Concentration:
Case Study: Optimal Discovery

Master 2 Mathematics and Computer Science

Guillaume Aubrun, Aurélien Garivier, Rémi Gribonval

2020-2021

UNIVERSITÉ DE LYON | ENS DE LYON

Table of contents

1. Discovering dangerous contigencies in electrical systems
2. Estimating the Unseen
3. The Good-UCB Algorithm
4. Optimality results

Discovering dangerous contigencies in electrical systems

The problem

Power system security assessment

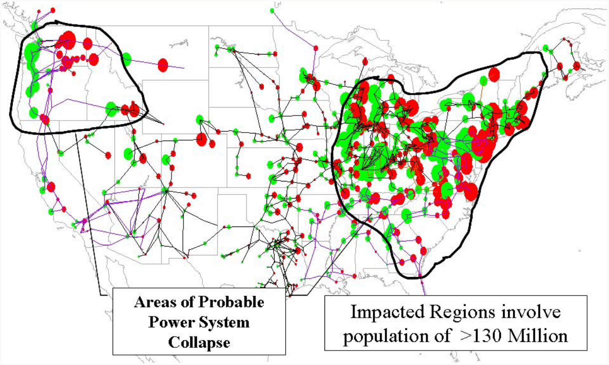

By Mark MacAlester, Federal Emergency Management Agency [Public domain], via Wikimedia Commons

Damien Ernst (Electrical Engineering, Liège): How to identify quickly contingencies/scenarios that could lead to unacceptable operating conditions (dangerous contingencies) if no preventive actions were taken?

The model

- Subset $A \subset \mathcal{X}$ of important items
- $|\mathcal{X}| \gg 1$, $|A| \ll |\mathcal{X}|$
- Access to $\mathcal{X}$ only by probabilistic experts $(P_i)_{1 \leq i \leq K}$: sequential independent draws

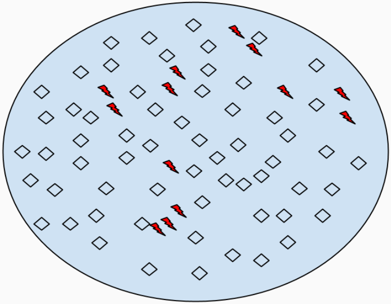

Goal: discover rapidly the elements of A

The model

- Subset $A \subset \mathcal{X}$ of important items
- $|\mathcal{X}| \gg 1$, $|A| \ll |\mathcal{X}|$
- Access to $\mathcal{X}$ only by probabilistic experts $(P_i)_{1 \le i \le K}$: sequential independent draws

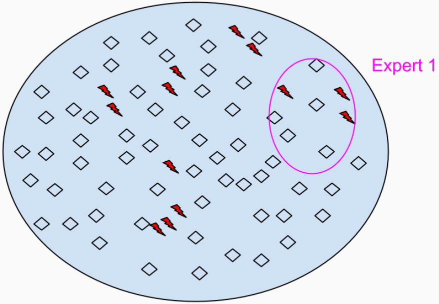

Goal: discover rapidly the elements of A

The model

- Subset $A \subset \mathcal{X}$ of important items
- $|\mathcal{X}| \gg 1$, $|A| \ll |\mathcal{X}|$
- Access to $\mathcal{X}$ only by probabilistic experts $(P_i)_{1 \leq i \leq K}$: sequential independent draws

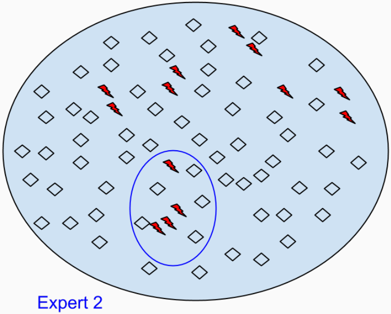

Goal: discover rapidly the elements of A

The model

- Subset $A \subset \mathcal{X}$ of important items
- $|\mathcal{X}| \gg 1$, $|A| \ll |\mathcal{X}|$
- Access to $\mathcal{X}$ only by probabilistic experts $(P_i)_{1 \leq i \leq K}$: sequential independent draws

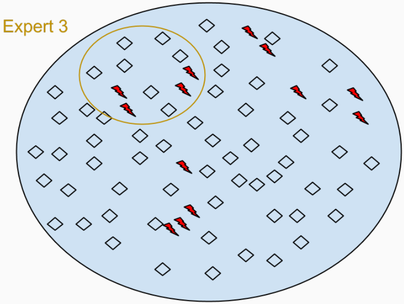

Goal: discover rapidly the elements of A

Estimating the Unseen

Enigma

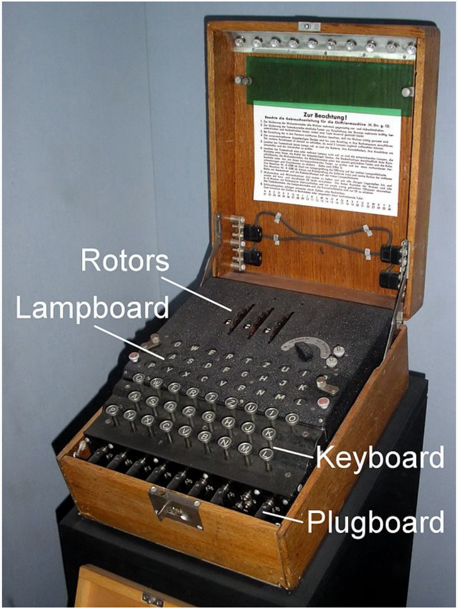

- Electro-mechanical rotor cipher machines, 26 characters
- Invented at the end of WW1 by Arthur Scherbius
- Commercial use, then German Army during WW2
- First cracked by Marian Rejewski in the 1930s (Bomb), then improved to 3.10^{114} configurations
- Read Simon Singh, *The Code Book*

