Minimax Policies for Combinatorial Prediction Games

Sébastien Bubeck[1]

joint work with Jean-Yves Audibert[2,3] and Gábor Lugosi[4]

[1] Centre de Recerca Matemàtica, Barcelona, Spain
[2] Imagine, Univ. Paris Est, and Sierra
[3] CNRS/ENS/INRIA, Paris, France
[4] ICREA and Pompeu Fabra University, Barcelona, Spain

Path planning

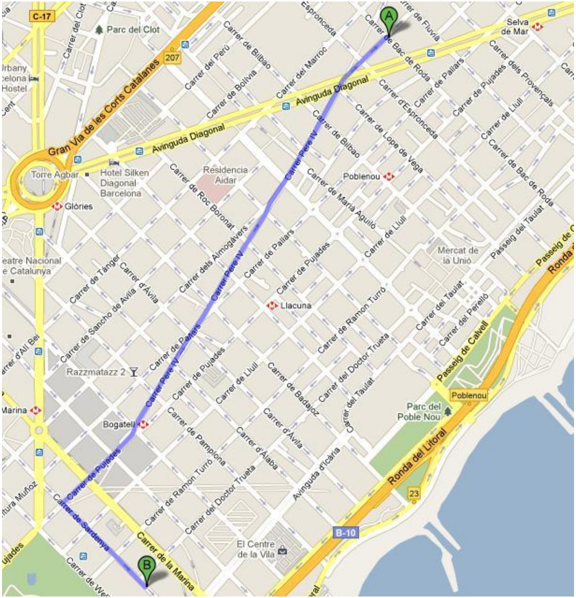

Combinatorial prediction game

Adversary

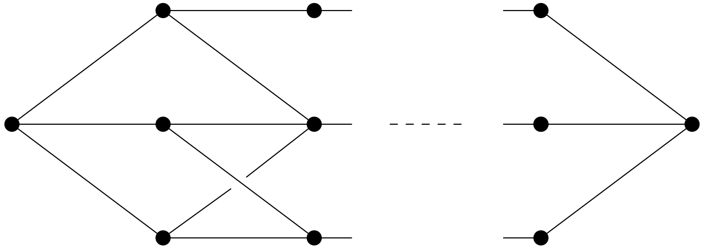

Player

Combinatorial prediction game

Adversary

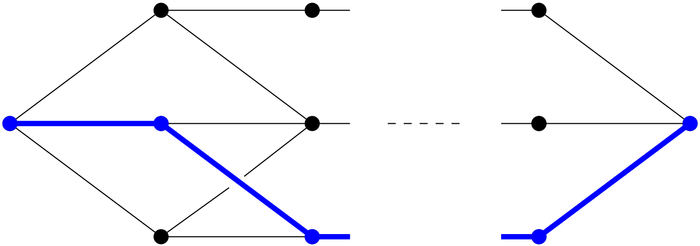

Player ⟶ 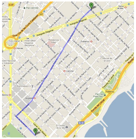

Combinatorial prediction game

Adversary ⟶

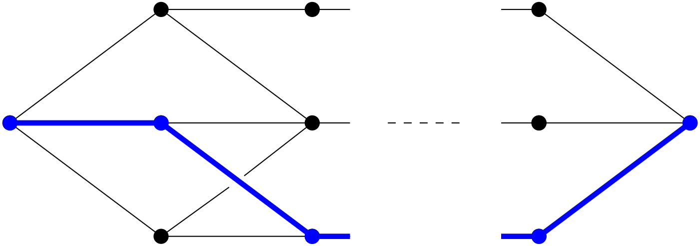

Player ⟶

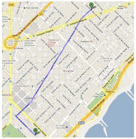

Combinatorial prediction game

Adversary ⟶

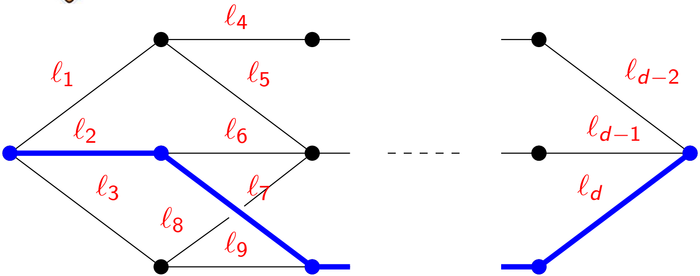

Player ⟶

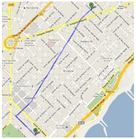

Combinatorial prediction game

Adversary ⟶ 

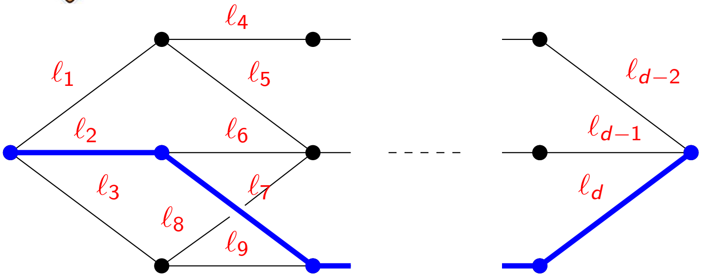

Player ⟶ 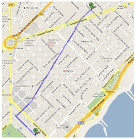

loss suffered: $\ell_2 + \ell_7 + \ldots + \ell_d$

Combinatorial prediction game

Adversary ⟶

Feedback: $\begin{cases} \text{Full Info: } \ell_1, \ell_2, \dots, \ell_d \\ \\ \end{cases}$

