Part 2, continued: Lower bounds on the regret for adversarial bandits

$\longleftrightarrow$ Minimax lower bounds on the regret for stochastic bandits

We discussed so far $\mathbb{E}_{\nu}[N_a(T)]$ in terms of distribution-dependent bounds like

$$C_a(\nu) \ln T$$

What about distribution-free / minimax bounds?

To that end we restrict our attention to the model $\mathcal{D} = \mathcal{P}([0,1])$, the set of all probability distributions over $[0,1]$.

Stochastic bandits

With each arm a is associated $\nu_a \in \mathcal{P}([0,1])$

For $t = 1, 2 \dots$

- The decision maker picks $I_t \in \{1..K\}$
- Her reward Y_t, which is such that $Y_t \mid I_t \sim \nu_{I_t}$, is her only piece of information

Aim: control the regret

$$\bar{R}_T = T \max_{a=1..K} E(\nu_a) - \mathbb{E}\left[\sum_{t=1}^{T} Y_t\right]$$

Adversarial bandits

An opponent selects the payoffs $g_{j,t}$

For $t = 1, 2 \dots$

- The opponent picks $(g_{1t} \dots g_{Kt}) \in [0,1]^K$
- while, simultaneously, The decision-maker picks $I_t \in \{1..K\}$
- Her payoff is $g_{I_t,t}$ and this is the only piece of information she gets

Aim: control the regret

$$R_T = \max_{k=1..K} \sum_{t=1}^{T} g_{kt} - \sum_{t=1}^{T} g_{I_t,t}$$

Typical adversarial results (Auer, Cesa-Bianchi, Freund, Schapire, 2002, later improved by Audibert and Bubeck, 2009):

Strategies such that for all opponents picking gains in $[0,1]$, for all $T \geq 1$,

$$\begin{cases} \text{with probability at least } 1-\delta, & R_T \leq C\sqrt{TK\ln(K/\delta)} \\ \mathbb{E}[R_T] \leq C\sqrt{TK\ln K} \end{cases}$$

for some numerical constant C

where the probability and $\mathbb{E}$ are wrt decision-maker's internal randomization

For "oblivious" opponents (ie, when the g_{kt} do not "react" to the decision-maker's actions): the $\sqrt{\ln K}$ can be dropped.

It is in particular the case when $g_{kt} \sim \nu_k$ $\forall t$ in an independent way

In this stochastic model:

$$\mathbb{E}[R_T] = \mathbb{E}\Big[\max_{k=1..K} \sum_{t=1}^{T} g_{k,t}\Big] - \mathbb{E}\Big[\sum_{t=1}^{T} g_{I_t,t}\Big] \qquad g_{I_t,t} \text{ is } Y_t$$

$$\geq T \max_{k=1..K} \mathbb{E}[g_{k,1}] - \mathbb{E}\Big[\sum_{t=1}^{T} Y_t\Big]$$

$$= T \max_{k=1..K} E(\nu_k) - \mathbb{E}\Big[\sum_{t=1}^{T} Y_t\Big] = \overline{R}_T$$

The adversarial results entail in particular that there exists a strategy of the decision-maker such that

$$\sup_{\nu_1 \dots \nu_K \in \mathcal{P}([0,1])} \overline{R}_T \leq \sup_{\substack{\text{opponents} \\ \text{picking } g_{j,t} \in [0,1]}} \mathbb{E}[R_T] \leq C\sqrt{TK}$$

for some numerical constant C

Lower bound:

<u>Theorem</u>: For all (randomized) strategies of the decision-maker, for all $K \geq 2$ and $T \geq K \ln 2 / 2$,

$$\sup_{\substack{\text{opponents} \\ \text{picking } g_{j,t} \in [0,1]}} \mathbb{E}[R_T] \geq \sup_{\substack{\text{individual} \\ \text{sequences } g_{j,t} \in [0,1]}} \mathbb{E}[R_T] \geq \sup_{\nu_1, \dots \nu_K \in \mathcal{P}([0,1])} \overline{R}_T \geq \frac{1}{20}\sqrt{TK}$$

the only inequality we need to prove!

Proof: (i) Bandit problem $\nu^{(0)} = (\mathrm{Ber}(1/2), \dots \mathrm{Ber}(1/2))$

versus bandit problems $\nu^{(i)} = (\mathrm{Ber}(1/2), \dots \mathrm{Ber}(1/2), \mathrm{Ber}(1/2 + \varepsilon), \mathrm{Ber}(1/2), \dots)$

for $i \in \{1..K\}$ and $\varepsilon \in (0, 1/2)$ — in i-th position

There exists $k \in \{1..K\}$ such that $\mathbb{E}_{\nu^{(0)}}[N_k(T)] \leq T/K$. For this k:

The fundamental inequality indicates that

$$\sum_{a=1}^{K} \mathbb{E}_{\nu^{(0)}}[N_a(T)]\, \mathrm{KL}(\nu_a^{(0)}, \nu_a^{(k)}) \geq \mathrm{kl}\left(\mathbb{E}_{\nu^{(0)}}\Big[\frac{N_k(T)}{T}\Big], \mathbb{E}_{\nu^{(k)}}\Big[\frac{N_k(T)}{T}\Big]\right)$$

$$= \mathbb{E}_{\nu^{(0)}}[N_k(T)]\, \mathrm{KL}(\nu_k^{(0)}, \nu_k^{(k)}) \leq T/K \;\; \mathrm{kl}(1/2, 1/2 + \varepsilon)$$