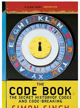

Enigma

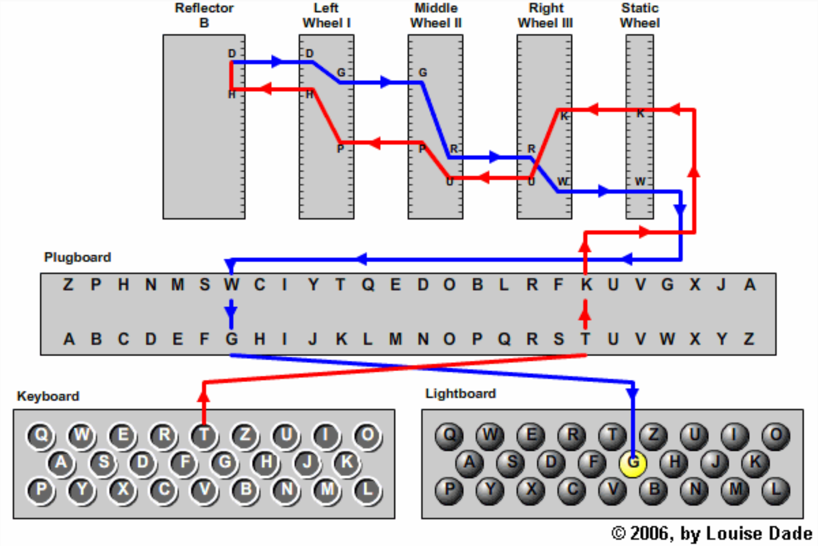

Src: http://enigma.louisedade.co.uk/

Battle of the Atlantic

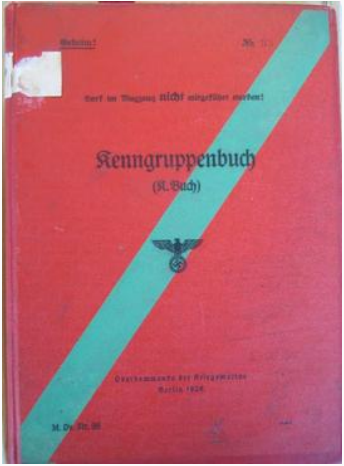

- Massively used by the German Kriegsmarine and Luftwaffe
- weakness: 3-letters setting to initiate communication, taken from the *Kenngruppenbuch*
- Government Code and Cypher School: Bletchley Park (on the train line between Cambridge and Oxford)
- Colossus (first programmable computers) in 1943

Estimating probabilities

- Discrete alphabet A.
- Unknown probability p on A
- Sample $X_1, \ldots, X_n$ of independent draws of p.
- Goal : use the sample to estimate $p(a)$ for all $a \in A$.

Natural idea:

$$\hat{p}(a) = \frac{N(a)}{n}, \quad \text{where } N(a) = \#\{i : X_i = a\}$$

Safari preparation

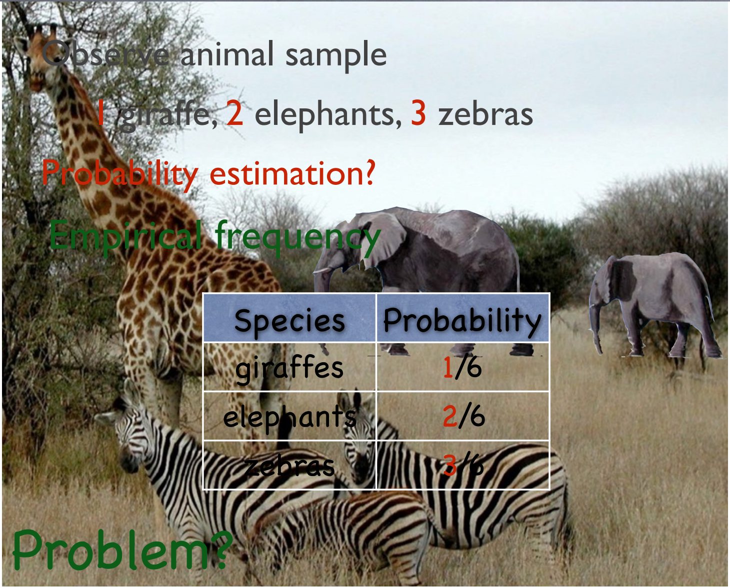

Observe animal sample

1 giraffe, 2 elephants, 3 zebras

Probability estimation?

Empirical frequency

Species	Probability
giraffes	1/6
elephants	2/6
zebras	3/6

Problem?

[Src: Alon Orlitsky, https://ic.epfl.ch/files/content/sites/ic/files/Inka/Orlitsky%20Talk%202016.pdf]

Bigram Model for NLP [Src: https://nlp.stanford.edu/~wcmac]

Learning set:

john read moby dick

mary read a different book

she read a book by cher

$$p(w_i|w_{i-1}) = \frac{c(w_{i-1}w_i)}{\sum_w c(w_{i-1}w)}$$

$$p(s) = \prod_{i=1}^{l+1} p(w_i|w_{i-1})$$

$$\begin{array}{cccccc}
p(& john & read & a & book &) \\
= & p(john|\cdot) & p(read|john) & p(a|read) & p(book|a) & p(\cdot|book) \\
= & \frac{c(\cdot\ john)}{\sum_w c(\cdot\ w)} & \frac{c(john\ read)}{\sum_w c(john\ w)} & \frac{c(reada)}{\sum_w c(read\ w)} & \frac{c(a\ book)}{\sum_w c(a\ w)} & \frac{c(book\ \cdot)}{\sum_w c(book\ w)} \\
= & \frac{1}{3} & \frac{1}{1} & \frac{2}{3} & \frac{1}{2} & \frac{1}{2} \\
\approx & 0.06 & & & &
\end{array}$$