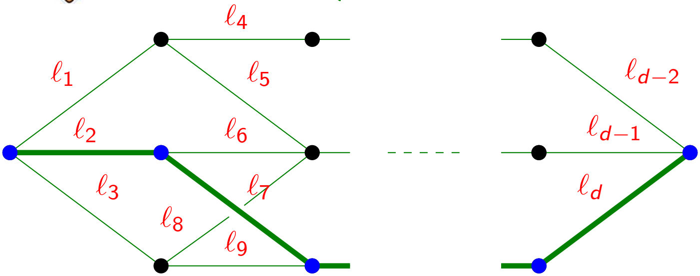

Player ⟶

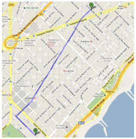

loss suffered: $\ell_2 + \ell_7 + \ldots + \ell_d$

Combinatorial prediction game

Adversary ⟶  Feedback:

- Full Info: $\ell_1, \ell_2, \dots, \ell_d$
- Semi-Bandit: $\ell_2, \ell_7, \dots, \ell_d$
- Bandit: $\ell_2 + \ell_7 + \dots + \ell_d$

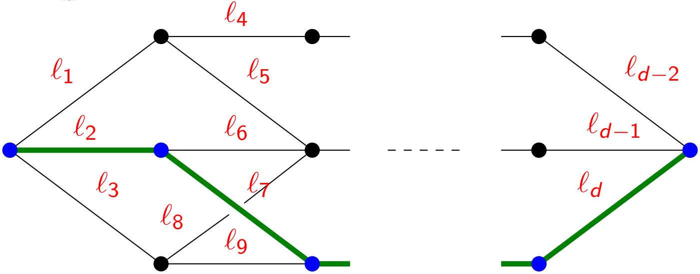

Player ⟶ 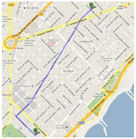

loss suffered: $\ell_2 + \ell_7 + \dots + \ell_d$

Combinatorial prediction game

Adversary ⟶ 

Feedback:

- Full Info: $\ell_1, \ell_2, \ldots, \ell_d$
- Semi-Bandit: $\ell_2, \ell_7, \ldots, \ell_d$
- Bandit: $\ell_2 + \ell_7 + \ldots + \ell_d$

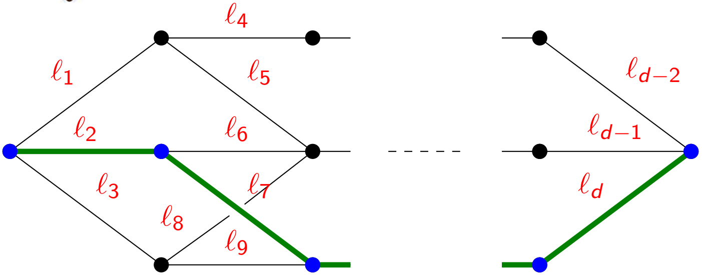

Player ⟶ 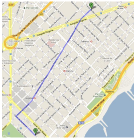

loss suffered: $\ell_2 + \ell_7 + \ldots + \ell_d$

Notation

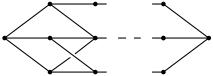

$\leftrightsquigarrow$ $\mathcal{S} \subset \{0,1\}^d$

$\leftrightsquigarrow$ $\ell_t \in \mathbb{R}_+^d$

$\leftrightsquigarrow$ $V_t \in \mathcal{S}$, loss suffered: $\ell_t^T V_t$

$$R_n = \mathbb{E} \sum_{t=1}^{n} \ell_t^T V_t - \min_{u \in S} \mathbb{E} \sum_{t=1}^{n} \ell_t^T u$$

Notation

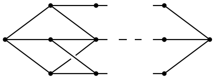

$\leftrightsquigarrow$ $\mathcal{S} \subset \{0,1\}^d$

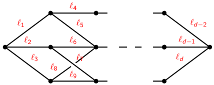

$\leftrightsquigarrow$ $\ell_t \in \mathbb{R}_+^d$

$\leftrightsquigarrow$ $V_t \in \mathcal{S}$, loss suffered: $\ell_t^T V_t$

$$R_n = \mathbb{E} \sum_{t=1}^{n} \ell_t^T V_t - \min_{u \in S} \mathbb{E} \sum_{t=1}^{n} \ell_t^T u$$