Using Pinsker's inequality:

$$\mathrm{kl}\left(\mathbb{E}_{\nu^{(0)}}\Big[\frac{N_k(T)}{T}\Big], \mathbb{E}_{\nu^{(k)}}\Big[\frac{N_k(T)}{T}\Big]\right) \geq 2\left(\mathbb{E}_{\nu^{(0)}}\Big[\frac{N_k(T)}{T}\Big] - \mathbb{E}_{\nu^{(k)}}\Big[\frac{N_k(T)}{T}\Big]\right)^2$$

Solving for $\mathbb{E}_{\nu^{(k)}}\left[\frac{N_k(T)}{T}\right]$:

$$\mathbb{E}_{\nu^{(k)}}\left[\frac{N_k(T)}{T}\right] \leq \mathbb{E}_{\nu^{(0)}}\left[\frac{N_k(T)}{T}\right] + \sqrt{\frac{T}{2K}\, kl\left(\tfrac{1}{2}, \tfrac{1}{2}+\varepsilon\right)}$$

$$\leq \frac{1}{K} + \sqrt{\frac{T}{2K}\, kl\left(\tfrac{1}{2}, \tfrac{1}{2}+\varepsilon\right)}$$

$$\leq \frac{1}{K} + \sqrt{(2\ln 2)\frac{T\varepsilon^2}{K}} \qquad \text{for } \varepsilon \in (0, 1/2\sqrt{2}]$$

as $\varepsilon \in (0, 1/2\sqrt{2}]$ entails $4\varepsilon^2 \leq 1/2$, as $\ln(1/1-u) \leq (2\ln 2)u \quad \forall u \in (0, 1/2]$

and $kl(1/2, 1/2-\varepsilon) = \frac{1}{2}\ln\left(\frac{1}{1-2\varepsilon}\right) + \frac{1}{2}\ln\left(\frac{1}{1+2\varepsilon}\right) = \frac{1}{2}\ln\frac{1}{1-4\varepsilon^2}$

$$\leq (4\ln 2)\,\varepsilon^2$$

(2) Let's go back to the regret.

In $\nu^{(k)}$: $\quad \overline{R}_T = \sum_{a\neq k} \underbrace{\varepsilon}_{\text{gap of arm } a} \underbrace{\mathbb{E}_{\nu^{(k)}}[N_a(T)]}_{\text{number of times } a \text{ is pulled}} = \varepsilon\left(T - \mathbb{E}_{\nu^{(k)}}[N_k(T)]\right)$

$$\geq T\varepsilon\left(\underbrace{1 - \frac{1}{K}}_{\geq 1/2} - \sqrt{(2\ln 2)\,T\varepsilon^2/K}\right)$$

to be optimized over $\varepsilon \in (0, 1/2\sqrt{2}]$

Optimal ε s.t. $\frac{1}{2} - 2\varepsilon\sqrt{(2\ln 2)T/K} = 0$, ie, $\varepsilon = \frac{1}{4} \times 1/\sqrt{(2\ln 2)T/K}$

which is $\leq 1/2\sqrt{2}$ as soon as $T\ln 2/K \geq 1/2$,

For this ε, the lower bound is ie $T \geq K\ln 2/2$

$$T\varepsilon\left(1/2 - \underbrace{\varepsilon\sqrt{(2\ln 2)T/K}}_{=1/4}\right) = T\varepsilon/4 = \frac{1}{16\sqrt{2\ln 2}}\sqrt{TK} \geq \frac{1}{20}\sqrt{TK}$$

<u>Note</u> about the constants: By further restricting the considered T (ie, by imposing $T \geq \gamma K$ with a larger γ), we can improve the $\frac{1}{20}\sqrt{TK}$ bound (ie, get $c'\sqrt{TK}$ with $c' > 1/20$).

This is done in the upper bound on $\ln(1/1-u)$: eg, $\leq 4\ln(4/3)\,u$ for $u \in [0, 1/4]$

↳ Entails: $\sup \overline{R}_T \geq \frac{\sqrt{2}-1}{\sqrt{32\ln(4/3)}}\sqrt{TK} \geq 0{,}136\sqrt{TK}$ for $T \geq \frac{1}{4\ln(4/3)}K \simeq 0.87\,K$

Open question:

So, the minimax rates for the regret for stochastic bandits or oblivious opponents are $\sqrt{TK}$.

What is the minimax rate against general, reactive opponents?

→ Should the upper bound be improved?

→ Should the proof technique be improved? (in particular, look for sequences of payoffs $g_{k,t}$ with real and strong correlations/ dependencies in the past).

