Mathematics of statistical decision making

and applications to bariatric surgery

Data science seminar M1-M2

Patrick Saux

supervised by **Odalric-Ambrym Maillard** and **Philippe Preux**

February 21, 2024

Towards recommender systems in healthcare

Online advertising

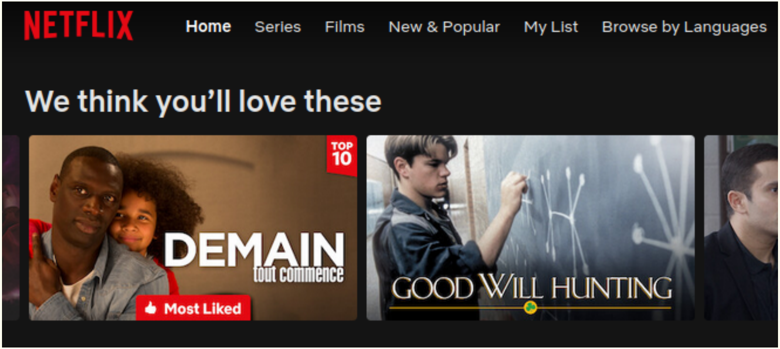

Online advertising

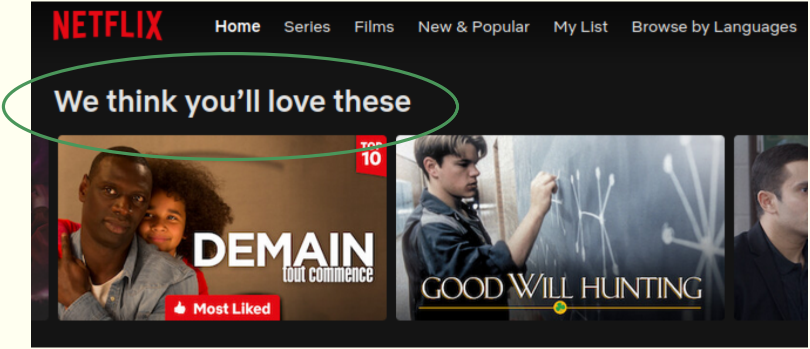

Online advertising

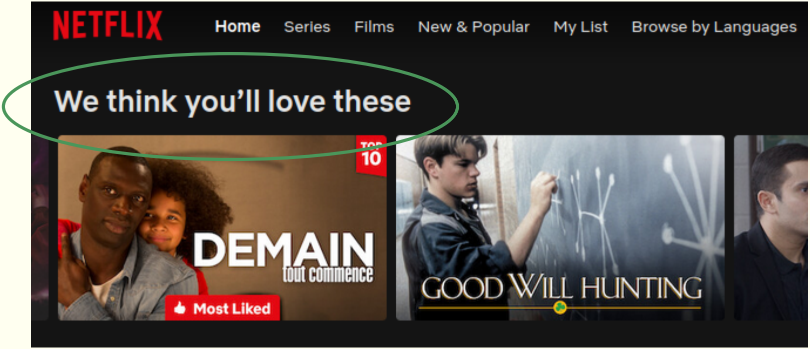

$\hookrightarrow$ from online marketing to healthcare recommendations.

Postoperative follow-up: where to go?

Patient at time t

Patient at time $t-1$

Patient at time $t+1$

Surgery department

General practitioner

Electronic consultation

Postoperative follow-up: when to go?

Patient at time t

Patient at time $t-1$

Patient at time $t+1$

Surgery department

follow-up visit after...

3 months?

6 months?

1 year?

Bariatric surgery

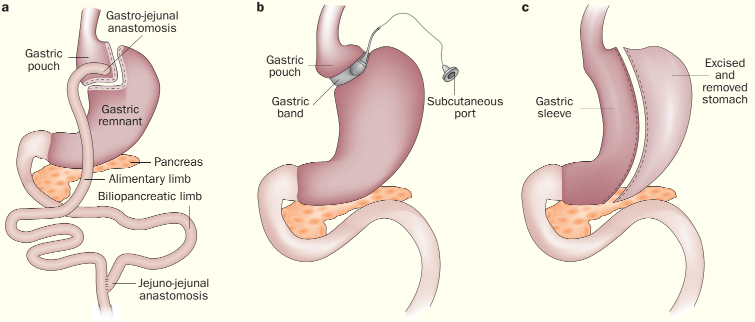

Roux-en-Y gastric bypass

Adjustable gastric band

Sleeve gastrectomy

$\hookrightarrow$ In France: 60 000 operations each year, 1% of adults.

