Bandits Games
and Combinatorial Problems in Statistics

Sébastien Bubeck,

Centre de Recerca Matemàtica, Barcelone

Standard prediction game

Adversary

1: CNN

2: NBC

- - -

K: ABC

Player

Standard prediction game

Adversary

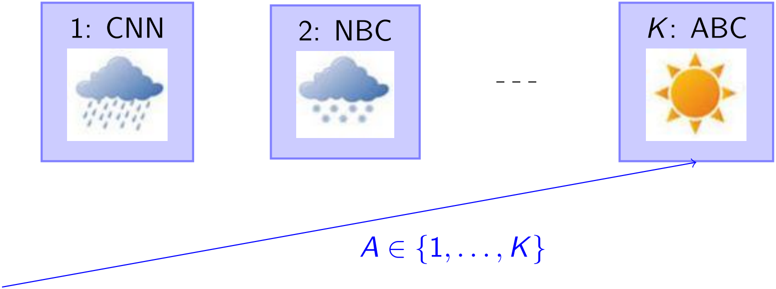

Player

Standard prediction game

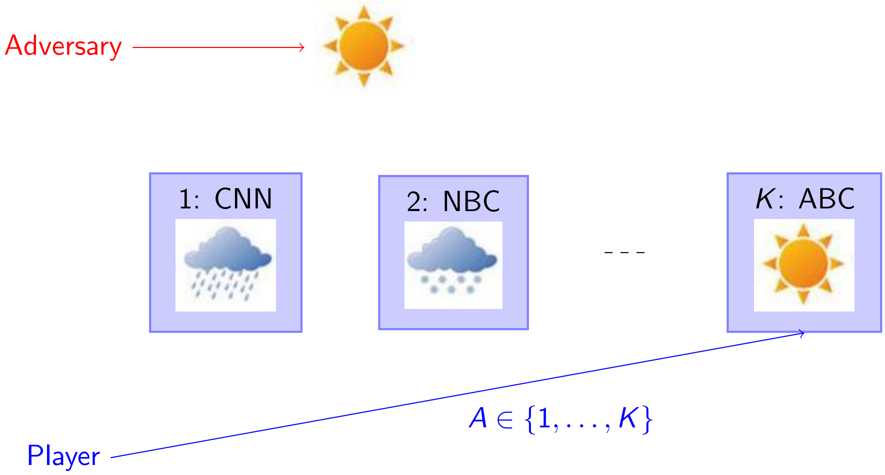

Standard prediction game

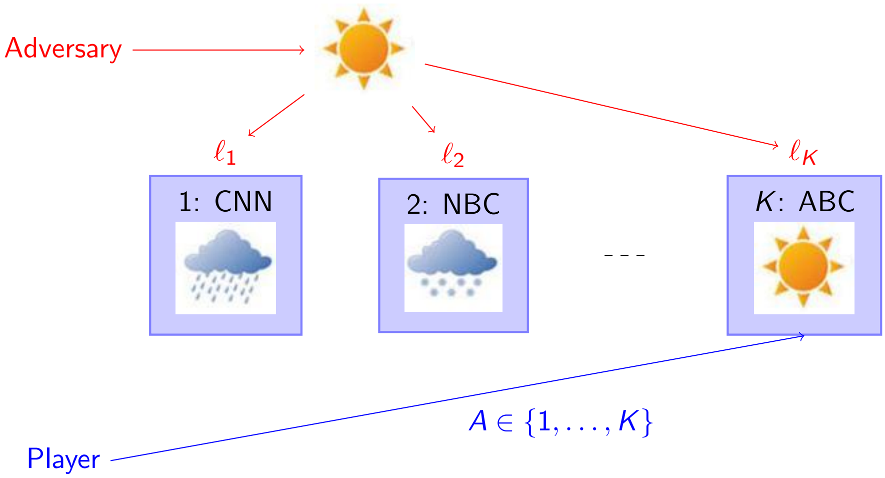

Standard prediction game

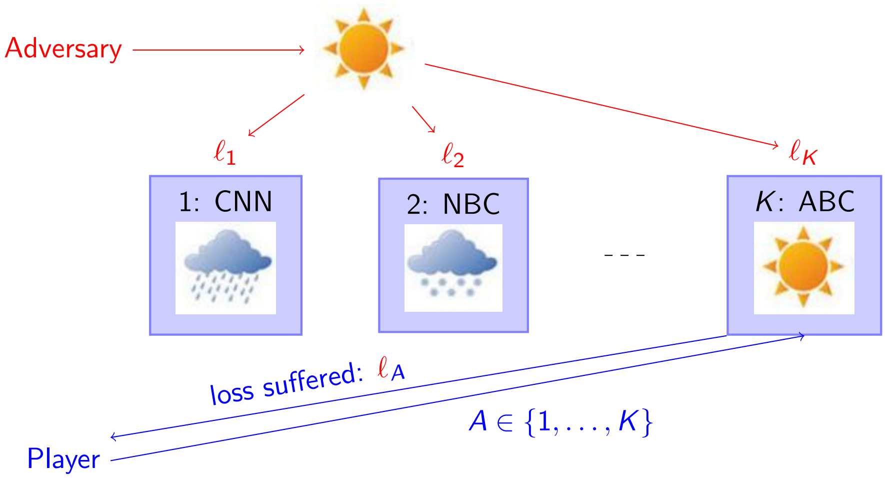

Standard prediction game

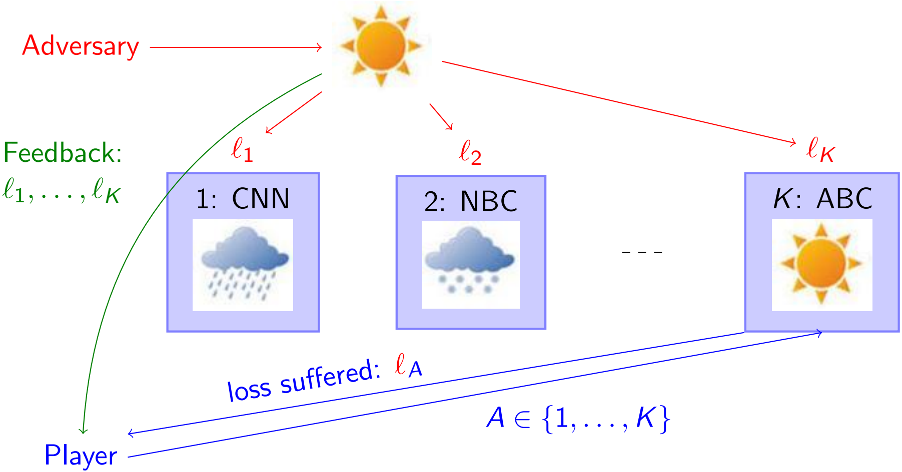

$$R_n = \mathbb{E} \sum_{t=1}^{n} \ell_{A_t,t} - \min_{a \in \{1,\ldots,K\}} \mathbb{E} \sum_{t=1}^{n} \ell_{a,t}.$$

Standard prediction game

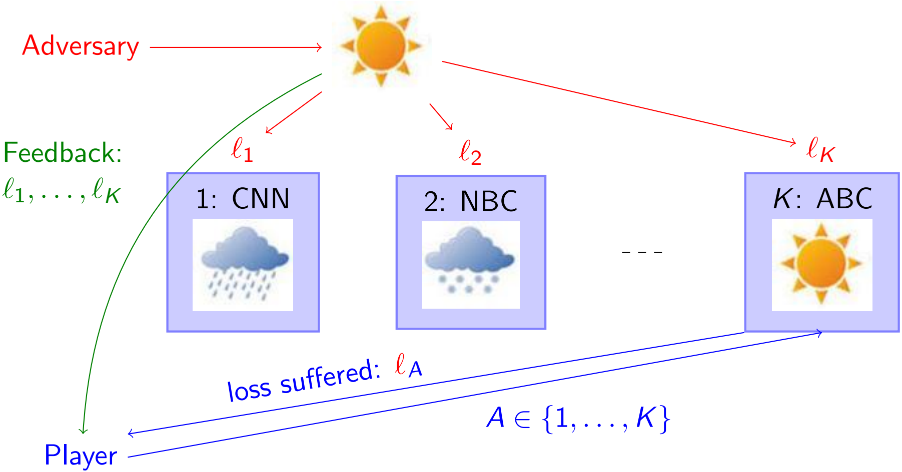

$$R_n = \mathbb{E} \sum_{t=1}^{n} \ell_{A_t,t} - \min_{a \in \{1,\ldots,K\}} \mathbb{E} \sum_{t=1}^{n} \ell_{a,t}.$$

Standard prediction game

Theorem (Hannan [1957])

There exists a strategy such that $R_n = o(n)$.

Standard prediction game

Theorem (Cesa-Bianchi, Freund , Haussler, Helmbold, Schapire and Warmuth [1997])

Hedge satisfies

$$R_n \leq \sqrt{\frac{n \log K}{2}}.$$

Moreover for any strategy,

$$\sup_{adversaries} R_n \geq \sqrt{\frac{n \log K}{2}} + o(\sqrt{n \log K}).$$

Standard prediction game

Theorem (Cesa-Bianchi, Freund , Haussler, Helmbold, Schapire and Warmuth [1997])

Hedge satisfies

$$R_n \leq \sqrt{\frac{n \log K}{2}}.$$

Moreover for any strategy,

$$\sup_{\textit{adversaries}} R_n \geq \sqrt{\frac{n \log K}{2}} + o(\sqrt{n \log K}).$$