Bigram Model for NLP

[Src: https://nlp.stanford.edu/~wcmac]

Learning set:

john read moby dick

mary read a different book

she read a book by cher

$$p(w_i|w_{i-1}) = \frac{c(w_{i-1}w_i)}{\sum_w c(w_{i-1}w)}$$

$$p(s) = \prod_{i=1}^{l+1} p(w_i|w_{i-1})$$

$$\begin{array}{cccccc}
p(& cher & read & a & book &) \\
= & p(cher|\cdot) & p(read|cher) & p(a|read) & p(book|a) & P(\cdot|book) \\
= & \frac{c(\cdot\ cher)}{\sum_w c(\cdot\ w)} & \frac{c(cher\ read)}{\sum_w c(cher\ w)} & \frac{c(reada)}{\sum_w c(read\ w)} & \frac{c(a\ book)}{\sum_w c(a\ w)} & \frac{c(book\ \cdot)}{\sum_w c(book\ w)} \\
= & \frac{0}{3} & \frac{0}{1} & \frac{2}{3} & \frac{1}{2} & \frac{1}{2} \\
= & 0 & & & &
\end{array}$$

$\Longrightarrow$ useless, the unseen **must** be treated correctly.

Bayesian Approach: Laplace Estimator

Pierre-Simon de Laplace (1749-1827), Thomas Bayes (1702-1761)

Will the sun rise tomorrow?

$$\hat{p}(a) = \frac{N(a) + 1}{n + |A|}$$

- good for small alphabets and many samples
- very bad when lots of items seen once (ex: DNA sequences)
- $|A|$ can be very large (or even infinite), but P concentrated on few items

$\Longrightarrow$ not a satisfying solution to the problem

Alan Turing Irving John Good

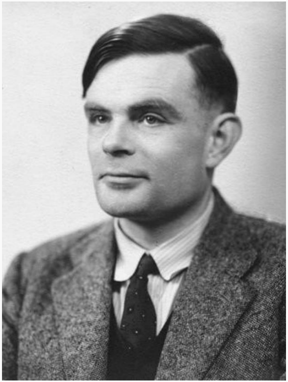

1912-1954
student of Godfrey Harold Hardy in Cambridge
PhD from Princeton with Alonzo Church

1916-2009
Graduated in Cambridge
Academic carrer in Bayesian statistics in Manchester and then in the University of Virginia (USA)

Missing mass estimation

$X_1, \ldots, X_n$ independent draws of $p \in \mathfrak{M}_1(A)$.

$$O_n(x) = \sum_{m=1}^{n} \mathbb{1}\{X_m = x\}$$

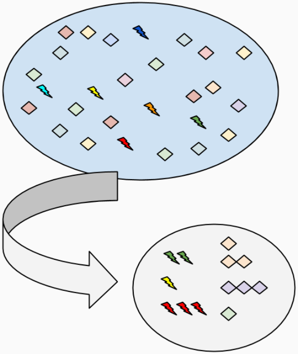

How to 'estimate' the total mass of the *unseen* items

$$M_n = \sum_{x \in A} p(x)\, \mathbb{1}\{O_n(x) = 0\}\ ?$$

Missing Mass

Let $A = \mathbb{N}$, let $p \in \mathcal{M}_1(\mathbb{N})$ and let $X_1, \dots, X_n \overset{iid}{\sim} p$.

For every $x \in \mathbb{N}$, let $O_n(x) = \sum_{i=1}^n \mathbb{1}\{X_i = x\}$.

Pb: estimate the mass of the unseen

$$M_n = \mathbb{P}\big(X_{n+1} \notin \{X_1, \dots, X_n\}\big) = \sum_{x=0}^{\infty} p(x)\, \mathbb{1}\big\{O_n(x) = 0\big\}$$

Idea: use *hapaxes* = symbols $x \in \mathbb{N}$ that appear once in the sample

$$\hat{M}_n = \frac{1}{n} \sum_{x=0}^{\infty} \mathbb{1}\big\{O_n(x) = 1\big\}$$

= Good-Turing 'estimator'

= *leave-one-out* estimator of M_n: if $X_{-i} = \{X_1, \dots, X_{i-1}, X_{i+1}, \dots, X_n\}$,

$$\hat{M}_n = \frac{1}{n} \sum_{i=1}^{n} \mathbb{1}\big\{X_i \notin X_{-i}\big\}$$

'Bias' of the Good-Turing estimator

Proposition [Good '1953]

Whatever the law p,

$$0 \leq \mathbb{E}\left[\hat{M}_n\right] - \mathbb{E}[M_n] \leq \frac{1}{n}$$

Proof:

$$\begin{aligned}
\mathbb{E}\left[\hat{M}_n\right] - \mathbb{E}[M_n] &= \frac{1}{n}\mathbb{E}\left[\sum_{x\in\mathbb{N}} \mathbb{1}\{O_n(x)=1\}\right] - \mathbb{E}\left[\sum_{x\in\mathbb{N}} p(X)\mathbb{1}\{O_n(x)=0\}\right] \\
&= \frac{1}{n}\sum_{x\in\mathbb{N}} \mathbb{P}\big(O_n(x)=1\big) - np(x)\,\mathbb{P}(O_n(x)=0) \\
&= \frac{1}{n}\sum_{x\in\mathbb{N}} np(x)\big(1-p(x)\big)^{n-1} - np(x)\big(1-p(x)\big)^{n} \\
&= \frac{1}{n}\sum_{x\in\mathbb{N}} p(x) \times np(x)\big(1-p(x)\big)^{n-1} \\
&= \frac{1}{n}\sum_{x\in\mathbb{N}} p(x)\,\mathbb{P}\big(O_n(x)=1\big) \\
&= \frac{1}{n}\mathbb{E}\left[\sum_{x\in\mathbb{N}} p(x)\mathbb{1}\big(O_n(x)=1\big)\right] \in \left[0, \frac{1}{n}\right]
\end{aligned}$$