Intermezzo: Another open problem — minimax order for internal regret

Setting: adversarial losses with full information (as when we discussed sparse vectors of losses)

For each $t = 1, 2 \dots$

- The environment picks $(\ell_{1t} \dots \ell_{Nt}) \in [0,1]^N$

  while simultaneously, the decision-maker picks $I_t \in \{1, \dots N\}$ possibly at random according to $U_t \sim \mathcal{U}_{[0,1]}$

- Both players observe $(\ell_{1t} \dots \ell_{Nt})$ and I_t

Aim of the decision-maker: minimize, i.e., control from above, the internal regret:

$$R_T^{int} = \max_{i \neq j} \sum_{t=1}^{T} (\ell_{it} - \ell_{jt}) \mathbb{1}_{\{I_t = i\}}$$

Stronger notion than the classical regret:

$$R_T = \sum_{t=1}^{T} \ell_{I_t, t} - \min_{k} \sum_{t=1}^{T} \ell_{kt}$$

$$= \sum_{i=1}^{N} \sum_{t=1}^{T} \mathbb{1}_{\{I_t = i\}} \ell_{it} - \min_{k=1..N} \sum_{i=1}^{N} \sum_{t=1}^{T} \ell_{kt} \mathbb{1}_{\{I_t = i\}}$$

min of a sum is larger than the sum of the min

$$\leq \sum_{i=1}^{N} \left(\sum_{t=1}^{T} \mathbb{1}_{\{I_t = i\}} \ell_{it} - \min_{k} \sum_{t=1}^{T} \ell_{kt} \mathbb{1}_{\{I_t = i\}} \right)$$

$$= \sum_{i=1}^{N} \Bigg(\underbrace{\max_{k} \sum_{t=1}^{T} (\ell_{it} - \ell_{kt}) \mathbb{1}_{\{I_t = i\}}}_{\leq R_T^{int}} \Bigg) \leq N R_T^{int}$$

Upper bounds:

high-probability bounds on R_T^{int} can be deduced via the Hoeffding-Azuma lemma

On $$\overline{R}_T^{int} = \max_{i \neq j} \sum_{t=1}^{T} (\ell_{it} - \ell_{jt}) \, \mathbb{P}\{I_t = i \mid H_{t-1}\}$$

where $H_{t-1} = \big((\ell_{js})_{j \leq N}, U_s \big)_{s \leq t-1}$

information available at the beginning of round t

* For small T, drawing I_t uniformly at random:

$$\overline{R}_T^{int} \leq \frac{T}{N}$$

* We will rather be interested in the large T regime:

<u>Theorem</u> (Foster & Vohra, 1998): There exists a strategy of the decision-maker such that

$$\forall N \geq 2, \quad \forall T \geq 1, \qquad \sup_{\substack{\text{opponents} \\ \text{picking } \ell_{jt} \in [0,1]}} \overline{R}_T^{\,int} \leq \sqrt{2T \ln N} \qquad (*)$$

* For fixed-in-advance sequences $(\ell_{jt})_{j,t}$ [not reacting to what the decision-maker plays: "individual sequences"]

$$\sigma(H_{t-1}) = \sigma\Big(\big((\ell_{js})_{j \leq N}, U_s \big)_{s \leq t-1} \Big) = \sigma(U_1 \dots U_{t-1})$$

$\perp\!\!\!\perp$ I_t, which is then U_t-measurable

Then:

$$\overline{R}_T^{\,int} = \max_{i \neq j} \sum_{t=1}^{T} (\ell_{it} - \ell_{jt}) \, \underbrace{\mathbb{P}\{I_t = i\}}_{\text{wrt } U_t \text{ only}},$$

which is the quantity we will control in the lower bound, not $\mathbb{E}[R_T^{int}]$

Note: $\mathbb{E}[R_T^{int}] = \mathbb{E}\Big[\max_{i \neq j} \sum_{t=1}^{T} (\ell_{it} - \ell_{jt}) \mathbb{1}_{\{I_t = i\}} \Big]$

$> \overline{R}_T^{\,int}$ in general, even for individual sequences of losses!

<u>Lower bound:</u>

<u>Theorem</u> (Stoltz, 2005): For all (randomized) strategies of the decision-maker, for all $N \geq 2$ and all $T \geq (8Nc)^2$,

$$\sup_{\substack{\text{individual sequences} \\ \ell_{jt} \in [0,1]}} \left\{ \max_{i \neq j} \sum_{t=1}^{T} (\ell_{it} - \ell_{jt}) \, \mathbb{P}\{I_t = i\} \right\} \geq c\sqrt{T} \qquad (**)$$

for some numerical constant c ↗ in particular for $c = \frac{\sqrt{\ln 2}}{128}$.

<u>Discussion:</u> A $\sqrt{\ln N}$ factor is missing between the upper bound $(*)$ and the lower bound $(**)$. The minimax order for internal regret against non-stochastic opponents (individual sequences or "true" opponents) is still unknown. What Gerchinovitz (2011) proved, however: there exists a strategy s.t.

$$\sup_{\substack{\nu_1 \dots \nu_N \in \mathcal{P}([0,1]) \\ \text{s.t. } \ell_{jt} \text{ iid} \sim \nu_j}} \mathbb{E}\Big[\max_{i \neq j} \sum_{t=1}^{T} (\ell_{it} - \ell_{jt}) \, \mathbb{P}\{I_t = i\} \Big] \leq C\sqrt{T}$$

stochastic opponents, as in bandit problems

wrt the random choices of the losses

wrt auxiliary randomization

for some numerical constant C

He did so because (**) is proved by lower bounding the sup over individual sequences in (**) by $\sup_{\nu_1 \dots \nu_k} \mathbb{E}[\quad]$. If the lower bound had to be improved (but who knows whether the upper bound or the lower bound should be improved?), then a completely different proof technique (non-iid sequences of losses!) would be required.