Multi-armed bandit game

Adversary

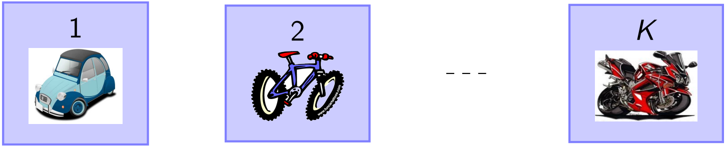

Player

Multi-armed bandit game

Adversary

1

2

- - -

K

$A \in \{1, \ldots, K\}$

Player

Multi-armed bandit game

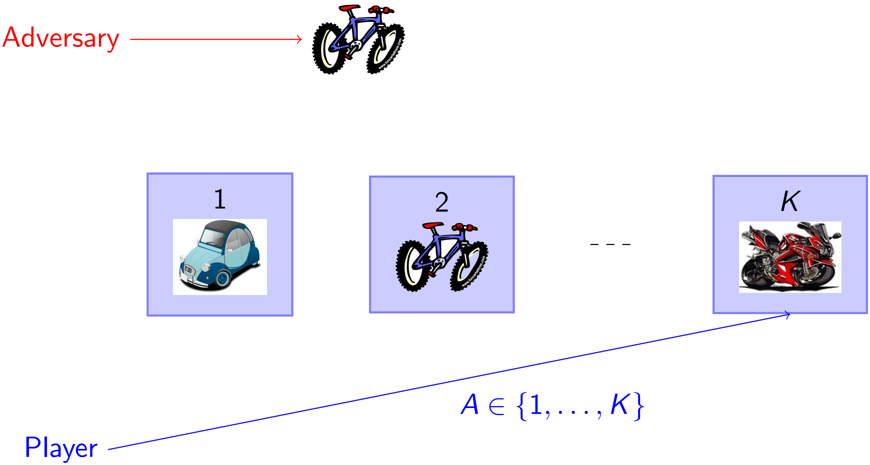

Multi-armed bandit game

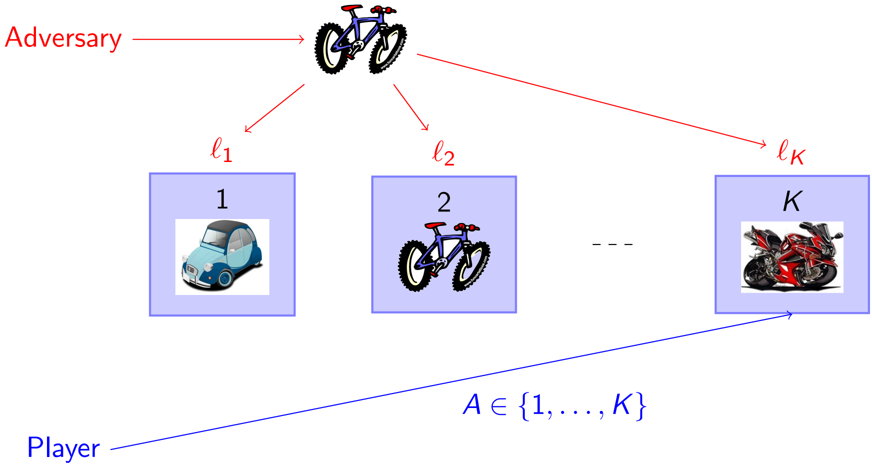

Multi-armed bandit game

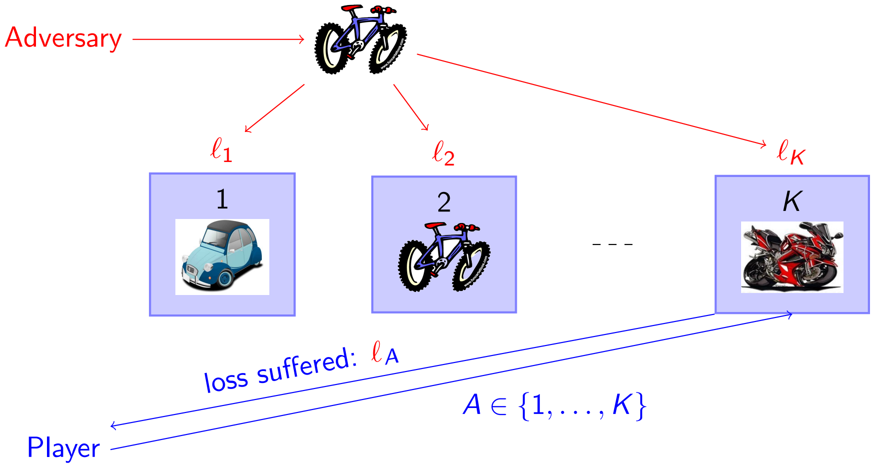

Multi-armed bandit game

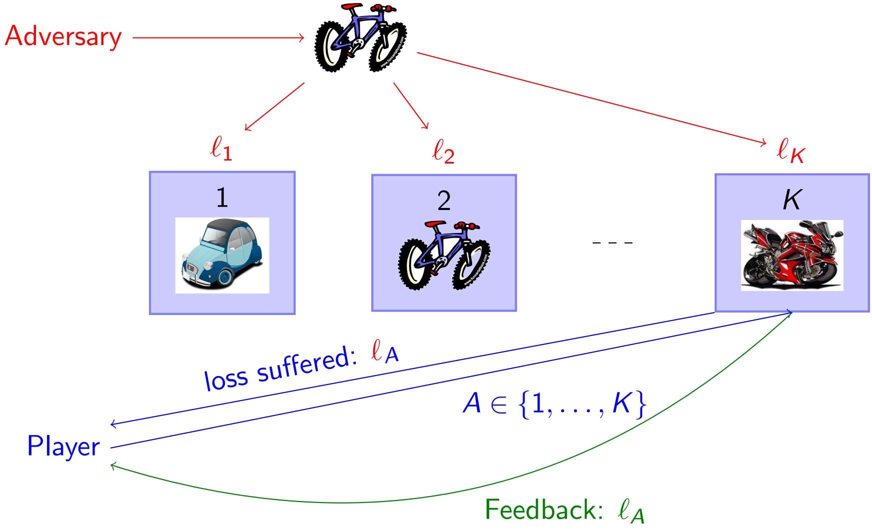

Minimax regret for the multi-armed bandit game

Theorem (Auer, Cesa-Bianchi, Freund and Schapire [1995])

Exp3 satisfies:

$$R_n \leq \sqrt{2nK \log K}.$$

Moreover for any strategy,

$$\sup_{\textit{adversaries}} R_n \geq \frac{1}{4}\sqrt{nK} + o(\sqrt{nK}).$$

Minimax regret for the multi-armed bandit game

Theorem (Auer, Cesa-Bianchi, Freund and Schapire [1995])

Exp3 satisfies:

$$R_n \leq \sqrt{2nK \log K}.$$

Moreover for any strategy,

$$\sup_{\text{adversaries}} R_n \geq \frac{1}{4}\sqrt{nK} + o(\sqrt{nK}).$$

Minimax regret for the multi-armed bandit game

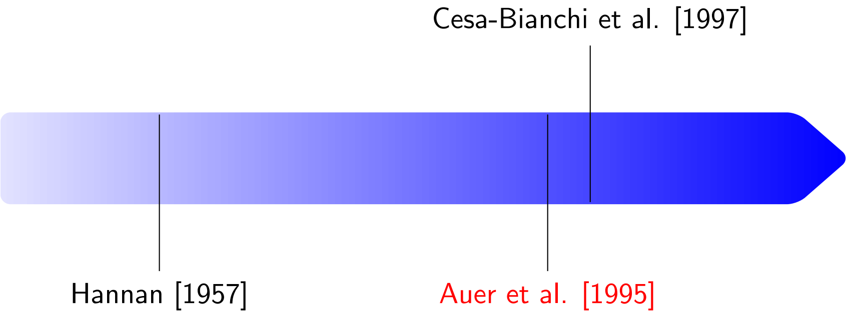

Theorem (Audibert and Bubeck [2009], Audibert and Bubeck [2010], Audibert, Bubeck and Lugosi [2011])

Poly INF satisfies:

$$R_n \leq 2\sqrt{2nK}.$$

Minimax regret for the multi-armed bandit game

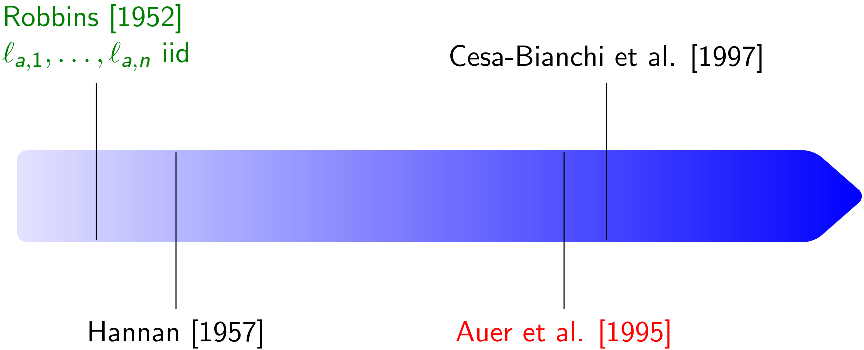

Theorem (Audibert and Bubeck [2009], Audibert and Bubeck [2010], Audibert, Bubeck and Lugosi [2011])

Poly INF satisfies:

$$R_n \leq 2\sqrt{2nK}.$$

Minimax regret for the multi-armed bandit game

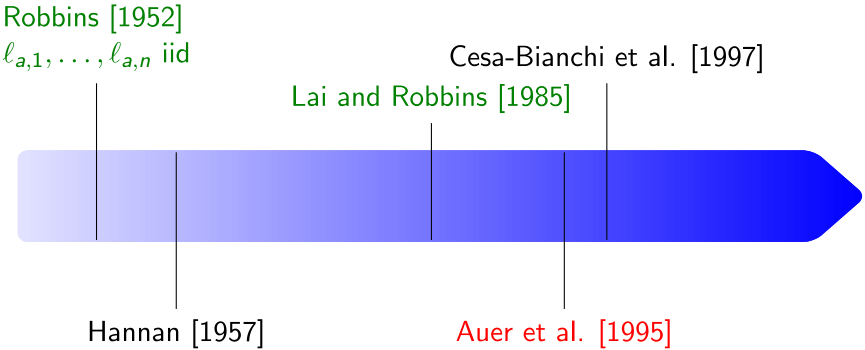

Theorem (Audibert and Bubeck [2009], Audibert and Bubeck [2010], Audibert, Bubeck and Lugosi [2011])

Poly INF satisfies:

$$R_n \leq 2\sqrt{2nK}.$$

Minimax regret for the multi-armed bandit game

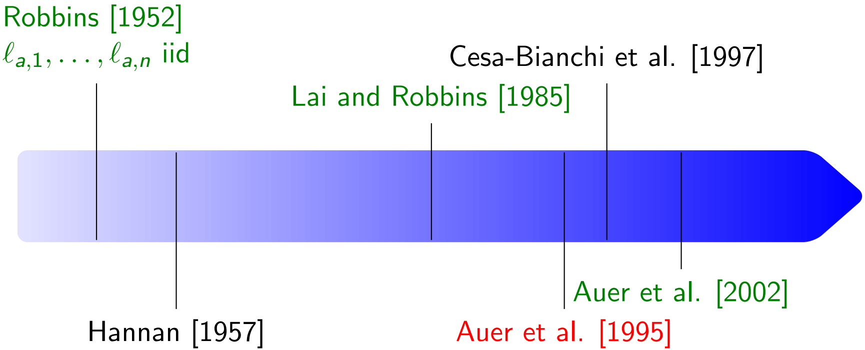

Theorem (Audibert and Bubeck [2009], Audibert and Bubeck [2010], Audibert, Bubeck and Lugosi [2011])

Poly INF satisfies:

$$R_n \leq 2\sqrt{2nK}.$$

Minimax regret for the multi-armed bandit game

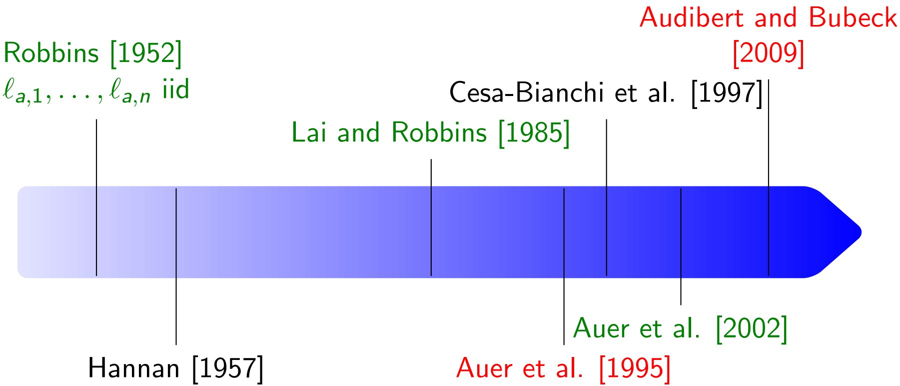

Theorem (Audibert and Bubeck [2009], Audibert and Bubeck [2010], Audibert, Bubeck and Lugosi [2011])

Poly INF satisfies:

$$R_n \leq 2\sqrt{2nK}.$$

Minimax regret for the multi-armed bandit game

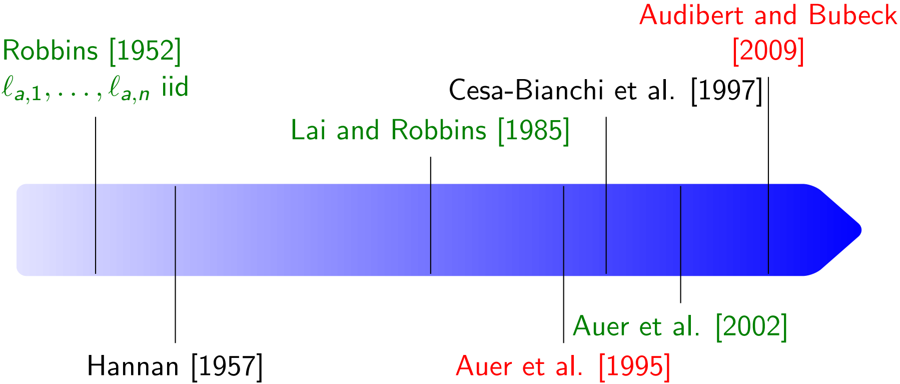

Theorem (Audibert and Bubeck [2009], Audibert and Bubeck [2010], Audibert, Bubeck and Lugosi [2011])

Poly INF satisfies:

$$R_n \leq 2\sqrt{2nK}.$$

High level idea of the proof

- Start with an Abel transform on the regret
- Then multivariate Taylor expansion on the instantaneous regrets, using the implicit function theorem
- Control the main term in the expansion with Hölder's inequality
- Control the second order terms with concentration inequalities for supermartingales

High level idea of the proof

- Start with an Abel transform on the regret
- Then multivariate Taylor expansion on the instantaneous regrets, using the implicit function theorem
- Control the main term in the expansion with Hölder's inequality
- Control the second order terms with concentration inequalities for supermartingales

High level idea of the proof

- Start with an Abel transform on the regret
- Then multivariate Taylor expansion on the instantaneous regrets, using the implicit function theorem
- Control the main term in the expansion with Hölder's inequality
- Control the second order terms with concentration inequalities for supermartingales