Miras and Le Roux. Mechanisms underlying weight loss after bariatric surgery. Nature reviews, 2013.

Bariatric Weight Trajectory Prediction

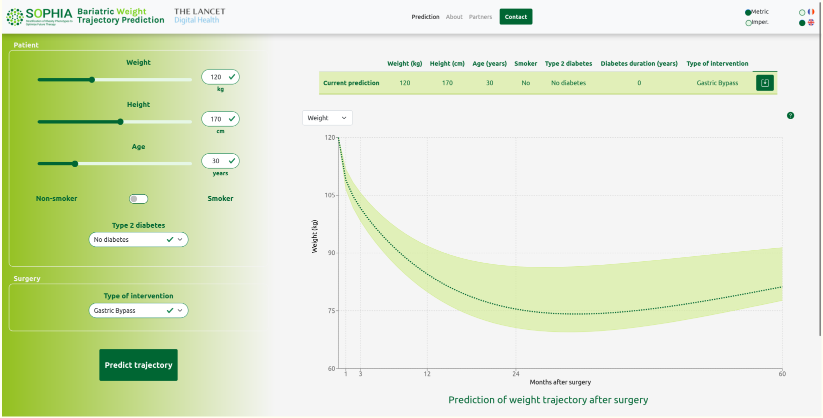

https://bwtp.univ-lille.fr/

PS, Bauvin et al., The Lancet Digital Health 2023.

Bariatric Weight Trajectory Prediction

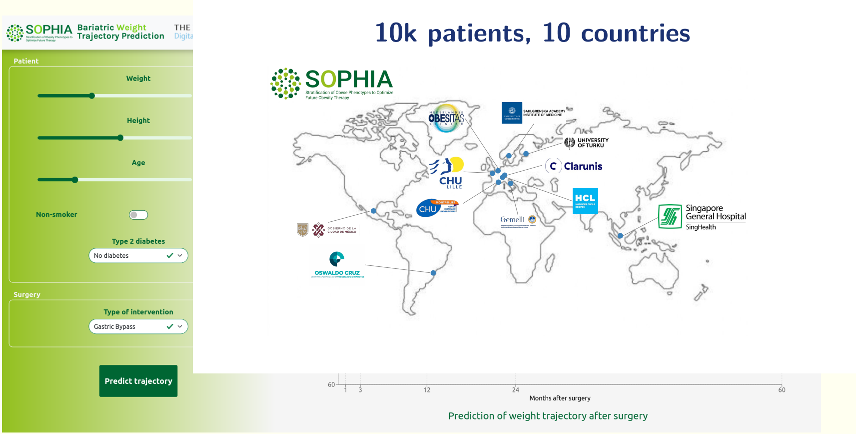

https://bwtp.univ-lille.fr/

PS, Bauvin et al., The Lancet Digital Health 2023.

Bariatric Weight Trajectory Prediction

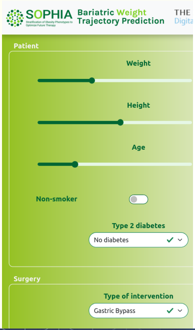

10k patients, 10 countries

obesityweek | THE OBESITY SOCIETY
November 1-4, 2022 • San Diego

IFSO NAPOLI 2023

https://bwtp.univ-lille.fr/

PS, Bauvin et al., The Lancet Digital Health 2023.

Bariatric Weight Trajectory Prediction

THE LANCET
Digital Health

10k patients, 10 countries

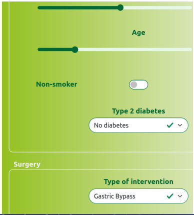

obesityweek | THE OBESITY SOCIETY
November 1-4, 2022 • San Diego

IFSO
NAPOLI
2023

https://bwtp.univ-lille.fr/

PS, Bauvin et al., The Lancet Digital Health 2023.

BWTP - training cohorts

ABOS

Prospective cohort
Retrospective study

- N=1147 (first intervention)
- > 500 baseline attributes

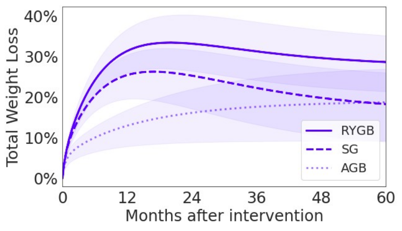

BAREVAL

- N=348
- SG only

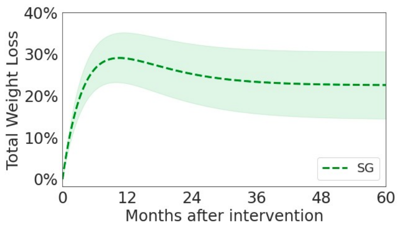

PS, Bauvin et al., The Lancet Digital Health 2023.