Concentration of $\hat{M}_n$

$\hat{M}_n = \frac{1}{n}\sum_{x=0}^{\infty} \mathbb{1}\{O_n(x) = 1\} = \phi(X_1, \ldots, X_n)$, where
$\forall k, \forall x_1, \ldots, x_n, x'_k \in \mathbb{N}$,

$$\left|\phi(x_1, \ldots, x_n) - \phi(x_1, \ldots, x_{k-1}, x'_k, x_{k+1}, \ldots, x_n)\right| \le \frac{2}{n} \; .$$

Hence, by McDiarmid's inequality,

$$\mathbb{P}\Big(\big|\hat{M}_n - \mathbb{E}[\hat{M}_n]\big| > x\Big) \le \exp\left(-\frac{n\,x^2}{2}\right)$$

and with probability at least $1-\delta$,

$$\hat{M}_n \in \left[\mathbb{E}[\hat{M}_n] \pm \sqrt{\frac{2\log(1/\delta)}{n}}\right]$$

Concentration of the missing mass

$M_n = \sum_{x=0}^{\infty} p(x)\,\mathbb{1}\big\{O_n(x) = 0\big\}$ is a sum of *dependent* random variables.

But the $\mathbb{1}\big\{O_n(x) = 0\big\}$ are negatively associated!

Indeed,

- By the 0-1 principle, for all $1 \le i \le n$ the $\big\{\mathbb{1}\{X_i = x\} : x \in \mathbb{N}\big\}$ are NA
- Hence, by the union property and by the fact that the X_i are independent, the $\big\{\mathbb{1}\{X_i = x\} : 1 \le i \le n, x \in \mathbb{N}\big\}$ are NA
- Hence, by the concordant monotone property for the monotonically increasing function $(u_1, \dots, u_n) \mapsto u_1 + \dots + u_n$, the $\big\{O_n(x) = \sum_{i=1}^{n} \mathbb{1}\{X_i = x\} : x \in \mathbb{N}\big\}$ are NA
- Hence, again by the concordant monotone property for the monotonically decreasing function $u \mapsto \mathbb{1}\{u = 0\}$ on $\mathbb{N}$, the $\big\{\mathbb{1}\big\{O_n(x) = 0\big\} : x \in \mathbb{N}\big\}$ are NA

$$\implies \mathbb{E}\Big[\exp\big(\lambda M_n\big)\Big] = \mathbb{E}\left[\prod_{x \in \mathcal{X}} \exp\big(\lambda p(x) \mathbb{1}\{O_n(x) = 0\}\big)\right] \le \mathbb{E}\Big[\exp\big(\lambda \tilde{M}_n\big)\Big]$$

where $\tilde{M}_n = \sum_{x \in \mathbb{N}} p(x) Z_x$ and where the $Z_x \sim \mathcal{B}\Big(q(x) := \mathbb{P}\big(O_n(x) = 0\big)\Big)$ are independent

Back to Chernoff's roots

Attempts with Hoeffding, Bernstein, McDiarmid, etc. fail without an assumption of $\max_{x\in\mathbb{N}} P(x)$. In what follows we just use that M_n is real-valued

For every $x > \mathbb{E}[M_n]$ and every $\lambda > 0$,

$$\mathbb{P}(M_n \geq x) \leq \int_{u=x}^{\infty} \frac{e^{\lambda u}}{e^{\lambda x}} dP_{M_n}(u) \leq e^{-\lambda x} \int_{u=0}^{\infty} e^{\lambda u} dP_{M_n}(u) = \exp\big(-(\lambda x - \Lambda(\lambda)\big)$$

where $\Lambda(\lambda) = \log\big(Z(\lambda) := \int_{u=0}^{\infty} e^{\lambda u} dP_{M_n}(u)\big)$, and hence

$$\mathbb{P}(M_n \geq x) \leq \exp\big(-I(x)\big)$$

where $I(x) = \sup_{\lambda>0} \lambda x - \Lambda(\lambda)$.

Similarly, for every $x < \mathbb{E}[M_n]$,

$$\mathbb{P}(M_n \leq x) \leq \exp\big(-I(x)\big)$$

where $I(x) = \sup_{\lambda<0} \lambda x - \Lambda(\lambda)$.

Chernoff's rate function and KL divergence

Let $P = P_{M_n}$ and for $\lambda \in \mathbb{R}$ let P_λ be defined by $\frac{dP_\lambda}{dP}(x) = \frac{e^{\lambda x}}{Z(\lambda)}$, ie for all measurable, non-negative function f: $\mathbb{E}_\lambda\big[f(X)\big] = \int_{\mathbb{R}} f(x)\frac{e^{\lambda x}}{Z(\lambda)}dP(x)$

Prop: $\mathrm{KL}(P_\lambda, P) = \lambda\mathbb{E}_\lambda[X] - \Lambda(\lambda) = \inf\big\{\mathrm{KL}(Q,P) : \mathbb{E}_Q[X] \geq \mathbb{E}_\lambda[X]\big\}$