<u>Proof of (**)</u>: We take: $\nu_1 = \mathrm{Ber}(1/2)$

$\nu_3, \dots \nu_N = \delta_1$ Dirac mass at 1

only ν_2 varies: $\mathrm{Ber}(1/2)$, $\mathrm{Ber}(1/2 - \varepsilon)$, $\mathrm{Ber}(1/2 + \varepsilon)$

We denote by $\mathbb{E}_0$, $\mathbb{E}_{-\varepsilon}$, $\mathbb{E}_\varepsilon$ the corresponding expectations wrt random choices of the losses.

We have $\mathbb{E}_0 = \frac{1}{2}\left(\mathbb{E}_{-\varepsilon} + \mathbb{E}_\varepsilon\right)$

$$\sup_{\substack{\text{individual sequences}\\ \ell_{jt} \in [0,1]}} \left\{ \max_{i \neq j} \sum_{t=1}^T (\ell_{it} - \ell_{jt})\, \mathbb{P}\{I_t = i\} \right\}$$

$$\geq \mathbb{E}_0\left[\max_{i \neq j} \sum_{t=1}^T (\ell_{it} - \ell_{jt})\, \mathbb{P}\{I_t = i\}\right] \qquad (***)$$

$$= \frac{1}{2}\, \mathbb{E}_\varepsilon\left[\max_{i \neq j} \sum_{t=1}^T (\ell_{it} - \ell_{jt})\, \mathbb{P}\{I_t = i\}\right] + \frac{1}{2}\, \mathbb{E}_{-\varepsilon}\left[\max_{i \neq j} \sum_{t=1}^T (\ell_{it} - \ell_{jt})\, \mathbb{P}\{I_t = i\}\right]$$

$$\geq \frac{1}{2}\, \mathbb{E}_\varepsilon\left[\sum_{t=1}^T (\ell_{2t} - \ell_{1t})\, \mathbb{P}\{I_t = 1\}\right] + \frac{1}{2}\, \mathbb{E}_{-\varepsilon}\left[\sum_{t=1}^T (\ell_{1t} - \ell_{2t})\, \mathbb{P}\{I_t = 2\}\right]$$

By the tower rule:
$$\mathbb{E}_\varepsilon\Big[\underbrace{(\ell_{2t} - \ell_{1t})}_{\text{is } \perp\!\!\!\perp \sigma(H_{t-1})} \underbrace{\mathbb{P}\{I_t = 1\}}_{\text{is } \sigma(H_{t-1})\text{-measurable}} \Big] = \varepsilon \quad \mathbb{P}_\varepsilon \otimes \mathbb{P}\{I_t = 1\}$$

where $H_{t-1} = \left((\ell_{js})_{j \leq N}, U_s\right)_{s \leq t-1}$

We get the lower bound, using the notation $N_j(T) = \sum_{t=1}^T \mathbb{1}_{\{I_t = j\}}$:

$$\frac{1}{2}\left(\mathbb{E}_\varepsilon \otimes \mathbb{E}[N_1(T)] + \mathbb{E}_{-\varepsilon} \otimes \mathbb{E}[N_2(T)]\right)$$

$$= \frac{\varepsilon}{2}\Big(T - \mathbb{E}_\varepsilon \otimes \mathbb{E}[N_2(T)] - \sum_{j=3}^N \mathbb{E}_\varepsilon \otimes \mathbb{E}[N_j(T)]$$

$$+ T - \mathbb{E}_{-\varepsilon} \otimes \mathbb{E}[N_1(T)] - \sum_{j=3}^N \mathbb{E}_{-\varepsilon} \otimes \mathbb{E}[N_j(T)]\Big)$$

$$= \varepsilon\left(T - \frac{1}{2}\left(\mathbb{E}_{\varepsilon} \otimes \mathbb{E}[N_2(T)] + \mathbb{E}_{-\varepsilon} \otimes \mathbb{E}[N_1(T)]\right) - \sum_{j=3}^{N} \mathbb{E}_0 \otimes \mathbb{E}[N_j(T)]\right)$$

$\mathbb{E}_0$ is the average of $\mathbb{E}_{\varepsilon}$ and $\mathbb{E}_{-\varepsilon}$

Now, by the tower rule:

$$\frac{1}{2} \mathbb{E}_0 \otimes \mathbb{E}[N_j(T)] = \mathbb{E}_0 \otimes \mathbb{E}\left[\sum_{t=1}^{T} (\ell_{jt} - \ell_{1t}) \mathbb{1}_{\{I_t = j\}} \right]$$

$$\leq c\sqrt{T}$$

with no loss of generality; otherwise, $(\overset{*}{**})$ is larger than $c\sqrt{T}$ and the proof is concluded

Thus, $$(\overset{*}{**}) \geq \varepsilon\left(T - \frac{1}{2}\left(\mathbb{E}_{\varepsilon} \otimes \mathbb{E}[N_2(T)] + \mathbb{E}_{-\varepsilon} \otimes \mathbb{E}[N_1(T)]\right) - 2(N-2)c\sqrt{T} \right)$$