High level idea of the proof

- Start with an Abel transform on the regret
- Then multivariate Taylor expansion on the instantaneous regrets, using the implicit function theorem
- Control the main term in the expansion with Hölder's inequality
- Control the second order terms with concentration inequalities for supermartingales

Other contributions to bandit theory

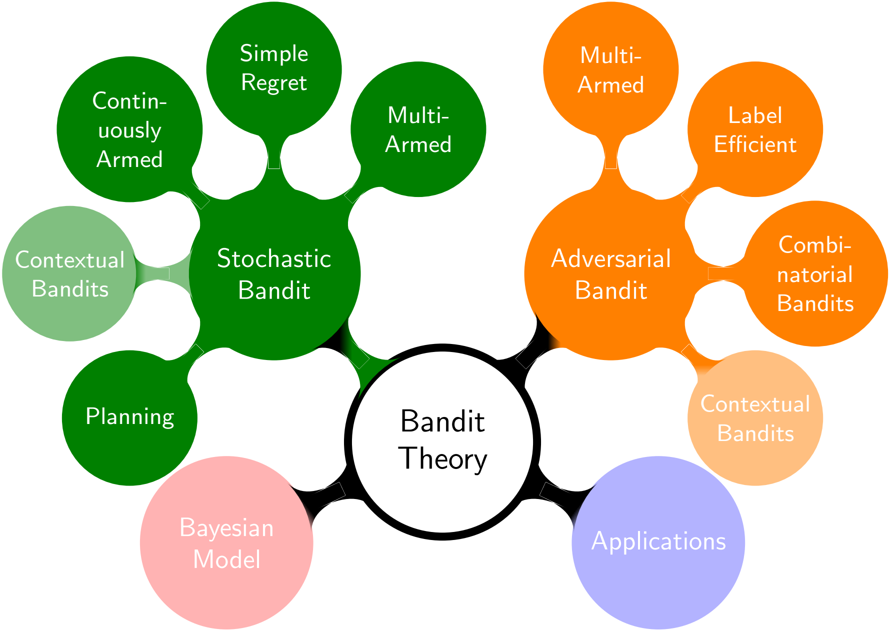

Other contributions to bandit theory

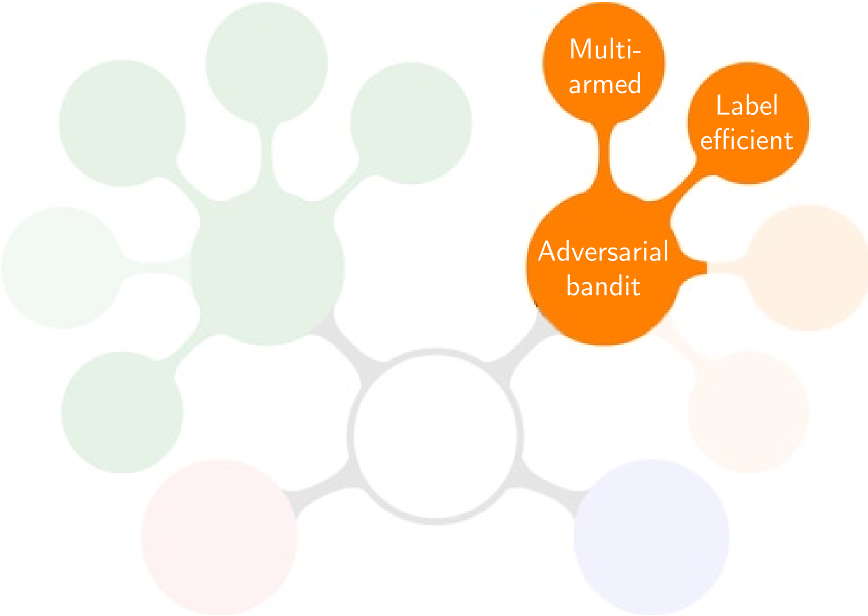

Other contributions to bandit theory

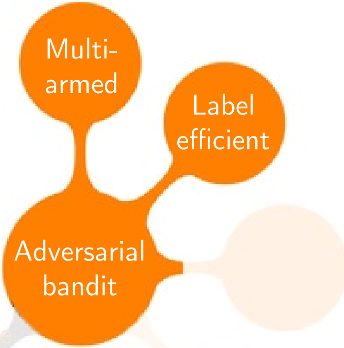

- Player has to ask for the feedback
- He can ask it at most m times

- Tools for the lower bound: Pinsker's inequality, Fano's lemma, chain rule for Kullback-Leibler divergence

Theorem (Audibert and Bubeck [2010])

Standard game: $\quad 0.03\, n\sqrt{\frac{\log K}{m}} \leq \inf \sup R_n \leq n\sqrt{\frac{\log K}{2m}}$

Bandit game: $\quad 0.04\, n\sqrt{\frac{K}{m}} \leq \inf \sup R_n \leq 8\, n\sqrt{\frac{K}{m}}$

Other contributions to bandit theory

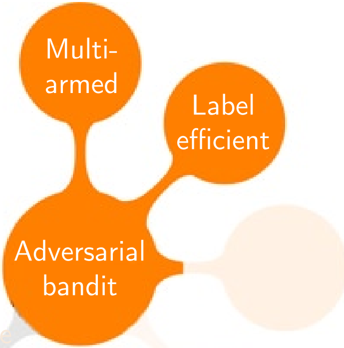

- Player has to ask for the feedback
- He can ask it at most m times

- Tools for the lower bound: Pinsker's inequality, Fano's lemma, chain rule for Kullback-Leibler divergence

Theorem (Audibert and Bubeck [2010])

Standard game: $0.03\, n\sqrt{\frac{\log K}{m}} \leq \inf\sup R_n \leq n\sqrt{\frac{\log K}{2m}}$

Bandit game: $0.04\, n\sqrt{\frac{K}{m}} \leq \inf\sup R_n \leq 8\, n\sqrt{\frac{K}{m}}$

Other contributions to bandit theory

- Player has to ask for the feedback
- He can ask it at most m times

- Tools for the lower bound:
  Pinsker's inequality, Fano's lemma,
  chain rule for Kullback-Leibler divergence

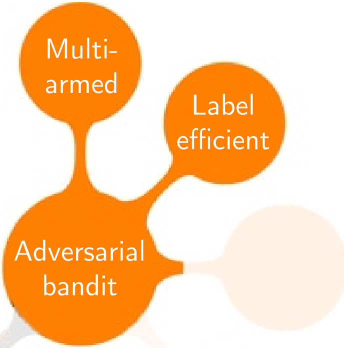

Theorem (Audibert and Bubeck [2010])

Standard game: $$0.03\ n\sqrt{\frac{\log K}{m}} \leq \inf\sup R_n \leq n\sqrt{\frac{\log K}{2m}}$$

Bandit game: $$0.04\ n\sqrt{\frac{K}{m}} \leq \inf\sup R_n \leq 8\ n\sqrt{\frac{K}{m}}$$

Other contributions to bandit theory

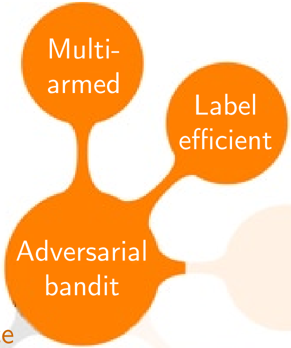

- Player has to ask for the feedback
- He can ask it at most m times

- Tools for the lower bound:
  Pinsker's inequality, Fano's lemma,
  chain rule for Kullback-Leibler divergence

Theorem (Audibert and Bubeck [2010])

Standard game: $0.03\; n\sqrt{\frac{\log K}{m}} \leq \inf \sup R_n \leq n\sqrt{\frac{\log K}{2m}}$

Bandit game: $0.04\; n\sqrt{\frac{K}{m}} \leq \inf \sup R_n \leq 8\; n\sqrt{\frac{K}{m}}$

Other contributions to bandit theory

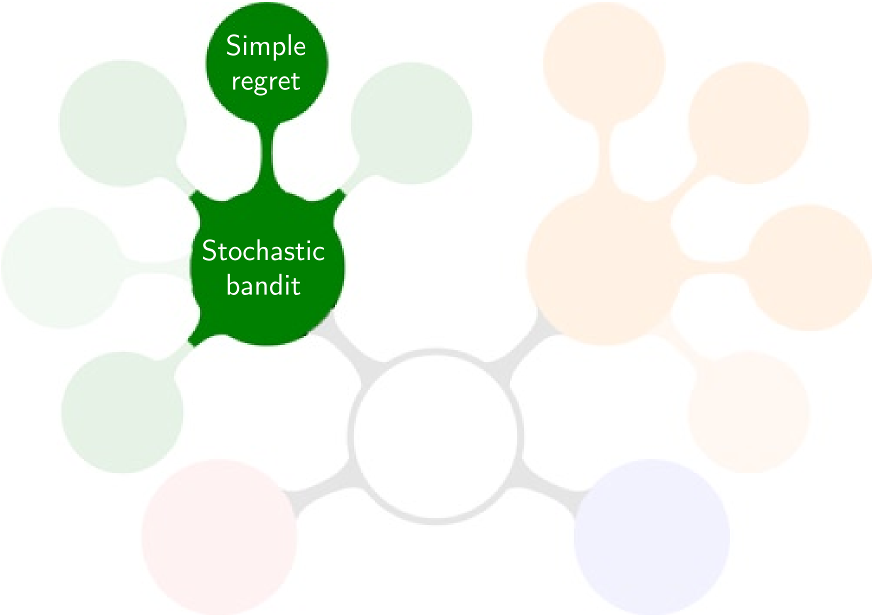

Other contributions to bandit theory

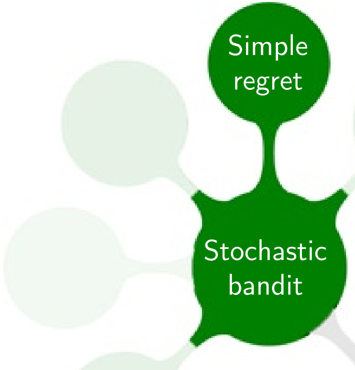

- Frequentist view on offline optimal learning, [Frazier and Powell, 2010]
- Bubeck, Munos and Stoltz [2009, 2010]: links between offline and online setting

Theorem (Audibert, Bubeck and Munos [2010])

Let μ_i be the expected loss of action i. Assume that there is a unique optimal action i^*. Let $H = \sum_{i \neq i^*} (\mu_i - \mu_{i^*})^{-2}$. Then

$$\exp\left(-c' \frac{n \log K}{H}\right) \leq \inf_{Player} \mathbb{P}(A_n \neq i^*) \leq K^2 \exp\left(-c \frac{n}{H \log K}\right).$$

Other contributions to bandit theory

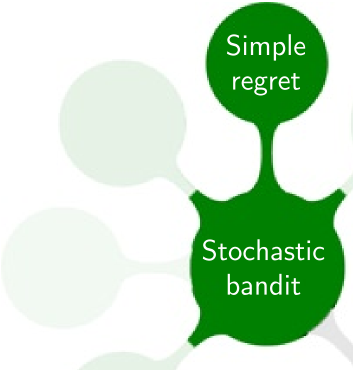

- Frequentist view on offline optimal learning, [Frazier and Powell, 2010]
- Bubeck, Munos and Stoltz [2009, 2010]: links between offline and online setting

Theorem (Audibert, Bubeck and Munos [2010])

Let μ_i be the expected loss of action i. Assume that there is a unique optimal action i^*. Let $H = \sum_{i \neq i^*} (\mu_i - \mu_{i^*})^{-2}$. Then

$$\exp\left(-c' \frac{n \log K}{H}\right) \leq \inf_{Player} \mathbb{P}(A_n \neq i^*) \leq K^2 \exp\left(-c \frac{n}{H \log K}\right).$$

Other contributions to bandit theory

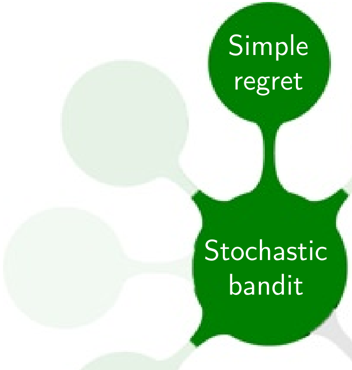

- Frequentist view on offline optimal learning, [Frazier and Powell, 2010]
- Bubeck, Munos and Stoltz [2009, 2010]: links between offline and online setting

Theorem (Audibert, Bubeck and Munos [2010])