BWTP - prediction pipeline

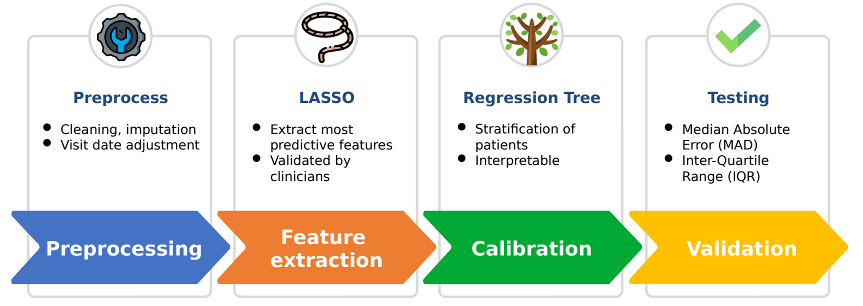

BWTP - prediction pipeline

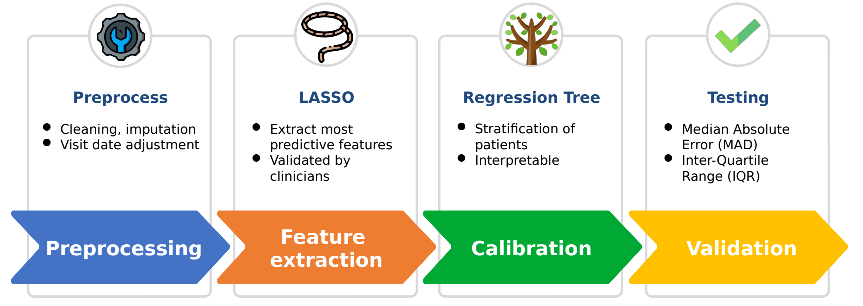

Interpretable regression trees (CART)

PS, Bauvin et al., The Lancet Digital Health 2023.

BWTP - prediction model

7 selected variables (among > 500)

- ✓ Weight
- ✓ Height
- ✓ Type of intervention
- ✓ Age
- ✓ Type II diabetes
- ✓ Diabetes duration
- ✓ Smoking

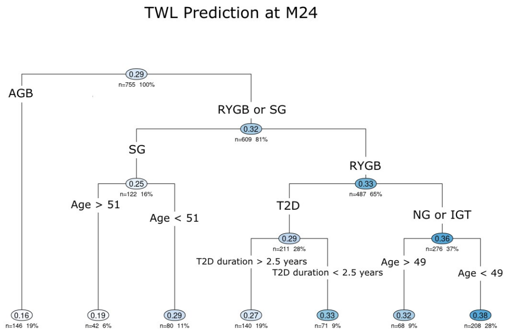

PS, Bauvin et al., The Lancet Digital Health 2023.

From preoperative to postoperative follow-up

Decision-making?

Preoperative

Bariatric Weight Trajectory Prediction

Postperative

Prospective tool to improve follow-up (complications, costs)