Notation

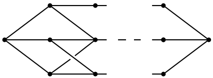

$\rightsquigarrow \quad \mathcal{S} \subset \{0,1\}^d$

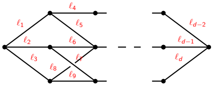

$\rightsquigarrow \quad \ell_t \in \mathbb{R}_+^d$

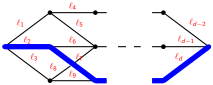

$\rightsquigarrow \quad V_t \in \mathcal{S}$, loss suffered: $\ell_t^T V_t$

$$R_n = \mathbb{E} \sum_{t=1}^{n} \ell_t^T V_t - \min_{u \in S} \mathbb{E} \sum_{t=1}^{n} \ell_t^T u$$

Notation

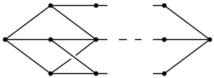

$\leftrightsquigarrow$ $\mathcal{S} \subset \{0,1\}^d$

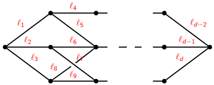

$\leftrightsquigarrow$ $\ell_t \in \mathbb{R}_+^d$

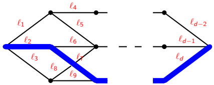

$\leftrightsquigarrow$ $V_t \in \mathcal{S}$, loss suffered: $\ell_t^T V_t$

$$R_n = \mathbb{E} \sum_{t=1}^{n} \ell_t^T V_t - \min_{u \in S} \mathbb{E} \sum_{t=1}^{n} \ell_t^T u$$

Loss assumptions

Definition (L_∞)

We say that the adversary statisfies the L_∞ **assumption:** if $\|\ell_t\|_\infty \leq 1$ for all $t = 1, \ldots, n$.

Definition (L_2)

We say that the adversary statisfies the L_2 **assumption:** if $\ell_t^T v \leq 1$ for all $t = 1, \ldots, n$ and $v \in \mathcal{S}$.

Loss assumptions

Definition (L_∞)

We say that the adversary statisfies the L_∞ **assumption:** if $\|\ell_t\|_\infty \leq 1$ for all $t = 1, \ldots, n$.

Definition (L_2)

We say that the adversary statisfies the L_2 **assumption:** if $\ell_t^T v \leq 1$ for all $t = 1, \ldots, n$ and $v \in \mathcal{S}$.

Key idea

$$V_t \sim p_t, \quad p_t \in \Delta(\mathcal{S})$$

Then, unbiased estimate $\tilde{\ell}_t$ of the loss ℓ_t:

- $\tilde{\ell}_t = \ell_t$ in the full information game,
- $\tilde{\ell}_{i,t} = \frac{\ell_{i,t}}{\sum_{V \in \mathcal{S}: V_i = 1} p_t(V)} V_{i,t}$ in the semi-bandit game,
- $\tilde{\ell}_t = P_t^+ V_t V_t^T \ell_t$, with $P_t = \mathbb{E}_{V \sim p_t}(V V^T)$ in the bandit game.

Key idea

$$V_t \sim p_t, \;\; p_t \in \Delta(\mathcal{S})$$

Then, unbiased estimate $\tilde{\ell}_t$ of the loss ℓ_t:

- $\tilde{\ell}_t = \ell_t$ in the full information game,
- $\tilde{\ell}_{i,t} = \frac{\ell_{i,t}}{\sum_{V \in \mathcal{S}: V_i = 1} p_t(V)} V_{i,t}$ in the semi-bandit game,
- $\tilde{\ell}_t = P_t^{+} V_t V_t^T \ell_t$, with $P_t = \mathbb{E}_{V \sim p_t}(V V^T)$ in the bandit game.

Key idea

$$V_t \sim p_t, \quad p_t \in \Delta(\mathcal{S})$$

Then, unbiased estimate $\tilde{\ell}_t$ of the loss ℓ_t:

- $\tilde{\ell}_t = \ell_t$ in the full information game,
- $\tilde{\ell}_{i,t} = \frac{\ell_{i,t}}{\sum_{V \in \mathcal{S}: V_i = 1} p_t(V)} V_{i,t}$ in the semi-bandit game,
- $\tilde{\ell}_t = P_t^+ V_t V_t^T \ell_t$, with $P_t = \mathbb{E}_{V \sim p_t}(V V^T)$ in the bandit game.

Key idea

$$V_t \sim p_t, \quad p_t \in \Delta(\mathcal{S})$$

Then, unbiased estimate $\tilde{\ell}_t$ of the loss ℓ_t:

- $\tilde{\ell}_t = \ell_t$ in the full information game,
- $\tilde{\ell}_{i,t} = \frac{\ell_{i,t}}{\sum_{V \in \mathcal{S}: V_i = 1} p_t(V)} V_{i,t}$ in the semi-bandit game,
- $\tilde{\ell}_t = P_t^{+} V_t V_t^T \ell_t$, with $P_t = \mathbb{E}_{V \sim p_t}(V V^T)$ in the bandit game.

Key idea

$$V_t \sim p_t, \quad p_t \in \Delta(\mathcal{S})$$

Then, unbiased estimate $\tilde{\ell}_t$ of the loss ℓ_t:

- $\tilde{\ell}_t = \ell_t$ in the full information game,
- $\tilde{\ell}_{i,t} = \frac{\ell_{i,t}}{\sum_{V \in \mathcal{S}: V_i = 1} p_t(V)} V_{i,t}$ in the semi-bandit game,
- $\tilde{\ell}_t = P_t^+ V_t V_t^T \ell_t$, with $P_t = \mathbb{E}_{V \sim p_t}(V V^T)$ in the bandit game.