(Fano-) Pinsker inequality: [with common alternative $\mathbb{P}_0 \otimes \mathbb{P}$]

$$\frac{1}{2}\left(\mathbb{E}_{\varepsilon} \otimes \mathbb{E}\left[\frac{N_2(T)}{T}\right] + \mathbb{E}_{-\varepsilon} \otimes \mathbb{E}\left[\frac{N_1(T)}{T}\right]\right)$$

$$\leq \frac{1}{2}\left(\mathbb{E}_0 \otimes \mathbb{E}\left[\frac{N_2(T)}{T}\right] + \mathbb{E}_0 \otimes \mathbb{E}\left[\frac{N_1(T)}{T}\right]\right)$$

we denote this by α

Classical Pinsker would be more efficient but I use Fano for the sake of comparison.

$$+ \sqrt{\frac{\overline{K}}{-\ln \alpha}}$$

where $\overline{K} = \frac{1}{2}\Big(KL\big(\text{law of } (\ell_{jt})_{j,t} \text{ under } \mathbb{P}_{\varepsilon}, \text{ law of } (\ell_{jt})_{j,t} \text{ under } \mathbb{P}_0\big) + KL\big(\text{law of } (\ell_{jt})_{j,t} \text{ under } \mathbb{P}_{-\varepsilon}, \text{ law of } (\ell_{jt})_{j,t} \text{ under } \mathbb{P}_0\big)\Big)$

cf. independence, and many components are equal, etc.

$$= \frac{1}{2}\Big(T\, KL\big(Ber(\tfrac{1}{2}+\varepsilon), Ber(\tfrac{1}{2})\big) + T\, KL\big(Ber(\tfrac{1}{2}-\varepsilon), Ber(\tfrac{1}{2})\big)\Big)$$

$$= \frac{T}{2}\Big((\tfrac{1}{2}+\varepsilon)\ln(1+2\varepsilon) + (\tfrac{1}{2}-\varepsilon)\ln(1-2\varepsilon)$$

twice the same quantity

$$+ (\tfrac{1}{2}-\varepsilon)\ln(1-2\varepsilon) + (\tfrac{1}{2}+\varepsilon)\ln(1+2\varepsilon)\Big)$$

$\ln(1+u) \leq u.$

$$\leq \frac{T}{2} \times 2 \times \Big((\tfrac{1}{2}+\varepsilon)(2\varepsilon) + (\tfrac{1}{2}-\varepsilon)(-2\varepsilon)\Big) = 4T\varepsilon^2$$

and where $$\alpha = \frac{1}{2} \mathbb{E}_0 \otimes \mathbb{E}\Bigg[\underbrace{\frac{N_2(T)+N_1(T)}{T}}_{\leq 1} \Bigg] \leq \frac{1}{2}$$

probably not improvable!

Substituting above: $$(\overset{*}{**}) \geq \varepsilon\left(T - \frac{T}{2} - T\sqrt{\frac{4T\varepsilon^2}{\ln 2}} - 2Nc\sqrt{T}\right)$$

$$= T\varepsilon\left(\frac{1}{2} - \sqrt{\frac{4T\varepsilon^2}{\ln 2}} - \underbrace{2Nc/\sqrt{T}}_{\leq 1/4}\right) \quad \text{if } T \geq (8Nc)^2$$

$$\geq T\varepsilon\left(\tfrac{1}{4} - \varepsilon\sqrt{4T/\ln 2}\right)$$

to be optimized over ε: ε^* s.t. $\frac{1}{4} - 2\varepsilon^*\sqrt{4T/\ln 2} = 0$, $\quad \varepsilon^* = \frac{\sqrt{\ln 2}}{16} \frac{1}{\sqrt{T}}$

Yields : $\binom{*}{**} \geq \frac{\sqrt{\ln 2}}{16} \sqrt{T} \left(1/4 - \underbrace{\frac{\sqrt{\ln 2}}{16} \sqrt{\frac{4}{\ln 2}}}_{=1/8} \right) = \underbrace{\frac{\sqrt{\ln 2}}{128}}_{=c} \sqrt{T}$

<u>ENSAE students needing a grade :</u>

I'll suggest you work on an extension of internal regret called the swap regret.

Bonus part : Discussion of the Fano-type inequalities and applications by Duchi & Wainwright (2013), Chen, Guntuboyina & Zhang (2016)

Generalized Fano's inequality of Chen, Guntuboyina & Zhang (2016)

Setting : $(\Theta, \mathcal{G})$ a parameter space with prior ν, family of distributions $(\mathbb{P}_\vartheta)_\vartheta$

Loss function $L : \Theta \times A \to [0,1]$, measurable

where A can be Θ or an action space

Observations $X \sim \mathbb{P}_\vartheta$, estimators $\hat{\vartheta} = \hat{\vartheta}(X) \in A$
or measurable actions