An abstract framework for decision-making

Stochastic bandit

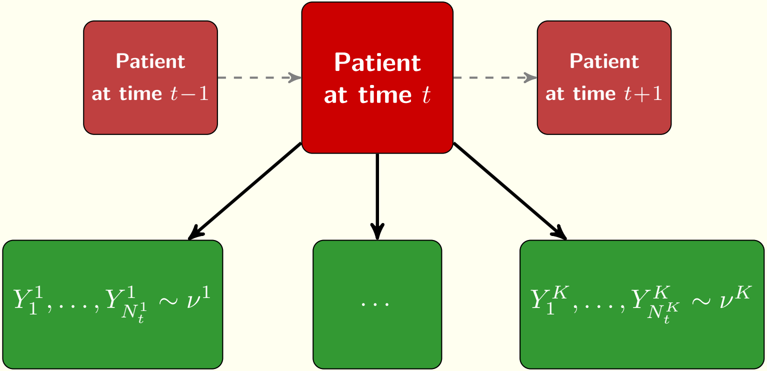

An abstract framework for decision-making

Stochastic bandit

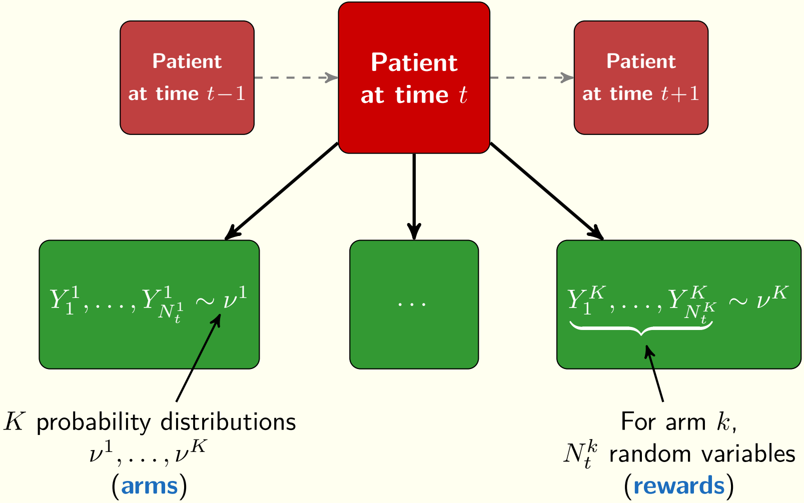

Stochastic bandits

$Y_1^1, \dots, Y_{N_t^1}^1 \sim \nu^1$ | $\dots$ | $Y_1^K, \dots, Y_{N_t^K}^K \sim \nu^K$

K probability distributions
$\nu^1, \dots, \nu^K$
(**arms**)

For arm k,
N_t^k random variables
(**rewards**)

- Active sampling of arm π_t.

Score (e.g. expectation)

- Minimise **regret** $\mathcal{R}_T = \sum_{t=1}^{T} \max_{k \in [K]} \rho(\nu^k) - \rho(\nu^{\pi_t})$.

- **Contextual bandits**: features X help predict rewards Y.

Lattimore and Szepesvári. Bandit algorithms. 2020.

Stochastic bandits

$$Y_1^1, \dots, Y_{N_t^1}^1 \sim \nu^1 \qquad \dots \qquad \underbrace{Y_1^K, \dots, Y_{N_t^K}^K} \sim \nu^K$$

K probability distributions
$\nu^1, \dots, \nu^K$
(**arms**)

For arm k,
N_t^k random variables
(**rewards**)

Regret lower bound ($T \to +\infty$):

$$\mathcal{R}_T = \Omega(\sqrt{T}) \quad \text{or} \quad \mathcal{R}_T = \Omega(\log T).$$

Worst case

Instance-dependent

Lattimore and Szepesvári. Bandit algorithms. 2020.

Exploration versus exploitation

$Y_1^1, \ldots, Y_{N_t^1}^1 \sim \nu^1$

$\ldots$

$\underbrace{Y_1^K, \ldots, Y_{N_t^K}^K} \sim \nu^K$

K probability distributions
$\nu^1, \ldots, \nu^K$
(**arms**)

For arm k,
N_t^k random variables
(**rewards**)

Greed is **not** good!

✖ $\pi_{t+1} \neq \underset{k \in [K]}{\operatorname{argmax}}\, \widehat{\rho}_t^k$.

Need to **explore**:

- ✔ Deterministic:
  $\widehat{\rho}_t^k$ + confidence set.
- ✔ Randomised:
  $\widehat{\rho}_t^k$ + noise.

Deterministic exploration, a.k.a. *optimism*

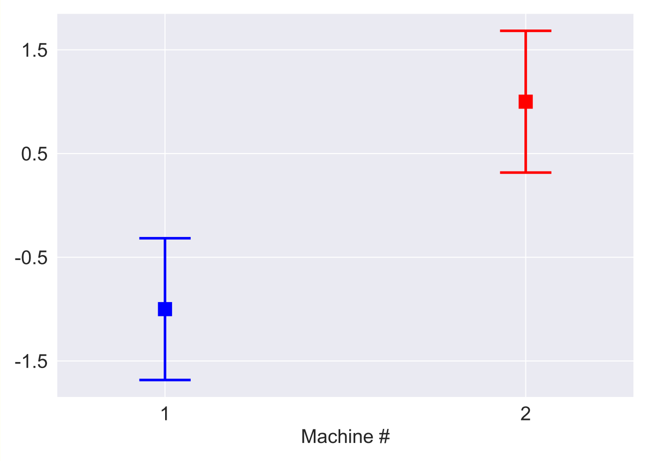

Hospital: 5 patients. GP: 5 patients.