Expanded Exponentially weighted average forecaster (Exp2)

$$p_t(v) = \frac{\exp\left(-\eta \sum_{s=1}^{t-1} \tilde{\ell}_s^T v\right)}{\sum_{u \in \mathcal{S}} \exp\left(-\eta \sum_{s=1}^{t-1} \tilde{\ell}_s^T u\right)}$$

- In the full information game, against L_2 adversaries, we have (for some η)

$$R_n \leq \sqrt{2dn},$$

  which is the optimal rate, Dani, Hayes and Kakade [2008].
- Thus against L_∞ adversaries we have

$$R_n \leq d^{3/2}\sqrt{2n}.$$

  But this is suboptimal, Koolen, Warmuth and Kivinen [2010].
- We show that, for any η, there exists a subset $S \subset \{0,1\}^d$ and an L_∞ adversary such that

$$R_n \geq 0.02\, d^{3/2}\sqrt{n}.$$

Expanded Exponentially weighted average forecaster (Exp2)

$$p_t(v) = \frac{\exp\left(-\eta \sum_{s=1}^{t-1} \tilde{\ell}_s^T v\right)}{\sum_{u \in \mathcal{S}} \exp\left(-\eta \sum_{s=1}^{t-1} \tilde{\ell}_s^T u\right)}$$

- In the full information game, against L_2 adversaries, we have (for some η)

$$R_n \leq \sqrt{2dn},$$

  which is the optimal rate, Dani, Hayes and Kakade [2008].

- Thus against L_∞ adversaries we have

$$R_n \leq d^{3/2}\sqrt{2n}.$$

  But this is suboptimal, Koolen, Warmuth and Kivinen [2010].

- We show that, for any η, there exists a subset $S \subset \{0,1\}^d$ and an L_∞ adversary such that

$$R_n \geq 0.02\, d^{3/2}\sqrt{n}.$$

Expanded Exponentially weighted average forecaster (Exp2)

$$p_t(v) = \frac{\exp\left(-\eta \sum_{s=1}^{t-1} \tilde{\ell}_s^T v\right)}{\sum_{u \in \mathcal{S}} \exp\left(-\eta \sum_{s=1}^{t-1} \tilde{\ell}_s^T u\right)}$$

- In the full information game, against L_2 adversaries, we have (for some η)

  $$R_n \leq \sqrt{2dn},$$

  which is the optimal rate, Dani, Hayes and Kakade [2008].
- Thus against L_∞ adversaries we have

  $$R_n \leq d^{3/2}\sqrt{2n}.$$

  But this is suboptimal, Koolen, Warmuth and Kivinen [2010].
- We show that, for any η, there exists a subset $S \subset \{0,1\}^d$ and an L_∞ adversary such that

  $$R_n \geq 0.02\, d^{3/2}\sqrt{n}.$$

Expanded Exponentially weighted average forecaster (Exp2)

$$p_t(v) = \frac{\exp\left(-\eta \sum_{s=1}^{t-1} \tilde{\ell}_s^T v\right)}{\sum_{u \in \mathcal{S}} \exp\left(-\eta \sum_{s=1}^{t-1} \tilde{\ell}_s^T u\right)}$$

- In the full information game, against L_2 adversaries, we have (for some η)

$$R_n \leq \sqrt{2dn},$$

  which is the optimal rate, Dani, Hayes and Kakade [2008].
- Thus against L_∞ adversaries we have

$$R_n \leq d^{3/2}\sqrt{2n}.$$

  But this is suboptimal, Koolen, Warmuth and Kivinen [2010].
- We show that, for any η, there exists a subset $S \subset \{0,1\}^d$ and an L_∞ adversary such that

$$R_n \geq 0.02\, d^{3/2}\sqrt{n}.$$

Legendre function

Definition

Let $\mathcal{D}$ be a convex subset of $\mathbb{R}^d$ with nonempty interior $\text{int}(\mathcal{D})$ and boundary $\partial\mathcal{D}$. We call Legendre any function $F : \mathcal{D} \to \mathbb{R}$ such that

- F is strictly convex and admits continuous first partial derivatives on $\text{int}(\mathcal{D})$,
- For any $u \in \partial\mathcal{D}$, for any $v \in \text{int}(\mathcal{D})$, we have

$$\lim_{s \to 0, s > 0} (u - v)^T \nabla F\big((1 - s)u + sv\big) = +\infty.$$

Legendre function

Definition

Let $\mathcal{D}$ be a convex subset of $\mathbb{R}^d$ with nonempty interior $\text{int}(\mathcal{D})$ and boundary $\partial\mathcal{D}$. We call Legendre any function $F : \mathcal{D} \to \mathbb{R}$ such that