Let μ_i be the expected loss of action i. Assume that there is a unique optimal action i^*. Let $H = \sum_{i \neq i^*} (\mu_i - \mu_{i^*})^{-2}$. Then

$$\exp\left(-c' \frac{n \log K}{H}\right) \leq \inf_{Player} \mathbb{P}(A_n \neq i^*) \leq K^2 \exp\left(-c \frac{n}{H \log K}\right).$$

Other contributions to bandit theory

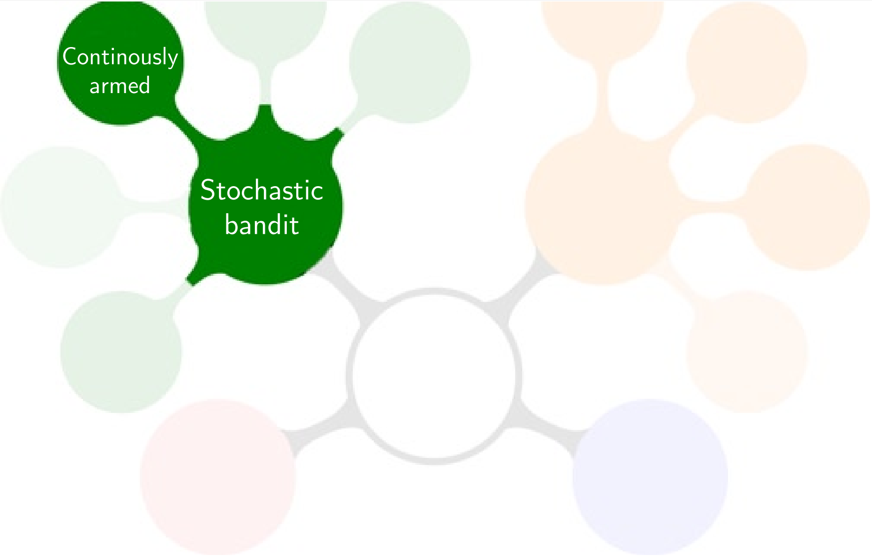

Other contributions to bandit theory

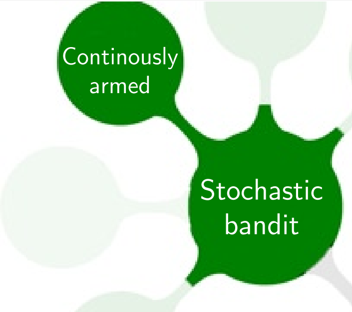

- $\{1, \ldots, K\}$ replaced by arbitrary set $\mathcal{X}$
- Tools: geometry in metric spaces, Hoeffding-Azuma's inequality for martingales

Theorem (Bubeck, Munos, Stoltz and Szepesvari [2009, 2010])

Let $\mathcal{X}$ be a compact subset of $\mathbb{R}^D$ and $\mathcal{F}$ be the set of bandits problems such that the mean-loss function is 1-Lipschitz (with respect to some norm). Then we have

$$\inf \sup_{\mathcal{F}} R_n = \tilde{\Theta}\left(n^{\frac{D+1}{D+2}}\right).$$

Other contributions to bandit theory

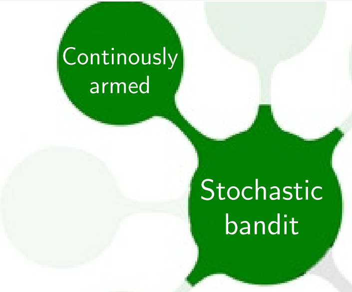

- $\{1, \ldots, K\}$ replaced by arbitrary set $\mathcal{X}$
- Tools: geometry in metric spaces, Hoeffding-Azuma's inequality for martingales

Theorem (Bubeck, Munos, Stoltz and Szepesvari [2009, 2010])

Let $\mathcal{X}$ be a compact subset of $\mathbb{R}^D$ and $\mathcal{F}$ be the set of bandits problems such that the mean-loss function is 1-Lipschitz (with respect to some norm). Then we have

$$\inf \sup_{\mathcal{F}} R_n = \tilde{\Theta}\left(n^{\frac{D+1}{D+2}}\right).$$

Other contributions to bandit theory

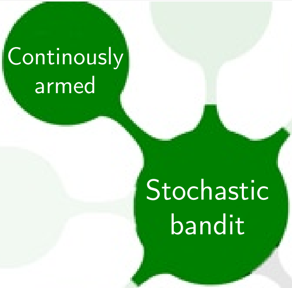

- $\{1, \ldots, K\}$ replaced by arbitrary set $\mathcal{X}$

- Tools: geometry in metric spaces, Hoeffding-Azuma's inequality for martingales

Theorem (Bubeck, Munos, Stoltz and Szepesvari [2009, 2010])

Let $\mathcal{X}$ be a compact subset of $\mathbb{R}^D$ and $\mathcal{F}$ be the set of bandits problems such that the mean-loss function is 1-Lipschitz (with respect to some norm). Then we have

$$\inf \sup_{\mathcal{F}} R_n = \tilde{\Theta}\left(n^{\frac{D+1}{D+2}}\right).$$

Other contributions to bandit theory

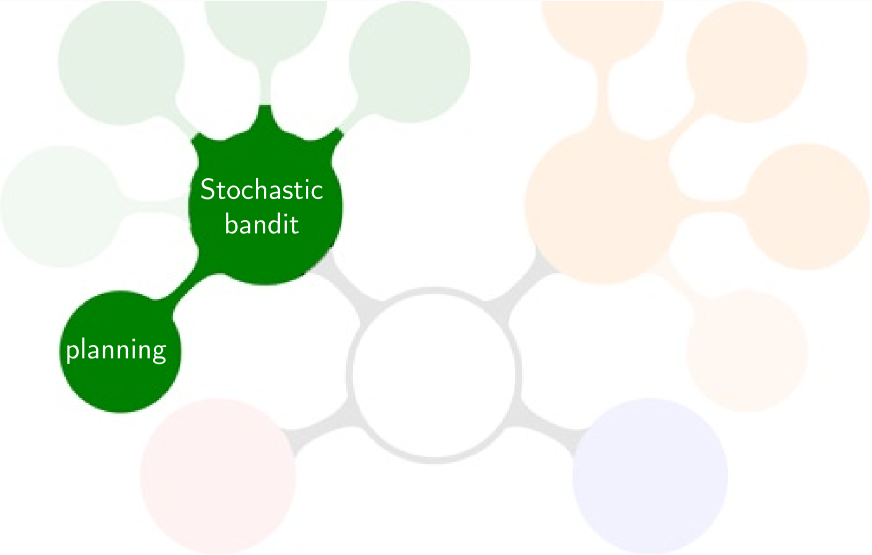

Other contributions to bandit theory

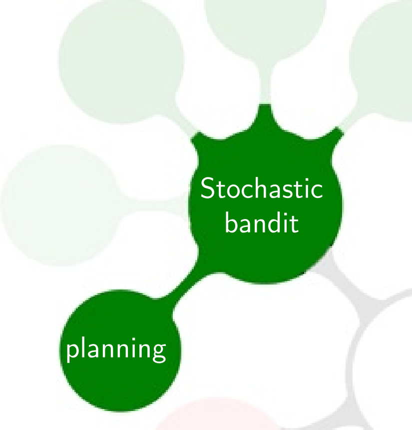

- $\{1, \ldots, K\}$ replaced by $\{1, \ldots, K\}^*$
- loss of t^{th} action discounted by γ^t

Theorem (Bubeck and Munos [2010])

$$\inf \sup R_n = \begin{cases} \tilde{\Theta}\left(n^{1-\frac{\log 1/\gamma}{\log K}}\right) & \text{if } \gamma\sqrt{K} > 1 \\ \tilde{\Theta}\left(\sqrt{n}\right) & \text{if } \gamma\sqrt{K} \leq 1 \end{cases}$$

Other contributions to bandit theory

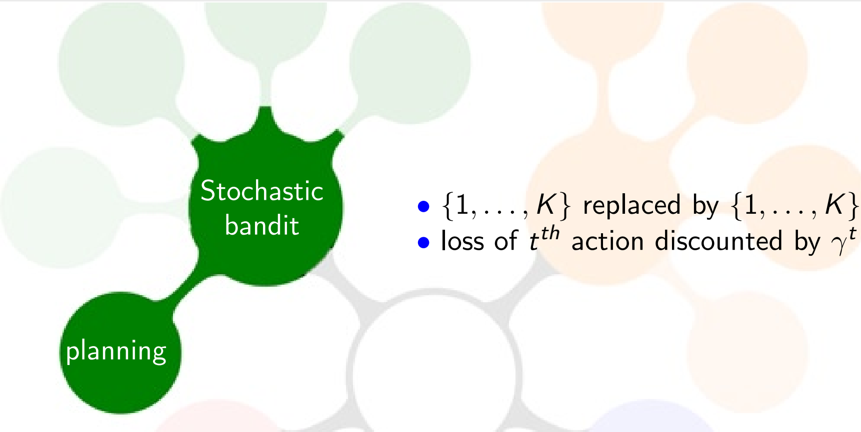

- $\{1, \ldots, K\}$ replaced by $\{1, \ldots, K\}^*$
- loss of t^{th} action discounted by γ^t

Theorem (Bubeck and Munos [2010])

$$\inf \sup R_n = \begin{cases} \tilde{\Theta}\left(n^{1-\frac{\log 1/\gamma}{\log K}}\right) & \textit{if } \gamma\sqrt{K} > 1 \\ \tilde{\Theta}\left(\sqrt{n}\right) & \textit{if } \gamma\sqrt{K} \leq 1 \end{cases}$$

Other contributions to bandit theory

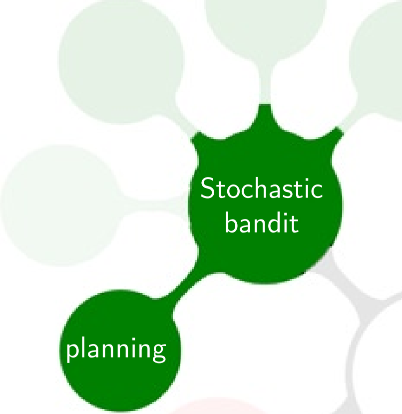

- $\{1, \ldots, K\}$ replaced by $\{1, \ldots, K\}^*$
- loss of t^{th} action discounted by γ^t

Theorem (Bubeck and Munos [2010])

$$\inf \sup R_n = \begin{cases} \tilde{\Theta}\left(n^{1-\frac{\log 1/\gamma}{\log K}}\right) & \text{if } \gamma\sqrt{K} > 1 \\ \tilde{\Theta}\left(\sqrt{n}\right) & \text{if } \gamma\sqrt{K} \leq 1 \end{cases}$$

Other contributions to bandit theory

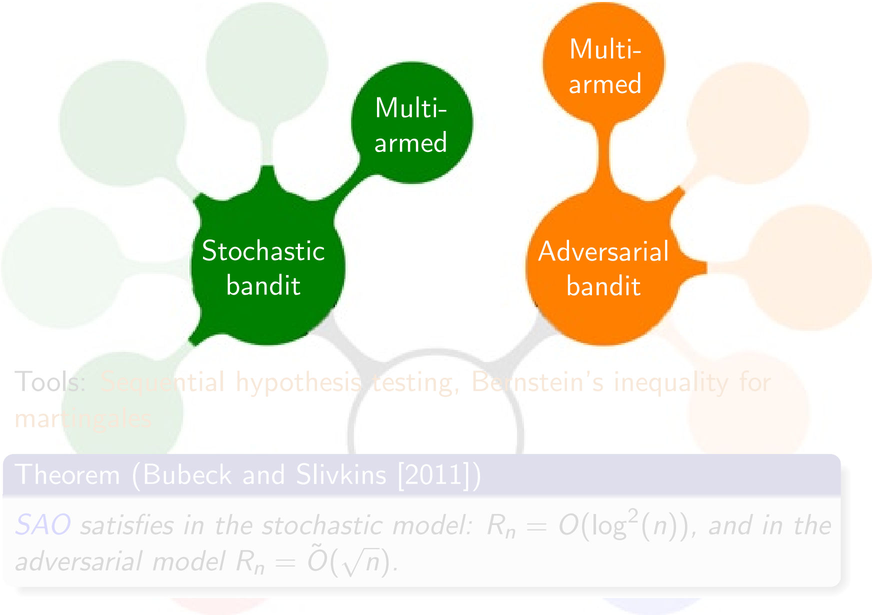

Tools: Sequential hypothesis testing, Bernstein's inequality for martingales

Theorem (Bubeck and Slivkins [2011])

SAO satisfies in the stochastic model: $R_n = O(\log^2(n))$, and in the adversarial model $R_n = \tilde{O}(\sqrt{n})$.