Play the highest confidence bar.

Deterministic exploration, a.k.a. *optimism*

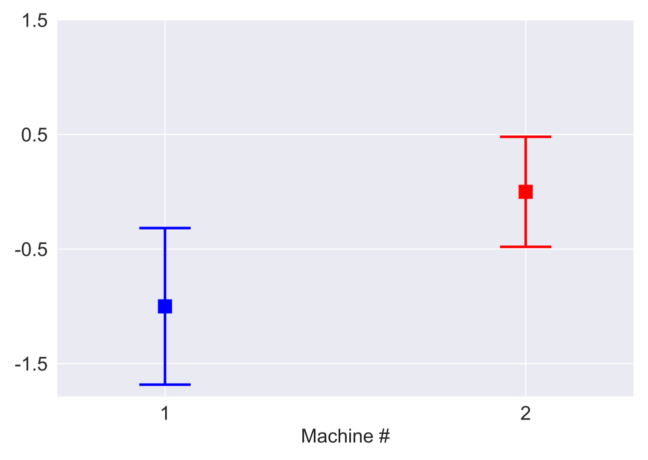

Hospital: 5 patients. GP: 10 patients.

Deterministic exploration, a.k.a. *optimism*

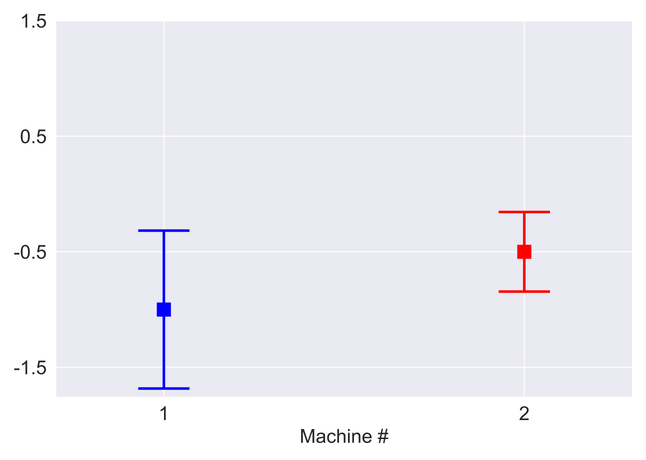

Hospital: 5 patients. GP: 20 patients.

Deterministic exploration, a.k.a. *optimism*

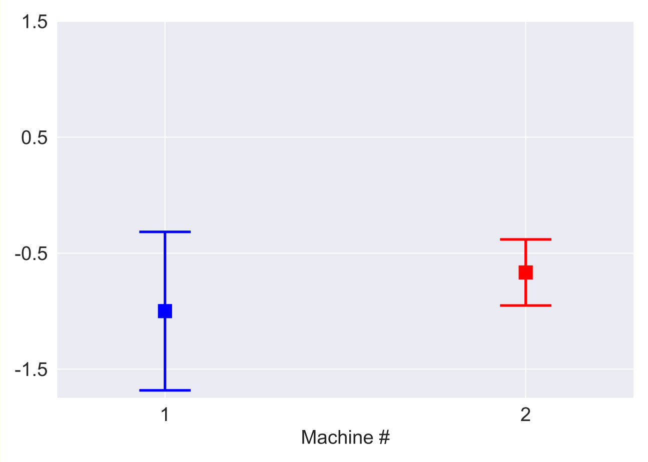

Hospital: 5 patients. GP: 30 patients.

Deterministic exploration, a.k.a. *optimism*

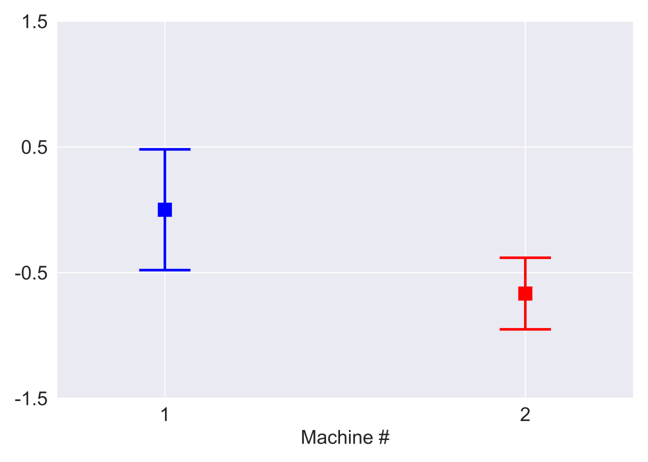

Hospital: 10 patients. GP: 30 patients.