- F is strictly convex and admits continuous first partial derivatives on $\text{int}(\mathcal{D})$,
- For any $u \in \partial\mathcal{D}$, for any $v \in \text{int}(\mathcal{D})$, we have

$$\lim_{s \to 0, s > 0} (u - v)^T \nabla F\big((1 - s)u + sv\big) = +\infty.$$

Legendre function

Definition

Let $\mathcal{D}$ be a convex subset of $\mathbb{R}^d$ with nonempty interior $\mathrm{int}(\mathcal{D})$ and boundary $\partial\mathcal{D}$. We call Legendre any function $F : \mathcal{D} \to \mathbb{R}$ such that

- F is strictly convex and admits continuous first partial derivatives on $\mathrm{int}(\mathcal{D})$,
- For any $u \in \partial\mathcal{D}$, for any $v \in \mathrm{int}(\mathcal{D})$, we have

$$\lim_{s \to 0, s > 0} (u - v)^T \nabla F\big((1 - s)u + sv\big) = +\infty.$$

Bregman divergence

Definition

The Bregman divergence $D_F : \mathcal{D} \times \text{int}(\mathcal{D})$ associated to a Legendre function F is defined by

$$D_F(u, v) = F(u) - F(v) - (u - v)^T \nabla F(v).$$

CLEB (Combinatorial LEarning with Bregman divergences)

Parameter: F Legendre on $\mathcal{D} \supset Conv(\mathcal{S})$

(1) $w'_{t+1} \in \mathcal{D}$:
$\nabla F(w'_{t+1}) = \nabla F(w_t) - \tilde{\ell}_t$

(2) $w_{t+1} \in \underset{w \in Conv(\mathcal{S})}{\operatorname{argmin}} D_F(w, w'_{t+1})$

(3) $p_{t+1} \in \Delta(\mathcal{S}) : w_{t+1} = \mathbb{E}_{V \sim p_{t+1}} V$

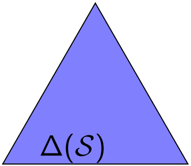

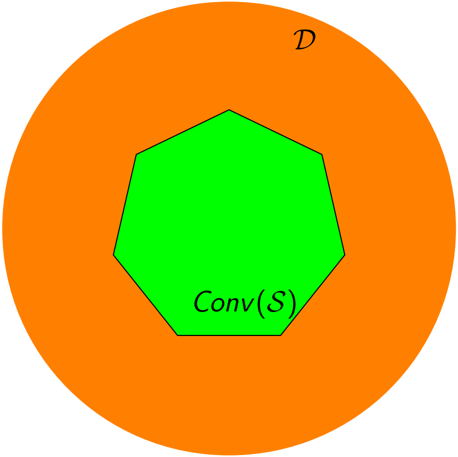

CLEB (Combinatorial LEarning with Bregman divergences)

Parameter: F Legendre on $\mathcal{D} \supset Conv(\mathcal{S})$

(1) $w'_{t+1} \in \mathcal{D}$:
$\nabla F(w'_{t+1}) = \nabla F(w_t) - \tilde{\ell}_t$

(2) $w_{t+1} \in \underset{w \in Conv(\mathcal{S})}{\operatorname{argmin}} D_F(w, w'_{t+1})$

(3) $p_{t+1} \in \Delta(\mathcal{S}) : w_{t+1} = \mathbb{E}_{V \sim p_{t+1}} V$

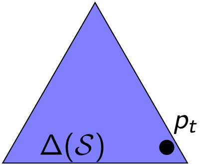

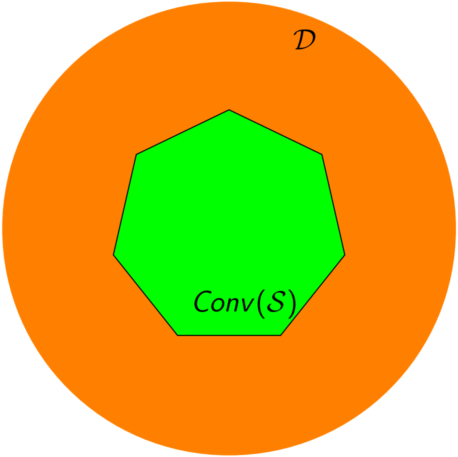

CLEB (Combinatorial LEarning with Bregman divergences)

Parameter: F Legendre on $\mathcal{D} \supset Conv(\mathcal{S})$

(1) $w'_{t+1} \in \mathcal{D}$:
$\nabla F(w'_{t+1}) = \nabla F(w_t) - \tilde{\ell}_t$

(2) $w_{t+1} \in \underset{w \in Conv(\mathcal{S})}{\operatorname{argmin}} D_F(w, w'_{t+1})$