Goal : Control the Bayes risk : $R_{bayes} = \inf_{\hat{\vartheta}} \int_\Theta \mathbb{E}_\vartheta [L(\vartheta, \hat{\vartheta})] \, d\nu(\vartheta)$

expectation wrt $\mathbb{P}_\vartheta$

Theorem : If L take values in $\{0,1\}$ only,

$$R_{bayes} \geq 1 + \frac{\left(\inf_{\mathbb{Q}} \int_\Theta KL(\mathbb{P}_\vartheta, \mathbb{Q}) \, d\nu(\vartheta) \right) + \ln\left(2 - \sup_{a \in A} \nu\{L(\cdot, a) = 0\}\right)}{\ln\left(\sup_{a \in A} \nu\{L(\cdot,a) = 0\}\right)}$$

We can actually prove a more general statement, for $[0,1]$-valued loss functions :

Theorem :
$$R_{bayes} \geq 1 + \frac{\left(\inf_{\mathbb{Q}} \int_\Theta KL(\mathbb{P}_\vartheta, \mathbb{Q}) \, d\nu(\vartheta) \right) + \ln\left(1 + \inf_{a \in A} \int_\Theta L(\vartheta, a) \, d\nu(\vartheta)\right)}{\ln\left(1 - \inf_{a \in A} \int_\Theta L(\vartheta,a) \, d\nu(\vartheta)\right)}$$

Chen, Guntuboyina & Zhang (2016) deal with the general $[0,1]$-valued case by replacing a general L by $L_t = t \, \mathbb{1}_{\{L \geq t\}}$, which is a lower bound on L :

$$R_{bayes}(L) \geq t \, R_{bayes}(L_t)$$

Then, they have to "optimize" over t...

They present some applications of their bound in their paper.

<u>Proof:</u> Fix any estimator $\hat{\theta}$,

fix any alternative $\mathbb{Q}$.

By the data-processing inequality, applied to $Z_\theta = 1 - L(\theta, \hat{\theta}) \in [0,1]$:

$$\mathrm{kl}\left(\int_\Theta \mathbb{E}_\theta\left[1 - L(\theta,\hat{\theta})\right] d\nu(\theta), \int_\Theta \mathbb{E}_{\mathbb{Q}}\left[1 - L(\theta,\hat{\theta})\right] d\nu(\theta) \right) \leqslant \int_\Theta \mathrm{KL}(\mathbb{P}_\theta, \mathbb{Q})\, d\nu(\theta)$$

Via $\mathrm{kl}(p,q) \geqslant p \ln \frac{1}{q} - \ln(2-q)$ we have $p \leqslant \dfrac{\mathrm{kl}(p,q) + \ln(2-q)}{\ln(1/q)}$

that is, here:

$$1 - \int_\Theta \mathbb{E}_\theta\left[L(\theta,\hat{\theta})\right] d\nu(\theta) \leqslant \frac{\int_\Theta \mathrm{KL}(\mathbb{P}_\theta, \mathbb{Q})\, d\nu(\theta) + \ln(2 - q_{\hat{\theta}})}{\ln(1/q_{\hat{\theta}})}$$

with $q = \int_\Theta \mathbb{E}_{\mathbb{Q}}\left[1 - L(\theta,\hat{\theta})\right] d\nu(\theta)$

Now, $u \longmapsto \dfrac{K + \ln(2-u)}{\ln(1/u)}$ is increasing, as $u \mapsto \dfrac{1}{\ln(1/u)}$

and $u \mapsto \dfrac{\ln(2-u)}{\ln(1/u)}$ both are.

We thus take the supremum over $\hat{\theta}$ in both sides:

$$1 - R_{\mathrm{bayes}} = 1 - \inf_{\hat{\theta}} \int_\Theta \mathbb{E}_\theta\left[L(\theta,\hat{\theta})\right] d\nu(\theta) = \sup_{\hat{\theta}} \left\{ 1 - \int_\Theta \mathbb{E}_\theta\left[L(\theta,\hat{\theta})\right] d\nu(\theta) \right\}$$

$$\leqslant \sup_{\hat{\theta}} \frac{\int_\Theta \mathrm{KL}(\mathbb{P}_\theta, \mathbb{Q})\, d\nu(\theta) + \ln(2 - q_{\hat{\theta}})}{\ln(1/q_{\hat{\theta}})}$$

$$= \frac{\int_\Theta \mathrm{KL}(\mathbb{P}_\theta, \mathbb{Q})\, d\nu(\theta) + \ln(2 - q^*)}{\ln(1/q^*)}$$

where $q^* = \sup_{\hat{\theta}} q_{\hat{\theta}} = 1 - \inf_{\hat{\theta}} \int_\Theta \mathbb{E}_{\mathbb{Q}}\left[L(\theta,\hat{\theta})\right] d\nu(\theta)$

$$= 1 - \inf_{a \in A} \int_\Theta L(\theta, a)\, d\nu(\theta).$$

Indeed, by Tonelli's theorem,

$$\forall \hat{\theta}, \quad \int_\Theta \mathbb{E}_Q\left[L(\vartheta, \hat{\theta})\right] d\nu(\vartheta) = \int_A \left(\int_\Theta L(\vartheta, \hat{\theta}(x))\, d\nu(\vartheta) \right) dQ(x)$$

$$\geq \int_A \left(\inf_a \int_\Theta L(\vartheta, a)\, d\nu(\vartheta) \right) dQ(x)$$