Deterministic exploration, a.k.a. *optimism*

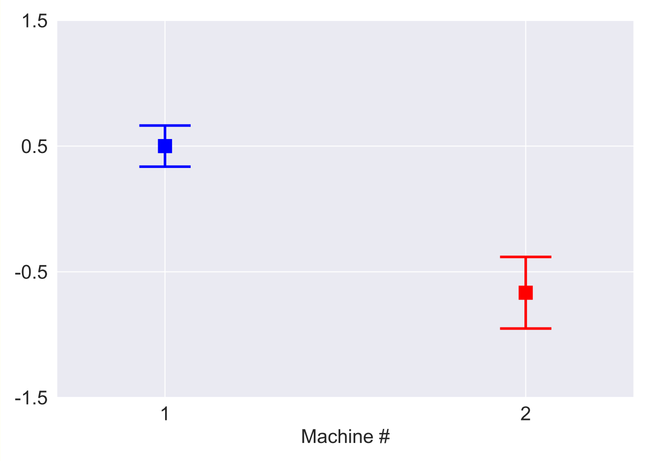

Hospital: 100 patients. GP: 30 patients.

☞ Correctly fixed the initial mistake!

Challenges of healthcare recommendations

Online advertising

- ✔ Huge volume of data.
- ✔ Easy to model (Bernoulli, logistic...).
- ✔ Limited risks ($\rho = \mathbb{E}$).

Healthcare

- ✖ Slow and scarce data.
- ✖ Nonparametric models.
- ✖ Risky.

Example distributions

Grain yield distributions (=rewards)

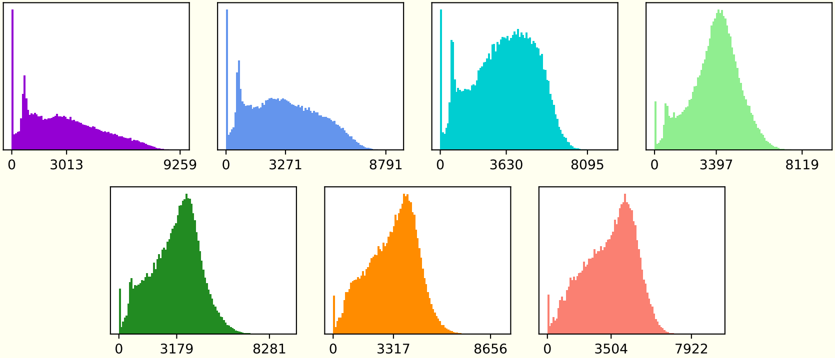

I. Efficient deterministic exploration

- Quantify uncertainty for the score ρ from i.i.d. $Y_1, \ldots, Y_N$.

- Confidence sets: $\forall n \in \mathbb{N},\ \mathbb{P}\left(\rho \in \widehat{\Theta}_n\right) \geqslant 1 - \delta$ (e.g. $\delta = 5\%$).

I. Efficient deterministic exploration

- Quantify uncertainty for the score ρ from i.i.d. $Y_1, \ldots, Y_N$.

- Confidence sets: $\forall n \in \mathbb{N},\ \mathbb{P}\left(\rho \in \widehat{\Theta}_n\right) \geqslant 1 - \delta$ (e.g. $\delta = 5\%$).
  - ✖ **p-hacking**: incompatible with random (active, data-dependent) N.

I. Efficient deterministic exploration

Quantify uncertainty for the score ρ from i.i.d. $Y_1, \ldots, Y_N$.

? Confidence sets: $\forall n \in \mathbb{N},\ \mathbb{P}\left(\rho \in \widehat{\Theta}_n\right) \geqslant 1 - \delta$ (e.g. $\delta = 5\%$).

- ✖ **p-hacking**: incompatible with random (active, data-dependent) N.

✔ **Anytime valid** confidence sequence:

$$\forall \text{ stopping time } N,\ \mathbb{P}\left(\rho \in \widehat{\Theta}_N\right) \geqslant 1 - \delta\,.$$

Ramdas, Grünwald, et al. Game-theoretic statistics and safe anytime-valid inference. Statistical Science, 2023.

I. Concentration bounds

Method	Assumption	Variance adaptive	Random N
t-test	Gaussian (or $N \to +\infty$)	✔	✖
Method of mixtures (MM) supermartingale inequality	Sub-Gaussian	✖	✔

Robbins and Pitman. Application of the method of mixtures to quadratic forms in normal variates. 1949.
Peña, Lai, and Shao. Self-normalized processes: Limit theory and statistical applications. 2009.