(3) $p_{t+1} \in \Delta(\mathcal{S}) : w_{t+1} = \mathbb{E}_{V \sim p_{t+1}} V$

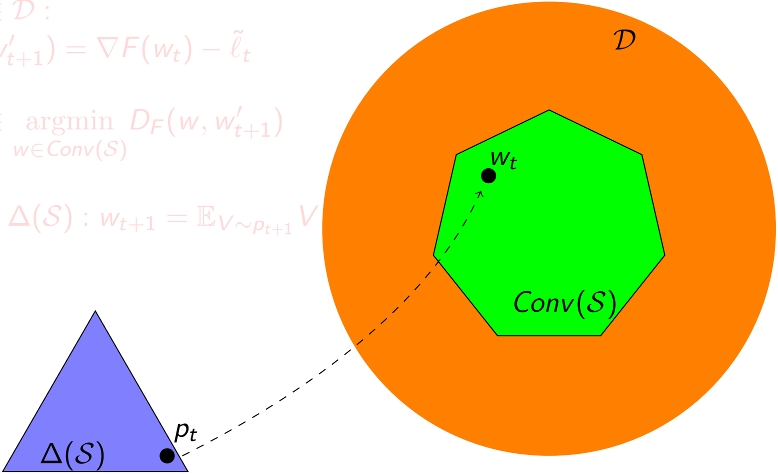

CLEB (Combinatorial LEarning with Bregman divergences)

Parameter: F Legendre on $\mathcal{D} \supset Conv(\mathcal{S})$

(1) $w'_{t+1} \in \mathcal{D}$:
$$\nabla F(w'_{t+1}) = \nabla F(w_t) - \tilde{\ell}_t$$

(2) $w_{t+1} \in \underset{w \in Conv(\mathcal{S})}{\operatorname{argmin}} D_F(w, w'_{t+1})$

(3) $p_{t+1} \in \Delta(\mathcal{S}) : w_{t+1} = \mathbb{E}_{V \sim p_{t+1}} V$

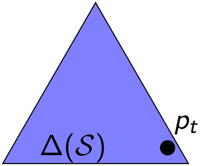

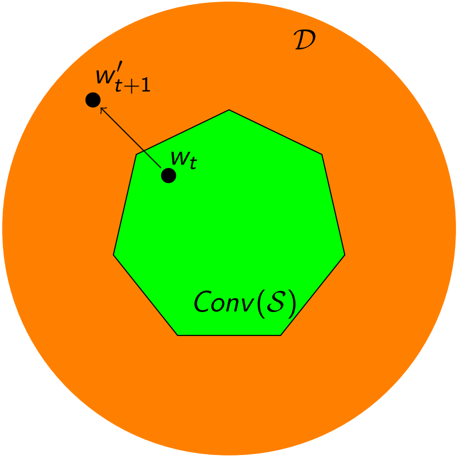

CLEB (Combinatorial LEarning with Bregman divergences)

Parameter: F Legendre on $\mathcal{D} \supset Conv(\mathcal{S})$

(1) $w'_{t+1} \in \mathcal{D}$:
$\nabla F(w'_{t+1}) = \nabla F(w_t) - \tilde{\ell}_t$

(2) $w_{t+1} \in \underset{w \in Conv(\mathcal{S})}{\operatorname{argmin}} D_F(w, w'_{t+1})$

(3) $p_{t+1} \in \Delta(\mathcal{S}) : w_{t+1} = \mathbb{E}_{V \sim p_{t+1}} V$

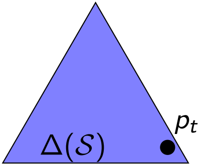

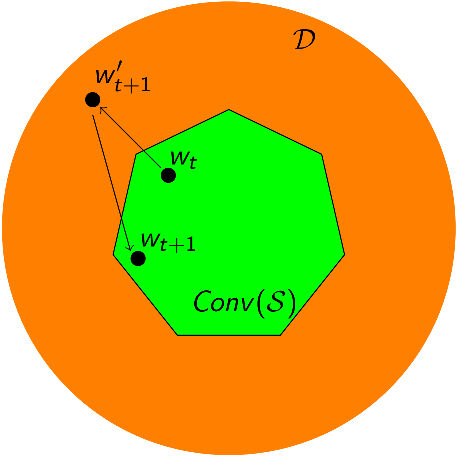

CLEB (Combinatorial LEarning with Bregman divergences)

Parameter: F Legendre on $\mathcal{D} \supset Conv(\mathcal{S})$

(1) $w'_{t+1} \in \mathcal{D}$:
$$\nabla F(w'_{t+1}) = \nabla F(w_t) - \tilde{\ell}_t$$

(2) $w_{t+1} \in \underset{w \in Conv(\mathcal{S})}{\operatorname{argmin}} D_F(w, w'_{t+1})$