Proof: For every $Q \ll P$ with $\mathbb{E}_Q[X] \geq x$,

$$\begin{aligned}
\mathrm{KL}(Q,P) &= \int_{\mathbb{R}} \log\left(\frac{dQ}{dP}(x)\right) dQ(x)\\
&= \int_{\mathbb{R}} \log\left(\frac{dQ}{dP_\lambda}(x)\frac{dP_\lambda}{dP}(x)\right) dQ(x)\\
&= \mathrm{KL}(Q,P_\lambda) + \int_{\mathbb{R}} \log\left(\frac{e^{\lambda x}}{Z(\lambda)}\right) dQ(x)\\
&= \mathrm{KL}(Q,P_\lambda) + \lambda\mathbb{E}_Q[X] - \log\big(Z(\lambda)\big)\\
&\geq \qquad 0 \qquad + \lambda\mathbb{E}_\lambda[X] - \quad \Lambda(\lambda) \qquad = \mathrm{KL}(P_\lambda, P)
\end{aligned}$$

Cor: if $\lambda(x)$ is such that $\mathbb{E}_{\lambda(x)}[X] = x$, then $I(x) = \mathrm{KL}(P_{\lambda(x)}, P)$

Chernoff's rate function and KL divergence

Let $P = P_{M_n}$ and for $\lambda \in \mathbb{R}$ let P_λ be defined by $\frac{dP_\lambda}{dP}(x) = \frac{e^{\lambda x}}{Z(\lambda)}$, ie for all measurable, non-negative function f: $\mathbb{E}_\lambda\big[f(X)\big] = \int_{\mathbb{R}} f(x) \frac{e^{\lambda x}}{Z(\lambda)} dP(x)$

Prop: $\mathrm{KL}(P_\lambda, P) = \lambda \mathbb{E}_\lambda[X] - \Lambda(\lambda) = \inf\big\{\mathrm{KL}(Q, P) : \mathbb{E}_Q[X] \geq \mathbb{E}_\lambda[X]\big\}$

Cor: if $\lambda(x)$ is such that $\mathbb{E}_{\lambda(x)}[X] = x$, then $I(x) = \mathrm{KL}(P_{\lambda(x)}, P)$

Since $\Lambda'(\lambda) = \dfrac{\mathbb{E}\left[Xe^{\lambda X}\right]}{\mathbb{E}\left[e^{\lambda X}\right]} = \mathbb{E}_\lambda[X]$ and

$$\Lambda''(\lambda) = \frac{\mathbb{E}\left[X^2 e^{\lambda X}\right]}{\mathbb{E}\left[e^{\lambda X}\right]} - \left(\frac{\mathbb{E}\left[X e^{\lambda X}\right]}{\mathbb{E}\left[e^{\lambda X}\right]}\right)^2 = \mathbb{V}\mathrm{ar}_\lambda[X] > 0,$$ the C^∞ mapping $\lambda \mapsto \lambda x - \Lambda(\lambda)$ is maximal where at $\lambda(x)$ where $x = \Lambda'\big(\lambda(x)\big) = \mathbb{E}_{\lambda(x)}[X]$ and then

$$\begin{aligned} I(x) &= \lambda(x)x - \Lambda\big(\lambda(x)\big) \\ &= \lambda(x)x - \Big(\lambda(x)\mathbb{E}_{\lambda(x)}[X] - \mathrm{KL}\big(P_{\lambda(x)}, P\big)\Big) \\ &= \mathrm{KL}\big(P_{\lambda(x)}, P\big) \end{aligned}$$

Kullback-Leibler divergence and variance

$$\mathrm{KL}\big(P_{\lambda(x)}, P\big) = \int_{\mathbb{E}[X]}^{x} \int_{\mathbb{E}[X]}^{t} \frac{1}{\mathbb{V}\mathrm{ar}_{\lambda(u)}[X]} du$$

Proof: If $g(x) = \mathrm{KL}\big(P_{\lambda(x)}, P\big) = \lambda(x)x - \Lambda(\lambda(x))$ then

$$g'(x) = \lambda'(x)x + \lambda(x) - \lambda'(x)\Lambda'\big(\lambda(x)\big) = \lambda(x)$$

and if $e(\ell) = \lambda^{-1}(\ell) = \mathbb{E}_\ell[X] = \Lambda'(\ell)$

$$g''(x) = \lambda'(x) = \frac{1}{e'\big(\lambda(x)\big)} = \frac{1}{\Lambda''\big(\lambda(x)\big)} = \frac{1}{\mathbb{V}\mathrm{ar}_{\lambda(x)}[X]}$$

The result follows since $g\big(\mathbb{E}[X]\big) = 0$ and $g'\big(\mathbb{E}[X]\big) = \lambda\big(\mathbb{E}[X]\big) = 0$.

Cor: if $\forall u \in \big[\mathbb{E}[X], x\big], \mathbb{V}\mathrm{ar}_{\lambda(u)}[X] \leq \sigma^2$ **then** $I\big(\mathbb{E}[X] + \epsilon\big) \geq \dfrac{\epsilon^2}{2\sigma^2}$

Similarly, if $\forall u \in [1, x], \mathbb{V}\mathrm{ar}_{\lambda(u)}[X] \leq u$ as for $\mathcal{P}(1)$ then $\forall x \geq 0$

$$I(1+x) \geq \int_1^{1+x} \int_1^t \frac{du}{u} = (1+x)\log(1+x) - x$$

For the missing mass

$\tilde{M}_n = \sum_{x \in X} p(x) Z_x$ where the $Z_x \sim \mathcal{B}\Big(q(x) := \mathbb{P}\big(O_n(x) = 0\big)\Big)$ are independent. Under P_λ, the $Z_x \overset{iid}{\sim} \mathcal{B}\left(q_\lambda(x) = \frac{q(x)e^{\lambda p(x)}}{1-q(x)+q(x)e^{\lambda p(x)}}\right)$

$$\mathbb{V}\mathrm{ar}_\lambda\big[\tilde{M}_n\big] = \sum_{x \in \mathbb{N}} p(x)^2 q_\lambda(x)\big(1 - q_\lambda(x)\big) \leq \sum_{x \in \mathbb{N}} p(x)^2 q_\lambda(x)$$