$$= \inf_{a \in A} \int_\Theta L(\vartheta, a)\, d\nu(\vartheta)$$

Thus,

$$\inf_{\hat{\theta}} \int_\Theta \mathbb{E}_Q\left[L(\vartheta, \hat{\theta})\right] d\nu(\vartheta) = \inf_{a \in A} \int_\Theta \mathbb{E}_Q\left[L(\vartheta, a)\right] d\nu(\vartheta)$$

$$= \inf_{a \in A} \int_\Theta L(\vartheta, a)\, d\nu(\vartheta)$$

Putting all things together, we proved:

$$\forall Q, \quad 1 - R_{bayes} \leq \frac{\int_\Theta KL(\mathbb{P}_\vartheta, Q)\, d\nu(\vartheta) + \ln\left(1 + \inf_{a \in A} \int_\Theta L(\vartheta, a)\, d\nu(\vartheta)\right)}{-\ln\left(1 - \inf_{a \in A} \int_\Theta L(\vartheta, a)\, d\nu(\vartheta)\right)}$$

Taking the infimum over Q in the right-hand side (the left-hand side is independent of Q) and rearranging concludes the proof.

A brief overview on how to prove lower bounds for nonparametric regression in $\mathbb{L}^2$

Without the need of any discretization but by picking the right prior ν (which in some sense, however, amounts to resorting to a discretization... !) } we adapt the arguments by Duchi & Wainwright (2013) Chen et al. (2016)

Setting: We observe $Y_1 \dots Y_n$ where $Y_i = f\left(\frac{i-1}{n}\right) + \varepsilon_i$

for some unknown function $f \in \mathcal{F}_L$, where $\mathcal{F}_L = \{ g : [0,1] \to \mathbb{R},\ g \text{ is } L\text{-Lipschitz} \}$

and for $\varepsilon_1 \dots \varepsilon_n$ iid $\sim \mathcal{N}(0, \sigma^2)$, with $\sigma^2 > 0$ unknown

↳ Parameters to be estimated: f, σ^2 based on $Y_1 \dots Y_n$ knowing $\mathcal{F}_L$

We focus on f and denote by $\mathbb{E}_f$, $\mathbb{P}_f$ the expectations and probability distributions when the parameters are f and σ^2.

Estimators: $\hat{f} \in \mathbb{L}^2([0,1])$

$\mathbb{L}^2$-risk: $L(f, \hat{f}) = \mathbb{E}_f\left[\|f - \hat{f}\|_2^2 \right] = \mathbb{E}_f\left[\int_0^1 (f(x) - \hat{f}(x))^2 \, dx \right]$

The lower bound was known and is tight; see Ibragimov and Has'minskii (1982, 1984) and Tsybakov (2009: the book on nonparametric regression).

Theorem: For all $n \geq \max\left\{ \dfrac{L}{\sigma\sqrt{\ln 2}},\ 64 \ln(2) \dfrac{\sigma^2}{L^2} \right\}$,

$$\inf_{\hat{f} \in \mathbb{L}^2([0,1])} \ \sup_{f \in \mathcal{F}_L} \mathbb{E}_f\left[\|f - \hat{f}\|_2^2 \right] \geq C \left(\frac{\sigma^2 L}{n} \right)^{2/3}$$

for some numerical constant $C > 0$, e.g., $C = 0.001$

Sketch of proof: (1) Let $\varphi : \mathbb{R} \to \mathbb{R}$ be

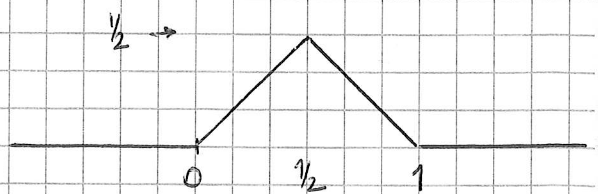

φ is 1-Lipschitz with $\|\varphi\|_2^2 = \frac{1}{12}$

For a well-chosen $d \geq 2$, the

$$f_j : x \in [0,1] \longmapsto \frac{L}{d}\, \varphi\left(d\left(x - \frac{j-1}{d}\right) \right), \qquad j = 1, \dots d$$

belong to $\mathcal{F}_L$ and form an orthogonal system in $\mathbb{L}^2([0,1])$ – since the f_j have pairwise disjoint supports.

(2) We now reduce the nonparametric problem to a parametric problem:

Define $f_\theta \overset{\text{def.}}{=} \sum_{j=1}^{d} \theta_j f_j$ for $\theta = (\theta_1, \dots \theta_d) \in \mathbb{R}^d$

Then $\theta \in \Theta \overset{\text{def.}}{=} \{ \theta \in \mathbb{R}^d : \|\theta\|_\infty \leq 1 \} \longmapsto f_\theta \in \mathcal{F}_L$ is injective

with $\forall \theta, \theta' \in \Theta$, $\|f_\theta - f_{\theta'}\|_2^2 = \sum_j (\theta_j - \theta_j')^2 \|f_j\|_2^2 = \frac{L^2}{12 d^3} \|\theta - \theta'\|_2^2$