(3) $p_{t+1} \in \Delta(\mathcal{S}) : w_{t+1} = \mathbb{E}_{V \sim p_{t+1}} V$

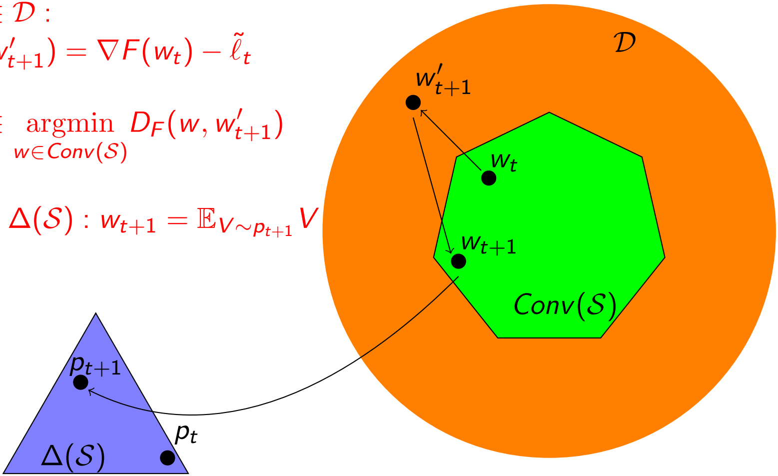

General regret bound for CLEB

Theorem

If F admits a Hessian $\nabla^2 F$ always invertible then,

$$R_n \lessapprox \mathit{diam}_{D_F}(\mathcal{S}) + \mathbb{E} \sum_{t=1}^{n} \tilde{\ell}_t^T \left(\nabla^2 F(w_t)\right)^{-1} \tilde{\ell}_t.$$

Different instances of CLEB: LinExp (Entropy Function)

$\mathcal{D} = [0, +\infty)^d$, $F(x) = \frac{1}{\eta} \sum_{i=1}^{d} x_i \log x_i$

Full Info: Hedge

Semi-Bandit=Bandit: Exp3
Auer et al. [2002]

Full Info: Component Hedge
Koolen, Warmuth and Kivinen [2010]

Semi-Bandit: MW
Kale, Reyzin and Schapire [2010]

Bandit: new algorithm

Different instances of CLEB: LinExp (Entropy Function)

$\mathcal{D} = [0, +\infty)^d$, $F(x) = \frac{1}{\eta} \sum_{i=1}^{d} x_i \log x_i$

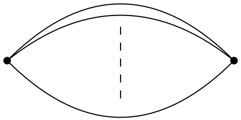

Full Info: Hedge

Semi-Bandit=Bandit: Exp3
Auer et al. [2002]

Full Info: Component Hedge
Koolen, Warmuth and Kivinen [2010]

Semi-Bandit: MW
Kale, Reyzin and Schapire [2010]

Bandit: new algorithm

Different instances of CLEB: LinExp (Entropy Function)

$\mathcal{D} = [0, +\infty)^d$, $F(x) = \frac{1}{\eta} \sum_{i=1}^{d} x_i \log x_i$

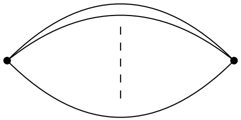

Full Info: Hedge

Semi-Bandit=Bandit: Exp3
Auer et al. [2002]

Full Info: Component Hedge
Koolen, Warmuth and Kivinen [2010]

Semi-Bandit: MW
Kale, Reyzin and Schapire [2010]

Bandit: new algorithm

Different instances of CLEB: LinExp (Entropy Function)

$\mathcal{D} = [0, +\infty)^d$, $F(x) = \frac{1}{\eta} \sum_{i=1}^{d} x_i \log x_i$

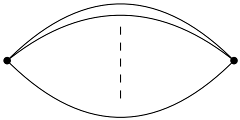

- Full Info: Hedge
- Semi-Bandit=Bandit: Exp3
  Auer et al. [2002]

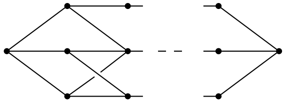

- Full Info: Component Hedge
  Koolen, Warmuth and Kivinen [2010]
- Semi-Bandit: MW
  Kale, Reyzin and Schapire [2010]
- Bandit: new algorithm

Different instances of CLEB: LinExp (Entropy Function)

$\mathcal{D} = [0, +\infty)^d$, $F(x) = \frac{1}{\eta} \sum_{i=1}^{d} x_i \log x_i$

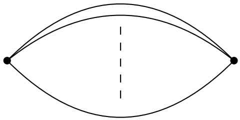

- Full Info: Hedge
- Semi-Bandit=Bandit: Exp3
  Auer et al. [2002]

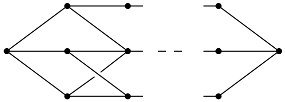

- Full Info: Component Hedge
  Koolen, Warmuth and Kivinen [2010]
- Semi-Bandit: MW
  Kale, Reyzin and Schapire [2010]
- Bandit: new algorithm

Different instances of CLEB: LinExp (Entropy Function)

$\mathcal{D} = [0, +\infty)^d$, $F(x) = \frac{1}{\eta} \sum_{i=1}^{d} x_i \log x_i$

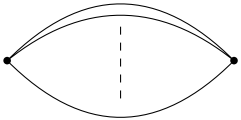

Full Info: Hedge

Semi-Bandit=Bandit: Exp3
Auer et al. [2002]

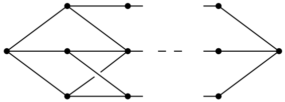

Full Info: Component Hedge
Koolen, Warmuth and Kivinen [2010]

Semi-Bandit: MW
Kale, Reyzin and Schapire [2010]

Bandit: new algorithm

Different instances of CLEB: LinINF (Exchangeable Hessian)