Hence, for $\lambda < 0$, $\mathbb{V}\mathrm{ar}_\lambda\big[\tilde{M}_n\big] \leq \sum_{x \in \mathbb{N}} p(x)^2 q(x)$ and since

$$p(x)q(x) \leq p(x) \exp\big(-np(x)\big) \leq \frac{1}{n} \sup_{u>0} \Big\{u\, e^{-u}\Big\} = \frac{1}{en}\ ,$$

$$\mathbb{V}\mathrm{ar}_\lambda\big[\tilde{M}_n\big] \leq \sum_{x \in \mathbb{N}} \frac{p(x)}{en} \leq \frac{1}{en}$$

which yields

For all $\epsilon > 0$**,** $I\big(\mathbb{E}[M_n] - \epsilon\big) \geq \frac{e\, n\, \epsilon^2}{2}$

Hence, with probability at least $1 - \delta$, $M_n \geq \mathbb{E}[M_n] - \sqrt{\frac{2\log(1/\delta)}{en}}$.

High confidence estimation of the missing mass

A similar bound can be obtained for the right-deviations of M_n. Putting everything together,

High confidence region

With probability at least $1-\delta$, whatever the law p,

$$\hat{M}_n - \frac{1}{n} - (1+\sqrt{2})\sqrt{\frac{\log(4/\delta)}{n}} \leq M_n \leq \hat{M}_n + (1+\sqrt{2})\sqrt{\frac{\log(4/\delta)}{n}}$$

$\Longrightarrow$ sub-Gaussian concentration despite the absence of independence and the absence of assumptions on p.

The Good-UCB Algorithm

The model

- Subset $A \subset \mathcal{X}$ of important items
- $|\mathcal{X}| \gg 1$, $|A| \ll |\mathcal{X}|$
- Access to $\mathcal{X}$ only by probabilistic experts $(P_i)_{1 \leq i \leq K}$: sequential independent draws

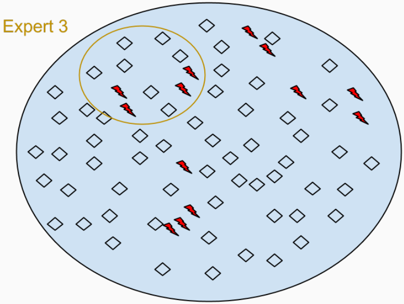

Goal: discover rapidly the elements of A

Goal

At each time step $t = 1, 2, \dots$:

- pick an index $I_t = \pi_t\big(I_1, Y_1, \dots, I_{s-1}, Y_{s-1}\big) \in \{1, \dots, K\}$ according to past observations
- observe $Y_t = X_{I_t, n_{I_t, t}} \sim P_{I_t}$, where

$$n_{i,t} = \sum_{s \le t} \mathbb{1}\{I_s = i\}$$

Goal: design the strategy $\pi = (\pi_t)_t$ so as to maximize the number of important items found after t requests

$$F^{\pi}(t) = \Big| A \cap \big\{ Y_1, \dots, Y_t \big\} \Big|$$

Assumption: non-intersecting supports

$$A \cap \mathrm{supp}(P_i) \cap \mathrm{supp}(P_j) = \emptyset \text{ for } i \neq j$$

Is it a Bandit Problem ?

It looks like a bandit problem. . .

- sequential choices among K options
- want to maximize cumulative rewards
- exploration vs exploitation dilemma

. . . but it is **not a bandit problem** !

- rewards are not i.i.d.
- destructive rewards: no interest to observe twice the same important item
- all strategies eventually equivalent

The oracle strategy

Proposition: Under the non-intersecting support hypothesis, the greedy oracle strategy

$$I_t^* \in \operatorname*{argmax}_{1 \leq i \leq K} P_i \left(A \setminus \{ Y_1, \ldots, Y_t \} \right)$$

is optimal: for every possible strategy π, $\mathbb{E}\big[F^\pi(t)\big] \leq \mathbb{E}\big[F^*(t)\big]$.

Remark: the proposition is false if the supports may intersect

$\Longrightarrow$ estimate the "*missing mass* of important items"!

Solution proposed in [Optimal Discovery with Probabilistic Expert Advice: Finite Time Analysis and Macroscopic Optimality, *by Sébastien Bubeck, Damien Ernst and Aurélien Garivier*, Journal of Machine Learning Research vol. 14 Feb. 2013, pp.601-623]

The Good-UCB algorithm

Estimator of the missing important mass for expert i:

$$\hat{R}_{i,n_{i,t-1}} = \frac{1}{n_{i,t-1}} \sum_{x \in A} \mathbb{1}\left\{ \sum_{s=1}^{n_{i,t-1}} \mathbb{1}\{X_{i,s} = x\} = 1 \right. \\ \left. \text{and } \sum_{j=1}^{K} \sum_{s=1}^{n_{j,t-1}} \mathbb{1}\{X_{j,s} = x\} = 1 \right\}$$

Good-UCB algorithm:

1: For $1 \leq t \leq K$ choose $I_t = t$.
2: **for** $t \geq K+1$ **do**
3: Choose $I_t = \text{argmax}_{1 \leq i \leq K} \left\{ \hat{R}_{i,n_{i,t-1}} + C\sqrt{\frac{\log(4t)}{n_{i,t-1}}} \right\}$
4: Observe Y_t distributed as P_{I_t}
5: Update the missing mass estimates accordingly
6: **end for**