Set $\hat{\theta} \in \underset{\theta \in \Theta}{\arg\min} \|\hat{f} - f_\theta\|_2^2$:

$$\sup_{f \in \mathcal{F}_L} \mathbb{E}_f \left[\|f - \hat{f}\|_2^2 \right] \geq \sup_{\theta \in \Theta} \mathbb{E}_{f_\theta} \left[\|f_\theta - \hat{f}\|_2^2 \right]$$

$$\geq \sup_{\theta \in \Theta} \mathbb{E}_{f_\theta} \left[\|f_\theta - f_{\hat{\theta}}\|_2^2 \right]$$

$$= \frac{L^2}{12 d^3} \sup_{\theta \in \Theta} \mathbb{E}_{f_\theta} \left[\|\theta - \hat{\theta}\|_2^2 \right]$$

(3) We apply a continuous Fano inequality after resorting to Markov's inequality first:

$$\sup_\theta \mathbb{E}_{f_\theta} \left[\|\theta - \hat{\theta}\|_2^2 \right] \geq \rho d \sup_\theta \mathbb{E}_{f_\theta} \left[\mathbb{1}_{\{ \|\theta - \hat{\theta}\|_2^2 \geq \rho d \}} \right] \quad \text{for } \rho > 0 \text{ to be determined}$$

$$\geq \rho d \left(1 - \inf_\theta \mathbb{P}_{f_\theta} \{ \|\theta - \hat{\theta}\|_2^2 \leq \rho d \} \right)$$

for any prior ν

$$\geq \rho d \left(1 - \int_\Theta \mathbb{P}_{f_\theta} \{ \|\theta - \hat{\theta}\|_2^2 \leq \rho d \} \, d\nu(\theta) \right)$$

Picking $\nu = \mathcal{U}_\Theta$ uniform prior over Θ and the alternative distribution $\mathbb{P}_{f_0}$ (where 0 is the null vector $(0, \dots 0)$ of Θ): Fano yields

$$\int_\Theta \mathbb{P}_{f_\theta} \{ \|\theta - \hat{\theta}\|_2^2 \leq \rho d \} \, d\nu(\theta) \leq \frac{\ln 2 + \int KL(\mathbb{P}_{f_\theta}, \mathbb{P}_{f_0}) \, d\nu(\theta)}{- \ln \int_\Theta \mathbb{P}_{f_0} \{ \|\theta - \hat{\theta}\|_2^2 \leq \rho d \} \, d\nu(\theta)}$$

for $\rho = 1/2\pi e$ we will show below that: $\leq \frac{1}{3} + \frac{(n+d) L^2}{8 d^3 \sigma^2 \ln 2}$

Indeed: * $\mathbb{P}_{f_\theta}$ can be identified with $\bigotimes_{i=1}^{n} \mathcal{N}\left(f_\theta\left(\frac{i-1}{n}\right), \sigma^2\right)$

$\mathbb{P}_{f_0}$ $\qquad \bigotimes_{i=1}^{n} \mathcal{N}(0, \sigma^2)$

so that by independence,
$$KL(\mathbb{P}_{f_\theta}, \mathbb{P}_{f_0}) = \sum_{i=1}^{n} KL\left(\mathcal{N}\left(f_\theta\left(\tfrac{i-1}{n}\right), \sigma^2\right), \mathcal{N}(0, \sigma^2)\right)$$
$$= \frac{1}{2\sigma^2} \sum_{i=1}^{n} \left(f_\theta\left(\tfrac{i-1}{n}\right)\right)^2$$

some omitted but simple calculations, using that the f_j have pairwise disjoint supports, that $\|f_j\|_\infty \leq L/2d$, etc.

$$\leq \dots \leq \frac{1}{2\sigma^2}\left(\frac{n}{d}+1\right)\frac{L^2}{4d}$$

* Second, $\displaystyle\int_\Theta \mathbb{P}_{f_0}\left\{\|\theta-\hat{\theta}\|_2^2 \leq \rho d\right\} d\nu(\theta)$

Tonelli's theorem
$$= \mathbb{E}_{f_0}\left[\nu\left(B(\hat{\theta}, \sqrt{\rho d}) \cap \Theta\right)\right]$$

where $B(\theta, u)$ = closed ball around θ with radius u

volumetric argument
$$\leq \frac{(\sqrt{\rho d})^d}{2^d} \operatorname{Vol}_{\mathbb{R}^d}\left(B(0,1)\right)$$
$$\leq \dots \leq 1/8 \quad \text{for } d \geq 2$$

(4) Putting all bounds together:

$$\sup_{f \in \mathcal{F}_L} \mathbb{E}_f\left[\|f-\hat{f}\|_2^2\right] \geq \frac{L^2}{12 d^3} \sup_\theta \mathbb{E}_{f_\theta}\left[\|\theta-\hat{\theta}\|_2^2\right]$$

$$\geq \frac{L^2}{24\pi e\, d^2}\left(1 - \int_\Theta \mathbb{P}_{f_0}\left\{\|\theta-\hat{\theta}\|_2^2 \leq \rho d\right\} d\nu(\theta)\right)$$

Fano
$$\geq \frac{L^2}{24\pi e d^2}\left(1 - \frac{1}{3} - \frac{(n+d)L^2}{8 d^3 \sigma^2 \ln 2}\right)$$

We carefully pick $d \in \{2, \dots, n\}$ as a function of n, L, σ^2 to conclude.