Other contributions to bandit theory

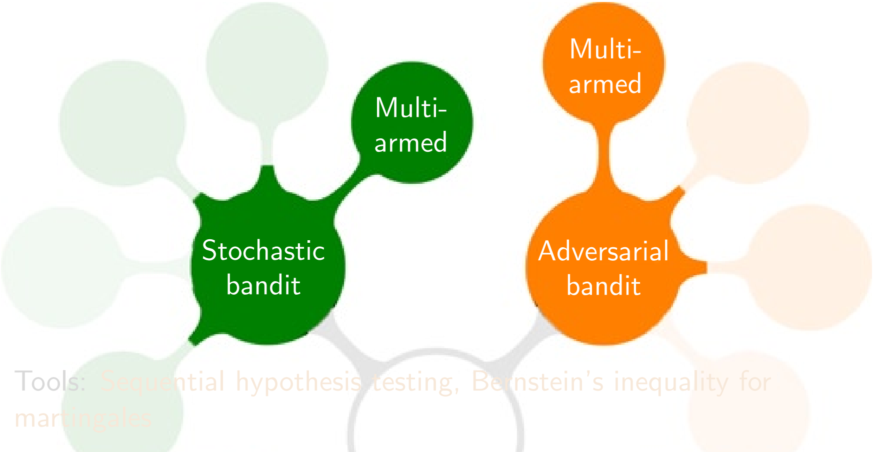

Tools: Sequential hypothesis testing, Bernstein's inequality for martingales

Theorem (Bubeck and Slivkins [2011])

SAO satisfies in the stochastic model: $R_n = O(\log^2(n))$*, and in the adversarial model* $R_n = \tilde{O}(\sqrt{n})$*.*

Other contributions to bandit theory

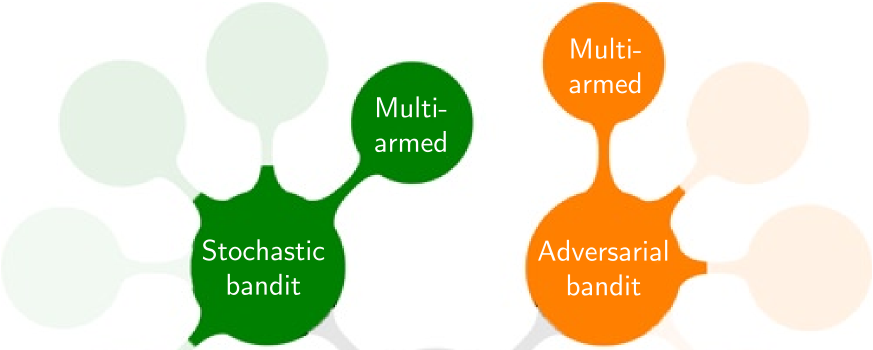

Tools: Sequential hypothesis testing, Bernstein's inequality for martingales

Theorem (Bubeck and Slivkins [2011])

SAO satisfies in the stochastic model: $R_n = O(\log^2(n))$, and in the adversarial model $R_n = \tilde{O}(\sqrt{n})$.

Other contributions to bandit theory

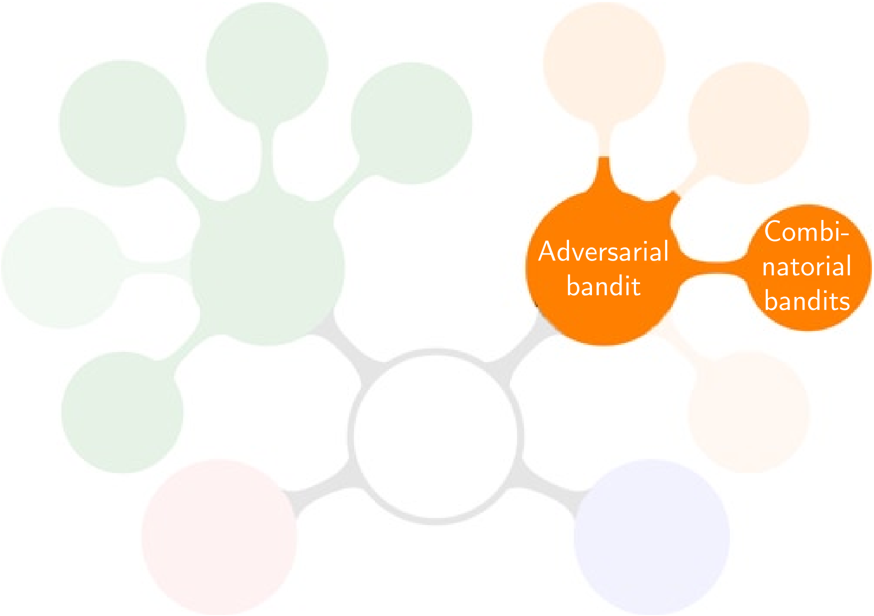

Other contributions to bandit theory

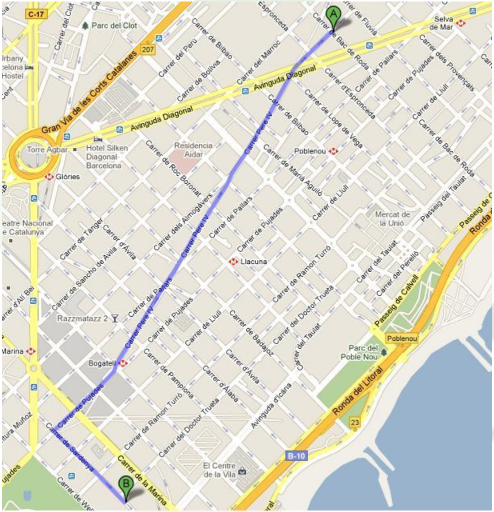

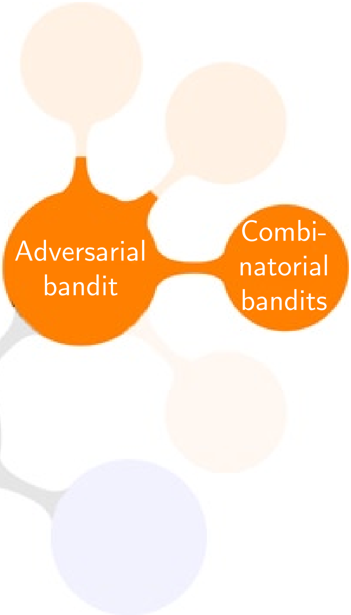

Combinatorial prediction game

Adversary

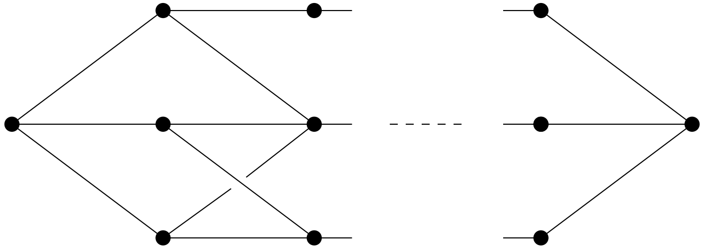

Player

Combinatorial prediction game

Adversary

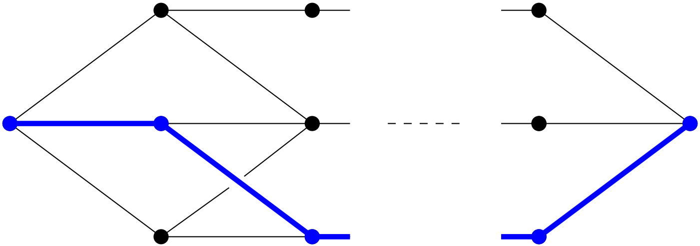

Player ⟶

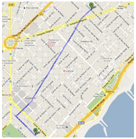

Combinatorial prediction game

Adversary ⟶ 

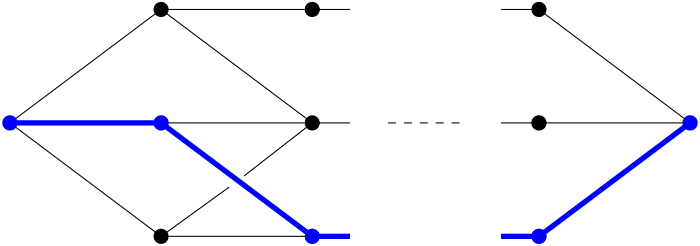

Player ⟶ 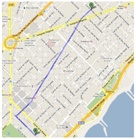

Combinatorial prediction game

Adversary ⟶ 

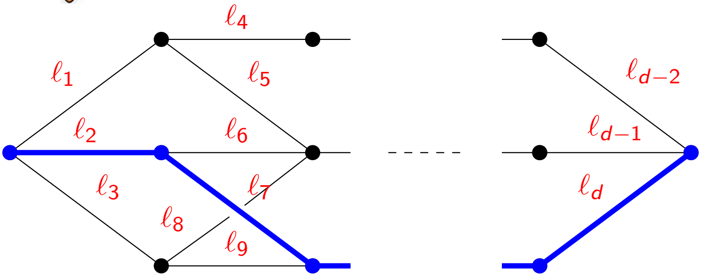

Player ⟶ 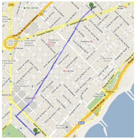

Combinatorial prediction game

Adversary ⟶

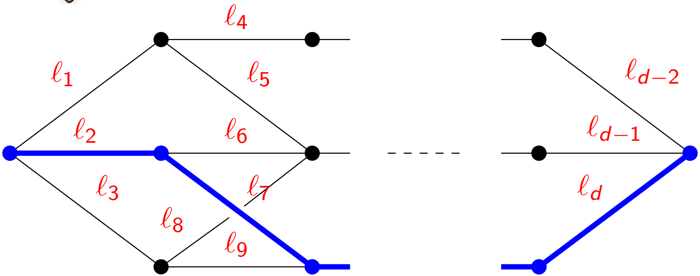

Player ⟶

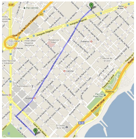

loss suffered: $\ell_2 + \ell_7 + \ldots + \ell_d$

Combinatorial prediction game

Adversary ⟶ 

Feedback: { Full Info: $\ell_1, \ell_2, \ldots, \ell_d$

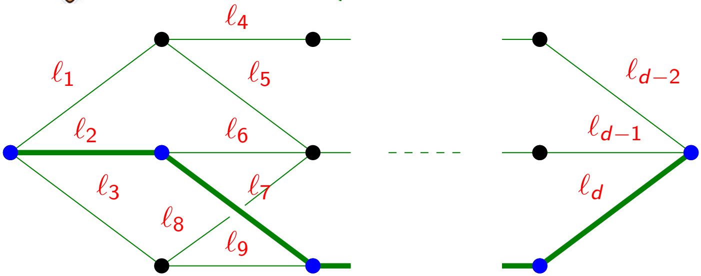

Player ⟶ 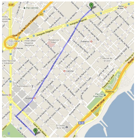

loss suffered: $\ell_2 + \ell_7 + \ldots + \ell_d$

Combinatorial prediction game

Adversary ⟶  Feedback:

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

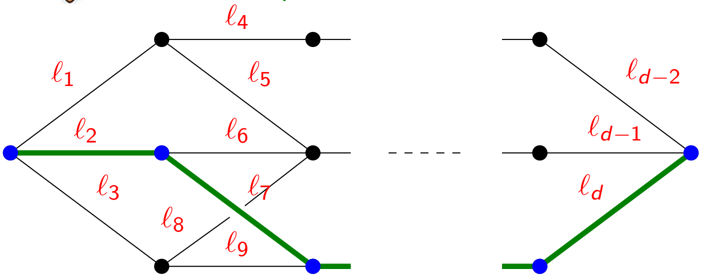

Player ⟶ 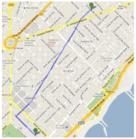

loss suffered: $\ell_2 + \ell_7 + \ldots + \ell_d$

Combinatorial prediction game

Adversary ⟶ 

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

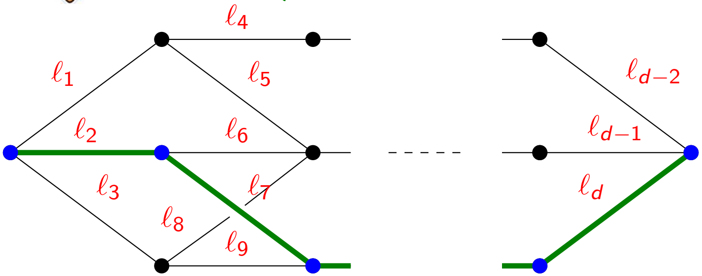

Player ⟶ 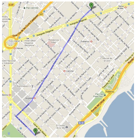

loss suffered: $\ell_2 + \ell_7 + \ldots + \ell_d$

Notations

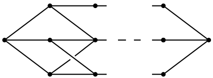

$\leftrightsquigarrow$ $\mathcal{S} \subset \{0,1\}^d$

$\leftrightsquigarrow$ $\ell_t \in [0,1]^d$

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

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.