Optimality results

Classical analysis

Theorem: For any $t \geq 1$, *under the non-intersecting support assumption*, Good-UCB (with constant $C = (1 + \sqrt{2})\sqrt{3}$) satisfies

$$\mathbb{E}\left[F^*(t) - F^{UCB}(t)\right] \leq 17\sqrt{Kt\log(t)} + 20\sqrt{Kt} + K + K\log(t/K)$$

Remark: Usual result for bandit problem, but not-so-simple analysis

Sketch of proof

1. On a set $\tilde{\Omega}$ of probability at least $1-\sqrt{\frac{K}{t}}$, the "confidence intervals" hold true simultaneously all $u \geq \sqrt{Kt}$
2. Let $\bar{I}_u = \operatorname{argmax}_{1\leq i\leq K} R_{i,n_{i,u-1}}$. On $\tilde{\Omega}$,

$$R_{I_u,n_{I_u,u-1}} \geq R_{\bar{I}_u,n_{\bar{I}_u,u-1}} - \frac{1}{n_{I_u,u-1}} - 2(1+\sqrt{2})\sqrt{\frac{3\log(4u)}{n_{I_u,u-1}}}$$

3. But one shows that $\mathbb{E}F^*(t) \leq \sum_{u=1}^t \mathbb{E}R_{\bar{I}_u,n^\pi_{\bar{I}_u,u-1}}$
4. Thus

$$\begin{aligned}
&\mathbb{E}\left[F^*(t) - F^{UCB}(t)\right] \\
&\leq \sqrt{Kt} + \mathbb{E}\left[\sum_{u=1}^t \frac{1}{n_{I_u,u-1}} + 2(1+\sqrt{2})\sqrt{\frac{3\log(4t)}{n_{I_u,u-1}}}\right] \\
&\leq \sqrt{Kt} + K + K\log(t/K) + 4(1+\sqrt{2})\sqrt{3Kt\log(4t)}
\end{aligned}$$

Experiment: restoring property

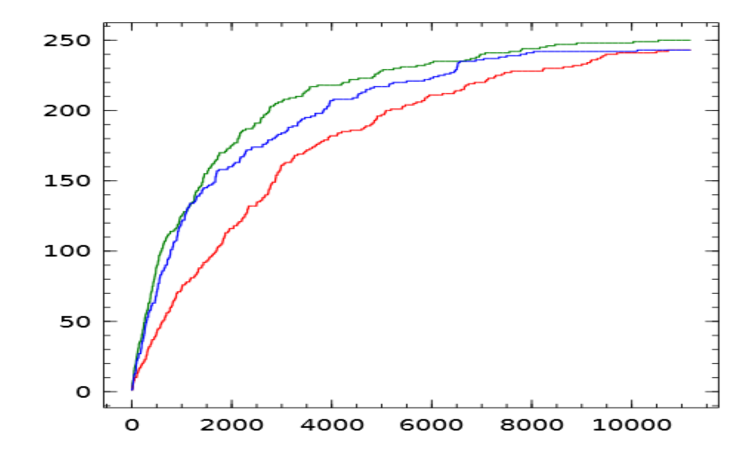

Figure 1: green: oracle, blue: Good-UCB, red: uniform sampling

Another analysis of Good-UCB

For $\lambda \in (0,1)$, $T(\lambda)$ = time at which missing mass of important items is smaller than λ on all experts:

$$T(\lambda) = \inf \left\{ t : \forall i \in \{1, \ldots, K\}, P_i\big(A \setminus \{Y_1, \ldots, Y_t\}\big) \le \lambda \right\}$$

Theorem: Let $c > 0$ and $S \ge 1$. Under the non-intersecting support assumption, for Good-UCB with $C = (1+\sqrt{2})\sqrt{c+2}$, with probability at least $1 - \frac{K}{cS^c}$, for any $\lambda \in (0,1)$,

$$T_{UCB}(\lambda) \le T^* + KS \log\left(8T^* + 16KS\log(KS)\right),$$

$$\text{where} \quad T^* = T^*\left(\lambda - \frac{3}{S} - 2(1+\sqrt{2})\sqrt{\frac{c+2}{S}}\right)$$

The macroscopic limit

- Restricted framework: $P_i = \mathcal{U}\{1, \ldots, N\}$
- $N \to \infty$
- $|A \cap \mathrm{supp}(P_i)|/N \to q_i \in (0,1)$, $q = \sum_i q_i$

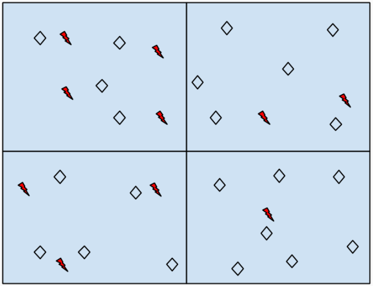

The macroscopic limit

- Restricted framework: $P_i = \mathcal{U}\{1, \dots, N\}$
- $N \to \infty$
- $|A \cap \mathrm{supp}(P_i)|/N \to q_i \in (0, 1)$, $q = \sum_i q_i$

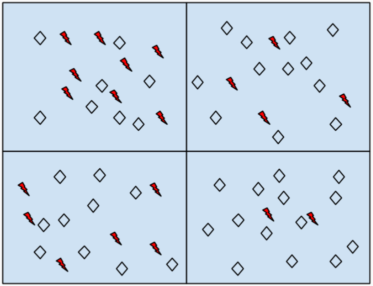

The macroscopic limit

- Restricted framework: $P_i = \mathcal{U}\{1, \ldots, N\}$
- $N \to \infty$
- $|A \cap \mathrm{supp}(P_i)|/N \to q_i \in (0, 1)$, $q = \sum_i q_i$