I. Concentration bounds

Method	Assumption	Variance adaptive	Random N
t-test	Gaussian (or $N \to +\infty$)	✔	✖
Method of mixtures (MM) supermartingale inequality	Sub-Gaussian	✖	✔
Bregman MM	Exp. families	✔	✔
Empirical Chernoff MM	Second order sub-Gaussian	✔	✔

Robbins and Pitman. Application of the method of mixtures to quadratic forms in normal variates. 1949.
Peña, Lai, and Shao. Self-normalized processes: Limit theory and statistical applications. 2009.

I. Bregman concentration

Parameter

- Exponential family: $p_\theta(y) = h(y) \exp\left(\langle \theta, F(y) \rangle - \mathcal{L}(\theta)\right)$.

Feature function Log-partition function

- Bregman divergence:

$$\mathcal{B}(\theta', \theta) = \mathcal{L}(\theta') - \mathcal{L}(\theta) - \langle \theta' - \theta, \nabla \mathcal{L}(\theta) \rangle = \mathrm{KL}(p_\theta \parallel p_{\theta'})\,.$$

Family	Known anytime valid concentration
Bernoulli, Gaussian (known variance)	✔
Gaussian, Chi-square, Poisson, Pareto, etc.	✖

Kaufmann and Koolen. Mixture martingales revisited. JMLR, 2021.

I. Bregman concentration

$$\widehat{\Theta}_n^\delta = \left\{\theta_0 \in \Theta : (n+c)\mathcal{B}_{\mathcal{L}}\left(\theta_0, \widehat{\theta}_{n,c}(\theta_0)\right) \leqslant \log\frac{1}{\delta} + \gamma_{n,c}(\theta_0)\right\}$$

is an anytime valid confidence sequence for this exponential family.

Chowdhury, **PS** et al., COLT 2023.

I. Bregman concentration

Regularised estimator

$$\widehat{\theta}_{n,c}(\theta_0) = (\nabla\mathcal{L})^{-1}\left(\frac{1}{n+c}\left(\sum_{j=1}^{n} F(Y_j) + c\nabla\mathcal{L}(\theta_0)\right)\right)$$

$$\widehat{\Theta}_n^{\delta} = \left\{\theta_0 \in \Theta : (n+c)\mathcal{B}_{\mathcal{L}}\left(\theta_0, \widehat{\theta}_{n,c}(\theta_0)\right) \leqslant \log\tfrac{1}{\delta} + \gamma_{n,c}(\theta_0)\right\}$$

Bregman divergence

Bregman information gain

$$\gamma_{n,c}(\theta_0) = \log\left(\frac{\int_{\Theta} \exp\left(-c\mathcal{B}_{\mathcal{L}}(\theta', \theta_0)\right) d\theta'}{\int_{\Theta} \exp\left(-(n+c)\mathcal{B}_{\mathcal{L}}(\theta', \widehat{\theta}_{n,c}(\theta_0))\right) d\theta'}\right)$$

is an anytime valid confidence sequence for this exponential family.

Chowdhury, **PS** et al., COLT 2023.

I. Bregman concentration

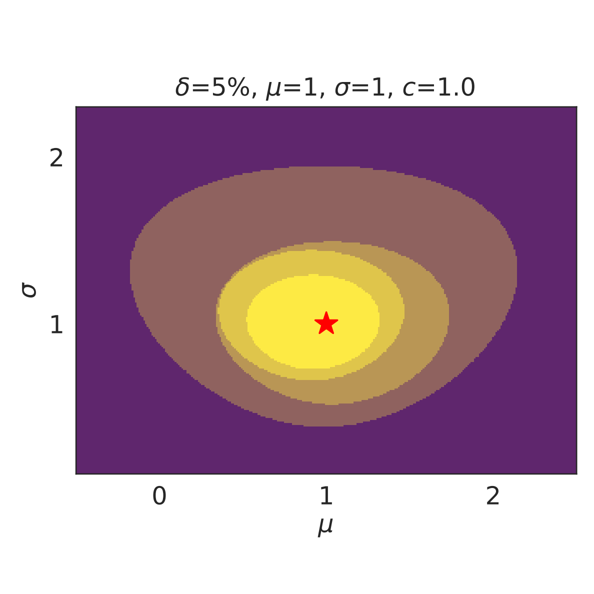

Chowdhury, **PS** et al., COLT 2023.