$\mathcal{D} = [0, +\infty)^d$, $F(x) = \sum_{i=1}^{d} \int_0^{x_i} \psi^{-1}(s) ds$

INF, Audibert and Bubeck [2009]

$$\begin{cases} \psi(x) = \exp(\eta x) : \text{LinExp} \\ \psi(x) = (-\eta x)^{-q}, q > 1 : \text{LinPoly} \end{cases}$$

Different instances of CLEB: LinINF (Exchangeable Hessian)

$\mathcal{D} = [0, +\infty)^d$, $F(x) = \sum_{i=1}^{d} \int_0^{x_i} \psi^{-1}(s) ds$

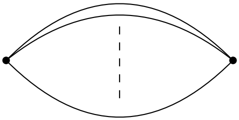

INF, Audibert and Bubeck [2009]

$\begin{cases} \psi(x) = \exp(\eta x) : \text{LinExp} \\ \psi(x) = (-\eta x)^{-q}, q > 1 : \text{LinPoly} \end{cases}$

Different instances of CLEB: LinINF (Exchangeable Hessian)

$\mathcal{D} = [0, +\infty)^d$, $F(x) = \sum_{i=1}^{d} \int_0^{x_i} \psi^{-1}(s) ds$

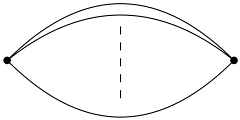

INF, Audibert and Bubeck [2009]

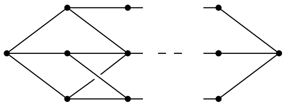

$\begin{cases} \psi(x) = \exp(\eta x) : \text{LinExp} \\ \psi(x) = (-\eta x)^{-q}, q > 1 : \text{LinPoly} \end{cases}$

Different instances of CLEB: LinINF (Exchangeable Hessian)

$\mathcal{D} = [0, +\infty)^d$, $F(x) = \sum_{i=1}^{d} \int_0^{x_i} \psi^{-1}(s)ds$

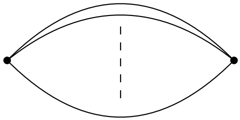

INF, Audibert and Bubeck [2009]

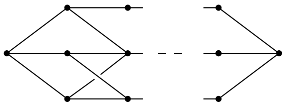

$$\begin{cases} \psi(x) = \exp(\eta x) : \text{LinExp} \\ \psi(x) = (-\eta x)^{-q}, q > 1 : \text{LinPoly} \end{cases}$$

Different instances of CLEB: Follow the regularized leader

$\mathcal{D} = Conv(\mathcal{S})$, then

$$w_{t+1} \in \operatorname*{argmin}_{w \in \mathcal{D}} \left(\sum_{s=1}^{t} \tilde{\ell}_s^T w + F(w) \right)$$

Particularly interesting choice: F self-concordant barrier function, Abernethy, Hazan and Rakhlin [2008]

Different instances of CLEB: Follow the regularized leader

$\mathcal{D} = Conv(\mathcal{S})$, then

$$w_{t+1} \in \operatorname*{argmin}_{w \in \mathcal{D}} \left(\sum_{s=1}^{t} \tilde{\ell}_s^T w + F(w) \right)$$

Particularly interesting choice: F self-concordant barrier function, Abernethy, Hazan and Rakhlin [2008]

Minimax regret for combinatorial prediction games

$$\overline{R}_{n,\infty,2} = \inf_{strategy} \max_{\mathcal{S}\subset\{0,1\}^d} \sup_{L_\infty, L_2 \, adversaries} R_n$$

Theorem

Let $n \geq d^2$. In the full information and semi-bandit games, we have:

$$0.008\, d\sqrt{n} \leq \overline{R}_{n,\infty} \leq d\sqrt{2n},$$

$$0.05\, \sqrt{dn} \leq \overline{R}_{n,2} \leq \sqrt{2edn\log(ed)},$$

and in the bandit game:

$$0.01\, d^{3/2}\sqrt{n} \leq \overline{R}_{n,\infty} \leq 2\, d^{5/2}\sqrt{2n}.$$

$$0.05\, d\sqrt{n} \leq \overline{R}_{n,2} \leq d^{3/2}\sqrt{2n}.$$

Minimax regret for combinatorial prediction games

$$\overline{R}_{n,\infty,2} = \inf_{strategy} \max_{\mathcal{S} \subset \{0,1\}^d} \sup_{L_\infty, L_2 \, adversaries} R_n$$

Theorem

Let $n \geq d^2$. In the full information and semi-bandit games, we have:

$$0.008\, d\sqrt{n} \leq \overline{R}_{n,\infty} \leq d\sqrt{2n},$$

$$0.05\, \sqrt{dn} \leq \overline{R}_{n,2} \leq \sqrt{2edn\log(ed)},$$

and in the bandit game:

$$0.01\, d^{3/2}\sqrt{n} \leq \overline{R}_{n,\infty} \leq 2\, d^{5/2}\sqrt{2n}.$$

$$0.05\, d\sqrt{n} \leq \overline{R}_{n,2} \leq d^{3/2}\sqrt{2n}.$$