Notations

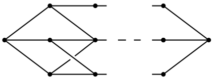

$\leftrightsquigarrow$ $\mathcal{S} \subset \{0,1\}^d$

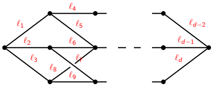

$\leftrightsquigarrow$ $\ell_t \in [0,1]^d$

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

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 $P_t = \mathbb{E}_{V \sim p_t}(VV^T)$ in the bandit game.

Notations

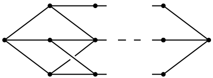

$\leftrightsquigarrow$ $\mathcal{S} \subset \{0,1\}^d$

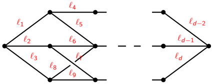

$\leftrightsquigarrow$ $\ell_t \in [0,1]^d$

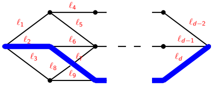

$\leftrightsquigarrow$ $V_t \in \mathcal{S}$, loss suffered: $\ell_t^T V_t$.

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 $P_t = \mathbb{E}_{V \sim p_t}(VV^T)$ in the bandit game.

Notations

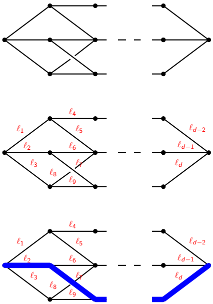

$\rightsquigarrow$ $\mathcal{S} \subset \{0,1\}^d$

$\rightsquigarrow$ $\ell_t \in [0,1]^d$

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

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.

Notations

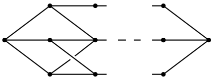

$\leftrightsquigarrow$ $\mathcal{S} \subset \{0,1\}^d$

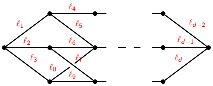

$\leftrightsquigarrow$ $\ell_t \in [0,1]^d$

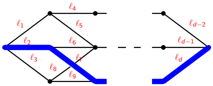

$\leftrightsquigarrow$ $V_t \in \mathcal{S}$, loss suffered: $\ell_t^T V_t$.

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 $P_t = \mathbb{E}_{V \sim p_t}(VV^T)$ in the bandit game.

Notations

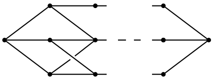

$\leftrightsquigarrow$ $\mathcal{S} \subset \{0,1\}^d$

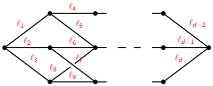

$\leftrightsquigarrow$ $\ell_t \in [0,1]^d$

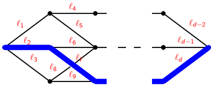

$\leftrightsquigarrow$ $V_t \in \mathcal{S}$, loss suffered: $\ell_t^T V_t$.

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 $P_t = \mathbb{E}_{V \sim p_t}(VV^T)$ in the bandit game.

Notations

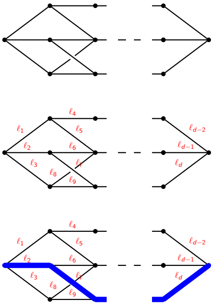

$\rightsquigarrow$ $\mathcal{S} \subset \{0,1\}^d$

$\rightsquigarrow$ $\ell_t \in [0,1]^d$

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

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.

Notations

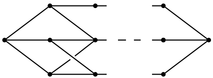

$\leftrightsquigarrow$ $\mathcal{S} \subset \{0,1\}^d$

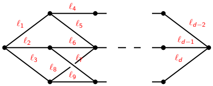

$\leftrightsquigarrow$ $\ell_t \in [0,1]^d$

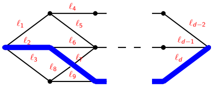

$\leftrightsquigarrow$ $V_t \in \mathcal{S}$, loss suffered: $\ell_t^T V_t$.

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 $P_t = \mathbb{E}_{V \sim p_t}(VV^T)$ in the bandit game.

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), Audibert, Bubeck and Lugosi [2011]

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

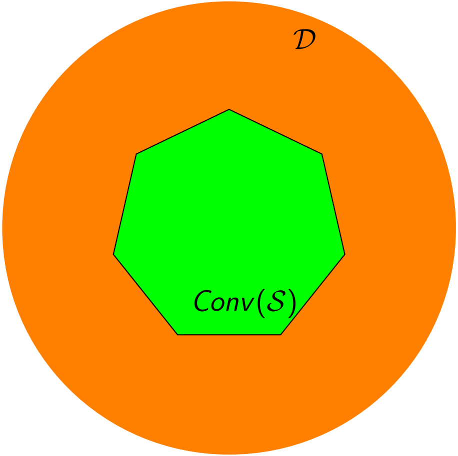

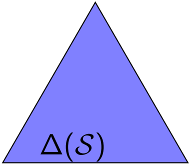

CLEB (Combinatorial LEarning with Bregman divergences), Audibert, Bubeck and Lugosi [2011]

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

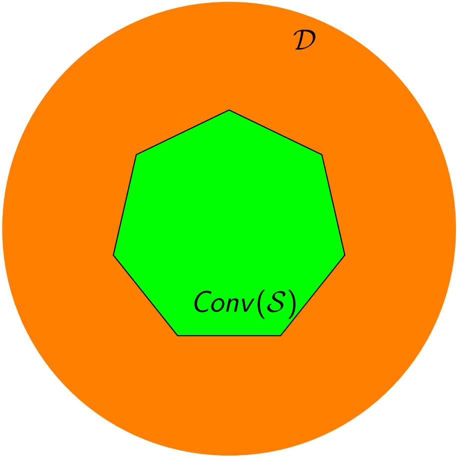

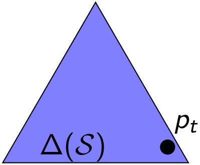

CLEB (Combinatorial LEarning with Bregman divergences), Audibert, Bubeck and Lugosi [2011]

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

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

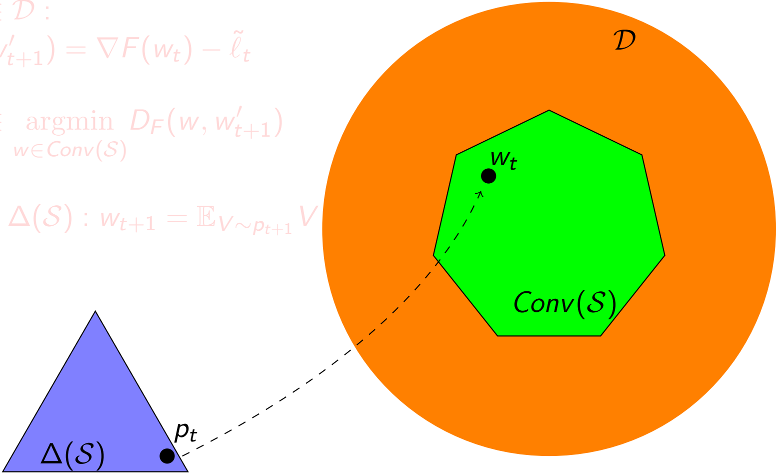

CLEB (Combinatorial LEarning with Bregman divergences), Audibert, Bubeck and Lugosi [2011]

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

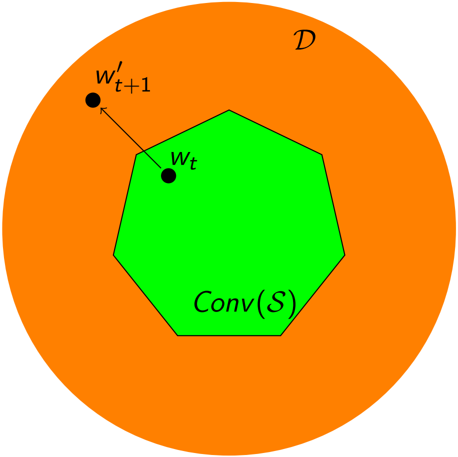

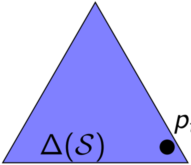

CLEB (Combinatorial LEarning with Bregman divergences), Audibert, Bubeck and Lugosi [2011]

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

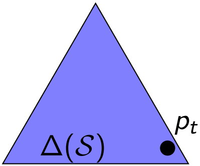

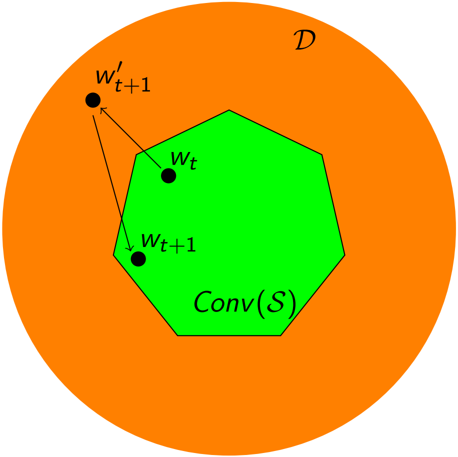

CLEB (Combinatorial LEarning with Bregman divergences), Audibert, Bubeck and Lugosi [2011]

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

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

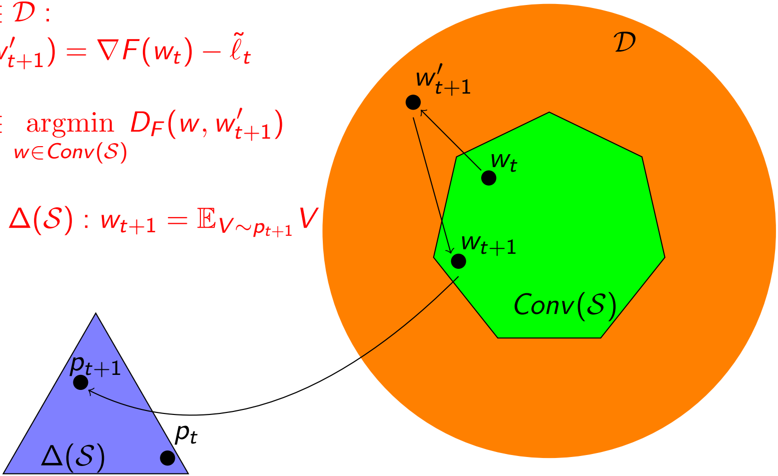

General regret bound for CLEB

Theorem (Audibert, Bubeck and Lugosi [2011])

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

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

Key tool: Pythagorean theorem for Bregman divergences

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

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)

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

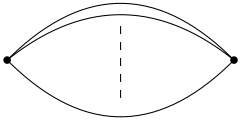

Full Info: Hedge
Semi-Bandit=Bandit: Exp3

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)

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

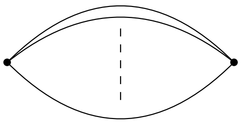

- Full Info: Hedge
- Semi-Bandit=Bandit: Exp3

- 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)

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

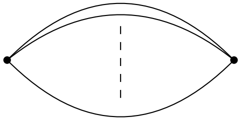

- Full Info: Hedge
- Semi-Bandit=Bandit: Exp3

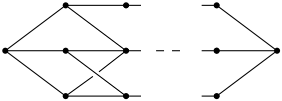

- Full Info: Component Hedge
  Koolen, Warmuth and Kivinen [2010]
- Semi-Bandit: MW
  Kale, Reyzin and Schapire [2010]
- Bandit: new algorithm

Different instances of CLEB: LinExp (Entropy Function)

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

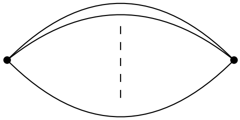

- Full Info: Hedge
- Semi-Bandit=Bandit: Exp3

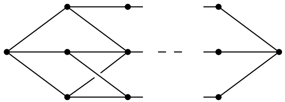

- Full Info: Component Hedge
  Koolen, Warmuth and Kivinen [2010]
- Semi-Bandit: MW
  Kale, Reyzin and Schapire [2010]
- Bandit: new algorithm

Different instances of CLEB: LinExp (Entropy Function)