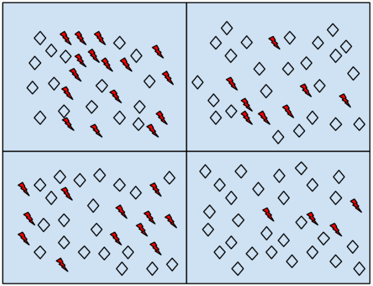

The Oracle behaviour

The limiting discovery process of the Oracle strategy is *deterministic*

Proposition: For every $\lambda \in (0, q_1)$, for every sequence $(\lambda^N)_N$ converging to λ as N goes to infinity, almost surely

$$\lim_{N\to\infty} \frac{T_*^N(\lambda^N)}{N} = \sum_i \left(\log \frac{q_i}{\lambda}\right)_+$$

Oracle vs. uniform sampling

Oracle: The proportion of important items not found after Nt draws tends to

$$q - F^*(t) = I(t)\underline{q}_{I(t)} \exp\left(-t/I(t)\right) \leq K\underline{q}_K \exp(-t/K)$$

with $\underline{q}_K = \left(\prod_{i=1}^K q_i\right)^{1/K}$ the geometric mean of the $(q_i)_i$.

Uniform: The proportion of important items not found after Nt draws tends to $K\bar{q}_K \exp(-t/K)$

$\implies$ Asymptotic ratio of efficiency

$$\rho(q) = \frac{\bar{q}_K}{\underline{q}_K} = \frac{\frac{1}{K}\sum_{i=1}^k q_i}{\left(\prod_{i=1}^k q_i\right)^{1/K}} \geq 1$$

larger if the $(q_i)_i$ are unbalanced

Macroscopic optimality

Theorem: Take $C = (1+\sqrt{2})\sqrt{c+2}$ with $c > 3/2$ in the Good-UCB algorithm.

- For every sequence $(\lambda^N)_N$ converging to λ as N goes to infinity, almost surely

$$\limsup_{N \to +\infty} \frac{T^N_{UCB}(\lambda^N)}{N} \leq \sum_i \left(\log \frac{q_i}{\lambda}\right)_+$$

- The proportion of items found after Nt steps $F^{GUCB}(Nt)$ converges *uniformly* to $F^*(Nt)$ as N goes to infinity

Experiment

Number of items found by Good-UCB (solid), the OCL (dashed), and uniform sampling (dotted) as a function of time for sizes $N = 128$, $N = 500$, $N = 1000$ and $N = 10000$ in a 7-experts setting.

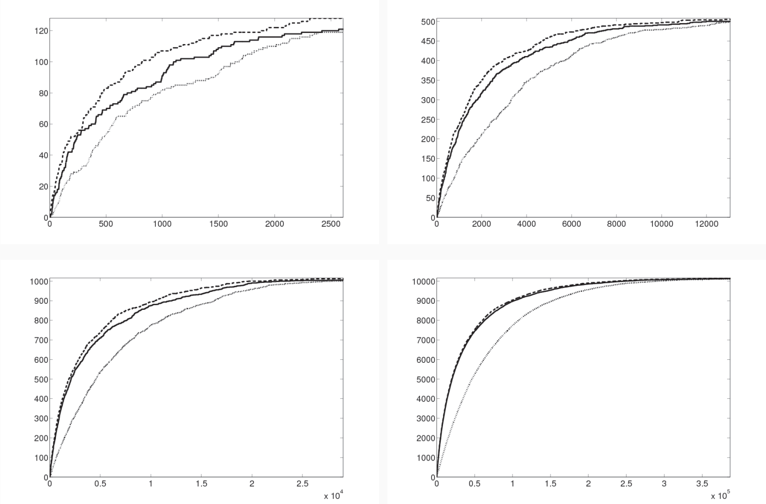

And when the assumptions are not satisfied?

Number of primes found by Good-UCB (solid), the oracle (dashed) and uniform sampling (dotted) using geometric experts with means 100, 300, 500, 700, 900, for $C = 0.1$ (top) and $C = 0.02$ (bottom).

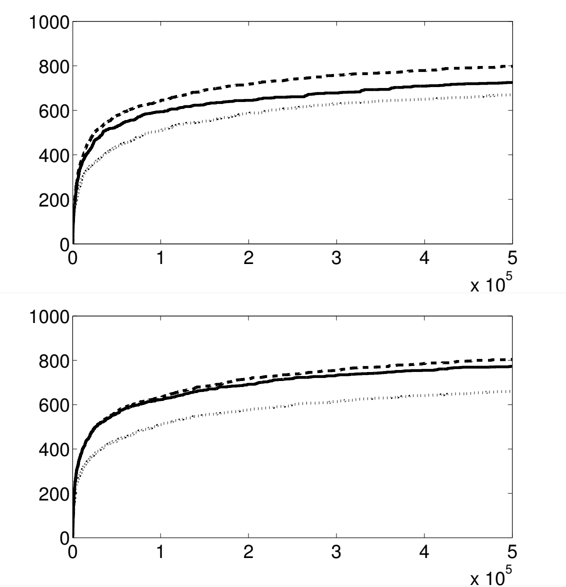

Conclusion and perspectives

- We propose an algorithm for the optimal discovery with probabilistic expert advice
- We give a standard regret analysis under the only assumption that the supports of the experts are non-overlapping
- We propose a different optimality result, which permits a macroscopic analysis in the uniform case
- Another interesting limit to consider is when the number of important items to find is fixed, but the total number of items tends to infinity (Poisson regime)
- Then, the behavior of the algorithm is not very good: too large confidence bonus because no tight deviations bounds for the Good-Turing estimator when the proportion of important items tends to 0. Improvement by better deviation bounds?