly convex wrt $\|\cdot\|_p$
- $\|g\|_q = \left(\sum_{i=1}^{d} \left|g^{(i)}\right|^q\right)^{1/q} \leqslant s^{1/q}$

$$R_T \leqslant \sqrt{\frac{T}{p-1}s^{1/q}}.$$

$$p = 1 + \frac{1}{2\log s - 1}$$

$$R_T \leqslant \sqrt{T \log s}$$

Minimax Regret for s-sparse Losses ?

Theorem (Littlestone–Warmuth (1994))

The Exponential Weights Algorithm against losses $\ell_t \in [0,1]^d$ guarantees:

$$R_T \quad \lesssim \quad \frac{\log d}{\eta} + \eta \cdot \min_{i \in [d]} \sum_{t=1}^{T} \ell_t^{(i)}.$$

$$sT \geqslant \sum_{t=1}^{T} \sum_{i=1}^{d} \ell_t^{(i)} = \sum_{i=1}^{d} \sum_{t=1}^{T} \ell_t^{(i)} \geqslant d \cdot \min_{i \in [d]} \sum_{t=1}^{T} \ell_t^{(i)}$$

$$R_T \quad \lesssim \quad \frac{\log d}{\eta} + \eta \frac{sT}{d} \quad = \quad \sqrt{Ts \frac{\log d}{d}}$$

Matching Lower Bound for s-sparse Losses

Define random i.i.d. s-sparse loss vectors $(\tilde{\ell}_t)_t$ as follows. For each $t \geqslant 1$

- Draw (uniformly) a subset I_t of $\{1, \ldots, d\}$ of cardinality s.
- Set

$$\tilde{\ell}_t^{(i)} = \begin{cases} 0 \text{ or } 1 & \text{if } i \in I_t \\ 0 & \text{if } i \notin I_t \end{cases}$$

$$\mathbb{E}\left[\tilde{\ell}_t^{(i)}\right] = \frac{s}{2d} \quad \textit{and} \quad \mathrm{Var}\,\tilde{\ell}_t^{(i)} = \frac{s}{2d}\left(1 - \frac{s}{2d}\right)$$

$$\begin{aligned}
\sqrt{T} \cdot \mathbb{E}\left[\max_{i \in [d]} \frac{1}{\sqrt{T}} \sum_{t=1}^{T} \left(\left\langle \tilde{\ell}_t \middle| x_t \right\rangle - \tilde{\ell}_t^{(i)}\right)\right] &\sim \sqrt{T} \cdot \mathbb{E}\left[\max_i X^{(i)}\right] \qquad (X \sim \mathcal{N}(0, \Sigma)) \\
&\geqslant \sqrt{T} \cdot \mathbb{E}\left[\max_i Y^{(i)}\right] \\
&\sim \sqrt{T} \cdot \sqrt{\log d} \cdot \sqrt{\frac{s}{2d}\left(1 - \frac{s}{2d}\right)} \\
&\gtrsim \sqrt{Ts\frac{\log d}{d}}
\end{aligned}$$

The Bandit Setting

For stages $t = 1, \ldots, T$,

- Player chooses action $i_t \in [d]$.
- Nature only reveals $g_t^{(i_t)}$.
- Player gets gain $g_t^{(i_t)}$.

Theorem

Minimax Regret is of order $\sqrt{Td}$

- **Upper bound**: Audibert and Bubeck (2009)
- **Lower bound**: Auer, Cesa-Bianchi, Freund and Schapire (2002)

Upper and Lower Bounds

Without sparsity: $\sqrt{Td}$

	Gains	Losses
Upper bound	$\sqrt{Td}$	$\sqrt{Ts\log\frac{d}{s}}$
Lower bound	$\sqrt{Ts}$	$\sqrt{Ts}$

- If the Player knows gain vectors are s-sparse, he can choose the right strategy to achieve $\sqrt{T \log s}$.
- What if is s is unknown ? Can he still take advantage of sparsity?
- The Player knows vectors are 1000-sparse. But if they actually turn out to be 10-sparse, ... ?

YES

Theorem (K. & Perchet (2015))

There exists a strategy which guarantees a $\sqrt{T \log s^*}$ *regret bound, where* $s^* = \max_{1 \leqslant t \leqslant T} \|g_t\|_0$.

- You don't know the sparsity level of the gain vectors.
- Just play the aforementionned strategy.
- If the gain vectors turn out to be s-sparse, then you will achieve:

$$R_T \lesssim \sqrt{T \log s}.$$

Analog result for losses