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

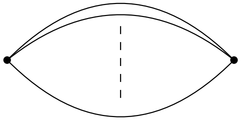

- Full Info: Hedge
- Semi-Bandit=Bandit: Exp3

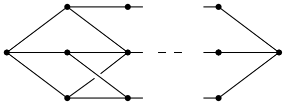

- Full Info: Component Hedge
  Koolen, Warmuth and Kivinen [2010]
- Semi-Bandit: MW
  Kale, Reyzin and Schapire [2010]
- Bandit: new algorithm

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)

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

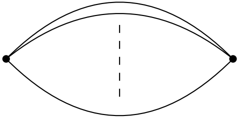

INF, Audibert and Bubeck [2009]

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

Different instances of CLEB: LinINF (Exchangeable Hessian)

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

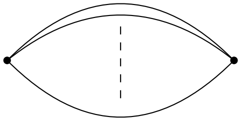

INF, Audibert and Bubeck [2009]

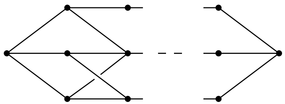

$\begin{cases} \psi(x) = \exp(\eta x) : \text{LinExp} \\ \psi(x) = (-\eta x)^{-q}, q > 1 : \text{LinPoly} \end{cases}$

Different instances of CLEB: LinINF (Exchangeable Hessian)

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

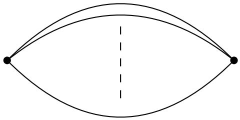

INF, Audibert and Bubeck [2009]

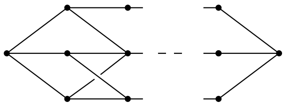

$$\begin{cases} \psi(x) = \exp(\eta x) : \text{LinExp} \\ \psi(x) = (-\eta x)^{-q}, q > 1 : \text{LinPoly} \end{cases}$$

Different instances of CLEB: Follow the regularized leader

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

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

Strong connections with interior-point methods

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

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

Strong connections with interior-point methods

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

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

Strong connections with interior-point methods

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

Minimax regret for combinatorial prediction games

$$\overline{R}_n = \inf_{strategy} \max_{\mathcal{S} \subset \{0,1\}^d} \sup_{adversaries} R_n$$

Theorem (Audibert, Bubeck and Lugosi [2011])

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

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

and in the bandit game:

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

Minimax regret for combinatorial prediction games

$$\overline{R}_n = \inf_{strategy} \max_{\mathcal{S} \subset \{0,1\}^d} \sup_{adversaries} R_n$$

Theorem (Audibert, Bubeck and Lugosi [2011])

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

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

and in the bandit game:

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

New project: Combinatorial testing

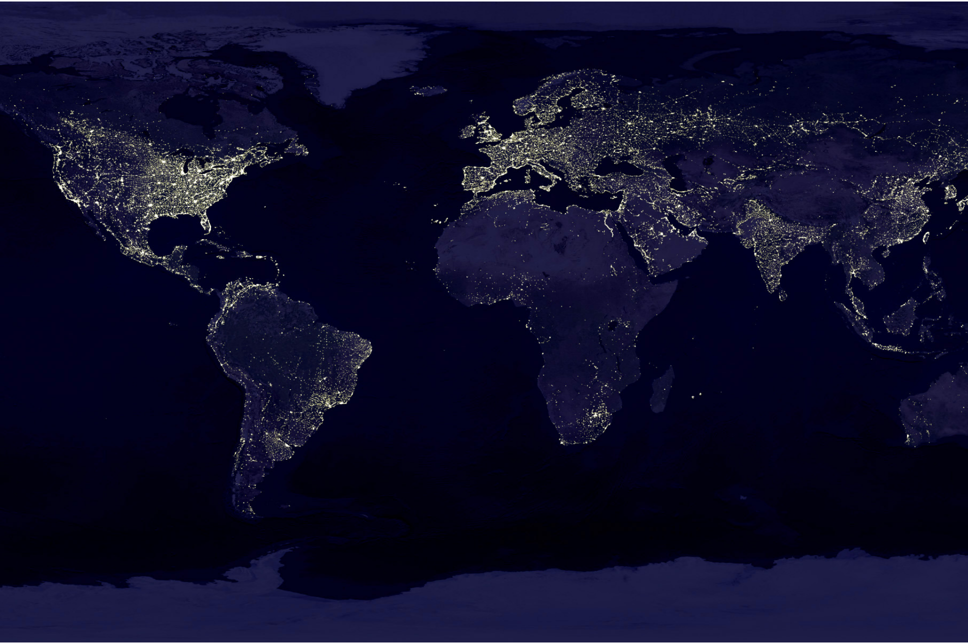

New project: Combinatorial testing

Set of concepts: $\mathcal{S} \subset \{0,1\}^d$

Paths

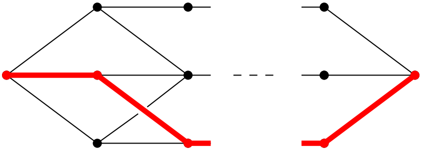

k-sets

k-sized intervals

Spanning trees

New project: Combinatorial testing

Set of concepts: $\mathcal{S} \subset \{0,1\}^d$

Paths

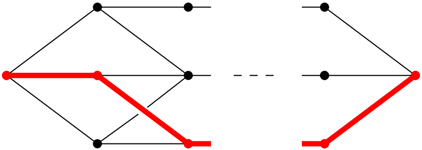

k-sets

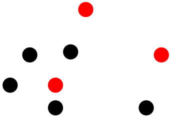

k-sized intervals

Spanning trees

New project: Combinatorial testing

Set of concepts: $\mathcal{S} \subset \{0,1\}^d$

Paths

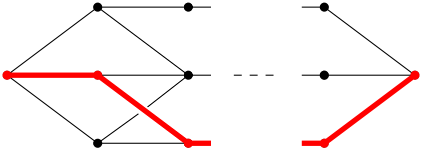

k-sets

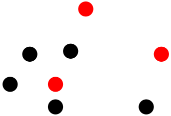

k-sized intervals

Spanning trees

New project: Combinatorial testing

Set of concepts: $\mathcal{S} \subset \{0,1\}^d$

Paths

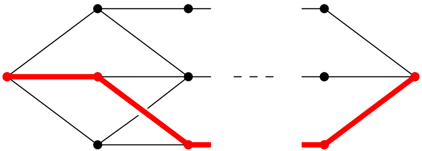

k-sets

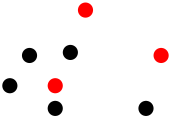

k-sized intervals

Spanning trees

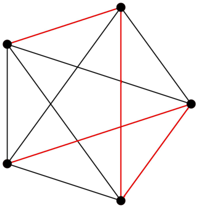

New project: Combinatorial testing

- Set of concepts: $\mathcal{S} \subset \{0,1\}^d$

Paths

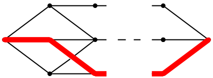

k-sets

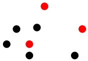

k-sized intervals

Spanning trees

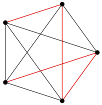

- Data: $X \in \mathbb{R}^d$
- Hypotheses:

$$H_0 : \text{"nothing special happens in } X\text{"}$$
$$H_1 : \exists C \in \mathcal{S} \text{ s.t "something special happens on } X|_C\text{"}$$

New project: Combinatorial testing

- Set of concepts: $\mathcal{S} \subset \{0,1\}^d$

Paths

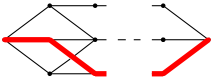

k-sets

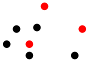

k-sized intervals

Spanning trees

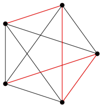

- Data: $X \in \mathbb{R}^d$
- Hypotheses:

$$H_0 : \text{“nothing special happens in } X\text{”}$$
$$H_1 : \exists C \in \mathcal{S} \text{ s.t “something special happens on } X|_C \text{”}$$

New project: Combinatorial testing

- Set of concepts: $\mathcal{S} \subset \{0,1\}^d$

Paths

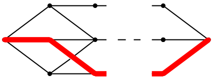

k-sets

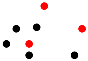

k-sized intervals

Spanning trees

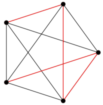

- Data: $X \in \mathbb{R}^d$
- Hypotheses:

$$H_0 : \text{“nothing special happens in } X\text{''}$$

$$H_1 : \exists C \in \mathcal{S} \text{ s.t “something special happens on } X|_C \text{''}$$

Two examples of combinatorial testing problems

- Simultaneous tests: $|\mathcal{S}| = 1$, Fan, Hall and Yao [2008]
- Detection of elevated mean:

$$H_0 : X \sim \mathcal{N}(0, I_d)$$
$$H_1 : \exists C \in \mathcal{S} \text{ such that } X \sim \mathcal{N}(\mu \mathbb{1}_C, I_d)$$

For k-sets: problem suggested by Tukey, analyzed in Donoho and Jin [2002].
General framework introduced in Arias-Castro, Candès, Helgason and Zeitouni [2008].

- Detection of combinatorial correlation, Arias-Castro, Bubeck and Lugosi [2011]: $X_i \sim \mathcal{N}(0,1), i \in \{1, \dots, d\}$

$$H_0 : \mathbb{E}(X_i X_j) = 0$$
$$H_1 : \exists C \in \mathcal{S} \text{ such that } \mathbb{E}(X_i X_j) = \rho \, \mathbb{1}_{i \neq j, i,j \in C}$$

Two examples of combinatorial testing problems

- Simultaneous tests: $|\mathcal{S}| = 1$, Fan, Hall and Yao [2008]
- Detection of elevated mean:

$$H_0 : X \sim \mathcal{N}(0, I_d)$$
$$H_1 : \exists C \in \mathcal{S} \text{ such that } X \sim \mathcal{N}(\mu \mathbb{1}_C, I_d)$$

For k-sets: problem suggested by Tukey, analyzed in Donoho and Jin [2002].
General framework introduced in Arias-Castro, Candès, Helgason and Zeitouni [2008].

- Detection of combinatorial correlation, Arias-Castro, Bubeck and Lugosi [2011]: $X_i \sim \mathcal{N}(0, 1), i \in \{1, \dots, d\}$

$$H_0 : \mathbb{E}(X_i X_j) = 0$$
$$H_1 : \exists C \in \mathcal{S} \text{ such that } \mathbb{E}(X_i X_j) = \rho \, \mathbb{1}_{i \neq j, i,j \in C}$$

Two examples of combinatorial testing problems

- Simultaneous tests: $|\mathcal{S}| = 1$, Fan, Hall and Yao [2008]
- Detection of elevated mean:

$$H_0 : X \sim \mathcal{N}(0, I_d)$$
$$H_1 : \exists C \in \mathcal{S} \text{ such that } X \sim \mathcal{N}(\mu \mathbb{1}_C, I_d)$$

For k-sets: problem suggested by Tukey, analyzed in Donoho and Jin [2002].
General framework introduced in Arias-Castro, Candès, Helgason and Zeitouni [2008].

- Detection of combinatorial correlation, Arias-Castro, Bubeck and Lugosi [2011]: $X_i \sim \mathcal{N}(0,1), i \in \{1, \ldots, d\}$

$$H_0 : \mathbb{E}(X_i X_j) = 0$$
$$H_1 : \exists C \in \mathcal{S} \text{ such that } \mathbb{E}(X_i X_j) = \rho \, \mathbb{1}_{i \neq j, i,j \in C}$$

Few tests for detection of combinatorial correlation

$$Z_C = X^T(A_C^{-1} - I_n)X, \qquad (A_C)_{i,j} = \mathbb{1}_{i=j} + \rho\, \mathbb{1}_{i\neq j, i,j\in C}$$

- Optimal test: Likelihood ratio test

$$\text{Reject if} \qquad \sum_{C\in\mathcal{S}} \exp\left(-\frac{1}{2}Z_C\right) > \text{threshold}$$

- Generalized Likelihood Ratio Test (GLRT):

$$\text{Reject if} \qquad \max_{C\in\mathcal{S}} -\frac{1}{2}Z_C > \text{threshold}$$

- Scan statistics:

$$\text{Reject if} \qquad \max_{C\in\mathcal{S}} \sum_{i\neq j, i,j\in C} X_i X_j > \text{threshold}$$

- Squared norm test:

$$\text{Reject if} \qquad ||X||_2 > \text{threshold}$$

Few tests for detection of combinatorial correlation

$$Z_C = X^T(A_C^{-1} - I_n)X, \qquad (A_C)_{i,j} = \mathbb{1}_{i=j} + \rho\, \mathbb{1}_{i \neq j, i,j \in C}$$

- Optimal test: Likelihood ratio test

$$\text{Reject if} \qquad \sum_{C \in \mathcal{S}} \exp\left(-\frac{1}{2} Z_C\right) > \text{threshold}$$

- Generalized Likelihood Ratio Test (GLRT):

$$\text{Reject if} \qquad \max_{C \in \mathcal{S}} -\frac{1}{2} Z_C > \text{threshold}$$

- Scan statistics:

$$\text{Reject if} \qquad \max_{C \in \mathcal{S}} \sum_{i \neq j, i,j \in C} X_i X_j > \text{threshold}$$

- Squared norm test:

$$\text{Reject if} \qquad ||X||_2 > \text{threshold}$$

Few tests for detection of combinatorial correlation

$$Z_C = X^T(A_C^{-1} - I_n)X, \qquad (A_C)_{i,j} = \mathbb{1}_{i=j} + \rho\, \mathbb{1}_{i \neq j, i,j \in C}$$

- Optimal test: Likelihood ratio test

$$\text{Reject if} \qquad \sum_{C \in \mathcal{S}} \exp\left(-\frac{1}{2} Z_C\right) > \text{threshold}$$

- Generalized Likelihood Ratio Test (GLRT):

$$\text{Reject if} \qquad \max_{C \in \mathcal{S}} -\frac{1}{2} Z_C > \text{threshold}$$

- Scan statistics:

$$\text{Reject if} \qquad \max_{C \in \mathcal{S}} \sum_{i \neq j, i,j \in C} X_i X_j > \text{threshold}$$

- Squared norm test:

$$\text{Reject if} \qquad ||X||_2 > \text{threshold}$$

Few tests for detection of combinatorial correlation

$$Z_C = X^T(A_C^{-1} - I_n)X, \qquad (A_C)_{i,j} = \mathbb{1}_{i=j} + \rho\, \mathbb{1}_{i\neq j, i,j\in C}$$

- Optimal test: Likelihood ratio test

$$\text{Reject if} \qquad \sum_{C\in\mathcal{S}} \exp\left(-\frac{1}{2}Z_C\right) > \text{threshold}$$

- Generalized Likelihood Ratio Test (GLRT):

$$\text{Reject if} \qquad \max_{C\in\mathcal{S}} -\frac{1}{2}Z_C > \text{threshold}$$

- Scan statistics:

$$\text{Reject if} \qquad \max_{C\in\mathcal{S}} \sum_{i\neq j, i,j\in C} X_i X_j > \text{threshold}$$

- Squared norm test:

$$\text{Reject if} \qquad ||X||_2 > \text{threshold}$$

Few tests for detection of combinatorial correlation

$$Z_C = X^T(A_C^{-1} - I_n)X, \qquad (A_C)_{i,j} = \mathbb{1}_{i=j} + \rho\, \mathbb{1}_{i\neq j, i,j\in C}$$

- Optimal test: Likelihood ratio test

$$\text{Reject if} \qquad \sum_{C\in\mathcal{S}} \exp\left(-\frac{1}{2}Z_C\right) > \text{threshold}$$

- Generalized Likelihood Ratio Test (GLRT):

$$\text{Reject if} \qquad \max_{C\in\mathcal{S}} -\frac{1}{2}Z_C > \text{threshold}$$

- Scan statistics:

$$\text{Reject if} \qquad \max_{C\in\mathcal{S}} \sum_{i\neq j, i,j\in C} X_i X_j > \text{threshold}$$

- Squared norm test:

$$\text{Reject if} \qquad ||X||_2 > \text{threshold}$$

Preliminary results for detection of combinatorial correlation

	k-sized intervals	k sets
Optimal test	Powerless if $\rho << \frac{\log(d/k)}{k}$	Conjecture: Powerless if $k << \sqrt{d}$
GLRT	Powerful if $\rho >> \frac{\log(d)}{k}$	Conjecture: Powerless if $k << d$
Scan statistics	-	Conjecture: Powerless if $k << d$
Squared norm test	Powerful iff $\rho >> \frac{\sqrt{d}}{k}$	Powerful iff $\rho >> \frac{\sqrt{d}}{k}$

Preliminary results for detection of combinatorial correlation

	k-sized intervals	k sets
Optimal test	Powerless if $\rho << \frac{\log(d/k)}{k}$	Conjecture: Powerless if $k << \sqrt{d}$
GLRT	Powerful if $\rho >> \frac{\log(d)}{k}$	Conjecture: Powerless if $k << d$
Scan statistics	-	Conjecture: Powerless if $k << d$
Squared norm test	Powerful iff $\rho >> \frac{\sqrt{d}}{k}$	Powerful iff $\rho >> \frac{\sqrt{d}}{k}$

Preliminary results for detection of combinatorial correlation

	k-sized intervals	k sets
Optimal test	Powerless if $\rho << \frac{\log(d/k)}{k}$	Conjecture: Powerless if $k << \sqrt{d}$
GLRT	Powerful if $\rho >> \frac{\log(d)}{k}$	Conjecture: Powerless if $k << d$
Scan statistics	-	Conjecture: Powerless if $k << d$
Squared norm test	Powerful iff $\rho >> \frac{\sqrt{d}}{k}$	Powerful iff $\rho >> \frac{\sqrt{d}}{k}$

Preliminary results for detection of combinatorial correlation

	k-sized intervals	k sets
Optimal test	Powerless if $\rho << \frac{\log(d/k)}{k}$	Conjecture: Powerless if $k << \sqrt{d}$
GLRT	Powerful if $\rho >> \frac{\log(d)}{k}$	Conjecture: Powerless if $k << d$
Scan statistics	-	Conjecture: Powerless if $k << d$
Squared norm test	Powerful iff $\rho >> \frac{\sqrt{d}}{k}$	Powerful iff $\rho >> \frac{\sqrt{d}}{k}$

Preliminary results for detection of combinatorial correlation

	k-sized intervals	k sets
Optimal test	Powerless if $\rho << \frac{\log(d/k)}{k}$	Conjecture: Powerless if $k << \sqrt{d}$
GLRT	Powerful if $\rho >> \frac{\log(d)}{k}$	Conjecture: Powerless if $k << d$
Scan statistics	-	Conjecture: Powerless if $k << d$
Squared norm test	Powerful iff $\rho >> \frac{\sqrt{d}}{k}$	Powerful iff $\rho >> \frac{\sqrt{d}}{k}$

Preliminary results for detection of combinatorial correlation

	k-sized intervals	k sets
Optimal test	Powerless if $\rho << \frac{\log(d/k)}{k}$	Conjecture: Powerless if $k << \sqrt{d}$
GLRT	Powerful if $\rho >> \frac{\log(d)}{k}$	Conjecture: Powerless if $k << d$
Scan statistics	-	Conjecture: Powerless if $k << d$
Squared norm test	Powerful iff $\rho >> \frac{\sqrt{d}}{k}$	Powerful iff $\rho >> \frac{\sqrt{d}}{k}$

Preliminary results for detection of combinatorial correlation

	k-sized intervals	k sets
Optimal test	Powerless if $\rho << \frac{\log(d/k)}{k}$	Conjecture: Powerless if $k << \sqrt{d}$
GLRT	Powerful if $\rho >> \frac{\log(d)}{k}$	Conjecture: Powerless if $k << d$
Scan statistics	-	Conjecture: Powerless if $k << d$
Squared norm test	Powerful iff $\rho >> \frac{\sqrt{d}}{k}$	Powerful iff $\rho >> \frac{\sqrt{d}}{k}$

Preliminary results for detection of combinatorial correlation

	k-sized intervals	k sets
Optimal test	Powerless if $\rho << \frac{\log(d/k)}{k}$	Conjecture: Powerless if $k << \sqrt{d}$
GLRT	Powerful if $\rho >> \frac{\log(d)}{k}$	Conjecture: Powerless if $k << d$
Scan statistics	-	Conjecture: Powerless if $k << d$
Squared norm test	Powerful iff $\rho >> \frac{\sqrt{d}}{k}$	Powerful iff $\rho >> \frac{\sqrt{d}}{k}$

Perspectives

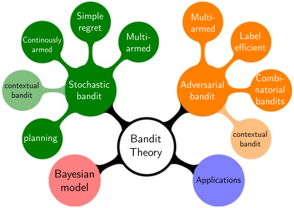

Lots of unexplored extensions, both important for applications and mathematically elegant

Bandit d-gap $\begin{cases} \sqrt{d} \text{ related to random walks on graphs} \\ \sqrt{d} \text{ related to interior-point methods} \end{cases}$

Perspectives

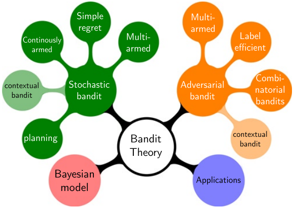

Lots of unexplored extensions, both important for applications and mathematically elegant

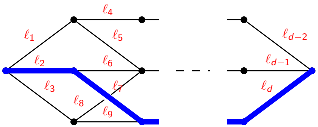

Bandit d-gap $\begin{cases} \sqrt{d} \text{ related to random walks on graphs} \\ \sqrt{d} \text{ related to interior-point methods} \end{cases}$

Perspectives

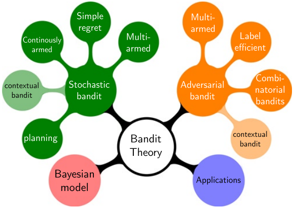

Lots of unexplored extensions, both important for applications and mathematically elegant

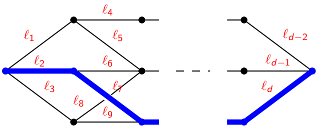

Bandit d-gap
- $\sqrt{d}$ related to random walks on graphs
- $\sqrt{d}$ related to interior-point methods

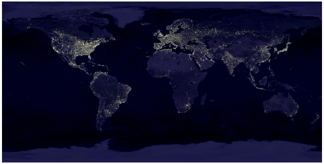

Detection of combinatorial correlation

Combinatorial LASSO?

Preprints

S. Bubeck and A. Slivkins, The best of both worlds: an adaptive strategy for stochastic and adversarial multi-armed bandits , submitted to **COLT 2011**

J.Y. Audibert, S. Bubeck and G. Lugosi, Minimax Policies for Combinatorial Prediction Games, submitted to **COLT 2011**

Journal Papers

S. Bubeck, N. Cesa-Bianchi and G. Lugosi, Bandit Theory–A Survey, to appear in **Foundations and Trends in Machine Learning**, 2011.

S. Bubeck, R. Munos, G. Stoltz and C. Szepesvari, X-Armed Bandits, **JMLR** (Journal of Machine Learning Research), 2011

J.Y. Audibert and S. Bubeck, Regret Bounds and Minimax Policies under Partial Monitoring, **JMLR**, 2010

S. Bubeck, R. Munos and G. Stoltz, Pure Exploration in Finitely-Armed and Continuously-Armed Bandits, **Theoretical Computer Science**, 2011

S. Bubeck and U. von Luxburg, Nearest Neighbor Clustering: A Baseline Method for Consistent Clustering with Arbitrary Objective Functions , **JMLR**, 2009

Conference Papers (Acceptance ratio NIPS $\sim 25\%$, COLT $\sim 35\%$)

J.Y. Audibert, S. Bubeck and R. Munos, Best Arm Identification in Multi-Armed Bandits, **COLT 2010**

S. Bubeck and R. Munos, Open-Loop Optimistic Planning, **COLT 2010**

J.Y. Audibert and S. Bubeck, Minimax Policies for Adversarial and Stochastic Bandits, **COLT 2009 (Best Student Paper Award)**

S. Bubeck, R. Munos and G. Stoltz, Pure Exploration in Multi-Armed Bandit Problems, **ALT 2009**

S. Bubeck, R. Munos, G. Stoltz and C. Szepesvari, Online Optimization in X-Armed Bandits, **NIPS 2008**

U. von Luxburg, S. Bubeck, S. Jegelka and M. Kaufmann, Consistent Minimization of Clustering Objective Functions, **NIPS 2007**

PhD Thesis, Book Chapters, Technical Reports

J.Y. Audibert, S. Bubeck and R. Munos, Bandit View on Noisy Optimization, in **Optimization for Machine Learning, MIT press, 2010**

S. Bubeck, Bandits Games and Clustering Foundations, **PhD dissertation**, 2010 (**runner-up for the Gilles Kahn prize 2010**)

S. Bubeck, M. Meila and U. von Luxburg, How the Initialization Affects the Stability of the k-means Algorithm, **ArXiv Report, 2009**