Challenges of healthcare recommendations

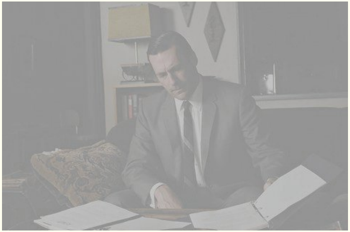

Online advertising

- ✔ Huge volume of data.
- ✔ Easy to model (Bernoulli, logistic...).
- ✔ Limited risks ($\rho = \mathbb{E}$).

Healthcare

- ✔ Slow and scarce data.
- ✘ Nonparametric models.
- ✘ Risky.

II. Nonparametric: bootstrapping from data only

Empirical average

- Greedy policy:

$$\pi_{t+1} = \operatorname*{argmax}_{k\in[K]} \frac{1}{N_t^k} \sum_{j=1}^{N_t^k} Y_j^k \,.$$

II. Nonparametric: bootstrapping from data only

- Greedy policy:

Empirical average

$$\pi_{t+1} = \operatorname*{argmax}_{k \in [K]} \sum_{j=1}^{N_t^k} w_{j,t}^k Y_j^k \, .$$

$w_{j,t}^k = \frac{1}{N_t^k}$

II. Nonparametric: bootstrapping from data only

- Greedy policy with fictitious rewards:

Empirical average

$$\pi_{t+1} = \underset{k \in [K]}{\operatorname{argmax}} \sum_{j=1}^{N_t^k} w_{j,t}^k Y_j^k + \widetilde{w}_t^k \widetilde{Y}_t^k \,.$$

$w_{j,t}^k, \widetilde{w}_t^k = \frac{1}{N_t^k + 1}$

Fictitious reward

Kveton, Szepesvari, et al. ICML, 2019.

II. Nonparametric: bootstrapping from data only

Randomised greedy policy with fictitious rewards:

Randomised average

$$\pi_{t+1} = \underset{k\in[K]}{\operatorname{argmax}} \sum_{j=1}^{N_t^k} w_{j,t}^k Y_j^k + \widetilde{w}_t^k \widetilde{Y}_t^k\ .$$

$w_{j,t}^k, \widetilde{w}_t^k \in [0,1]$ random

$$\sum_{j=1}^{N_t^k} w_{j,t}^k + \widetilde{w}_t^k = 1$$

Fictitious reward

Kveton, Szepesvari, et al. ICML, 2019.

II. Nonparametric: bootstrapping from data only

Dirichlet average

☞ Nonparametric Thompson sampling:

$$\pi_{t+1} = \underset{k\in[K]}{\operatorname{argmax}} \sum_{j=1}^{N_t^k} w_{j,t}^k Y_j^k + \widetilde{w}_t^k \widetilde{Y}_t^k \,.$$

$w_{j,t}^k, \widetilde{w}_t^k \sim \mathrm{Dir}(1,\ldots,1)$
(uniform distribution on the simplex)

Fictitious reward
($\approx$ reward upper bound)

Kveton, Szepesvari, et al. ICML, 2019.
Riou and Honda. ALT, 2020.

II. Nonparametric: bootstrapping from data only

Dirichlet average

Nonparametric Thompson sampling:

$$\pi_{t+1} = \operatorname*{argmax}_{k \in [K]} \sum_{j=1}^{N_t^k} w_{j,t}^k Y_j^k + \widetilde{w}_t^k \widetilde{Y}_t^k .$$

$w_{j,t}^k, \widetilde{w}_t^k \sim \mathrm{Dir}(1, \ldots, 1)$
(uniform distribution on the simplex)

Fictitious reward
($\approx$ reward upper bound)

Beyond bounded rewards (with known bounds)?

Kveton, Szepesvari, et al. ICML, 2019.
Riou and Honda. ALT, 2020.

II. Regret guarantees for Dirichlet sampling

$$\pi_{t+1} = \operatorname*{argmax}_{k \in [K]} \sum_{j=1}^{N_t^k} w_{j,t}^k Y_j^k + \widetilde{w}_t^k \widetilde{Y}_t^k \,.$$

- Regret guarantees (instance-dependent)
  - Optimal in several settings (bounded, bounded detectable, semibounded + quantile condition)

Light tailed: $\mathcal{R}_T = \mathcal{O}\left(\log(T) \log\log(T)\right)$.

Baudry, **PS** et al., NeurIPS 2021.

II. Dirichlet sampling

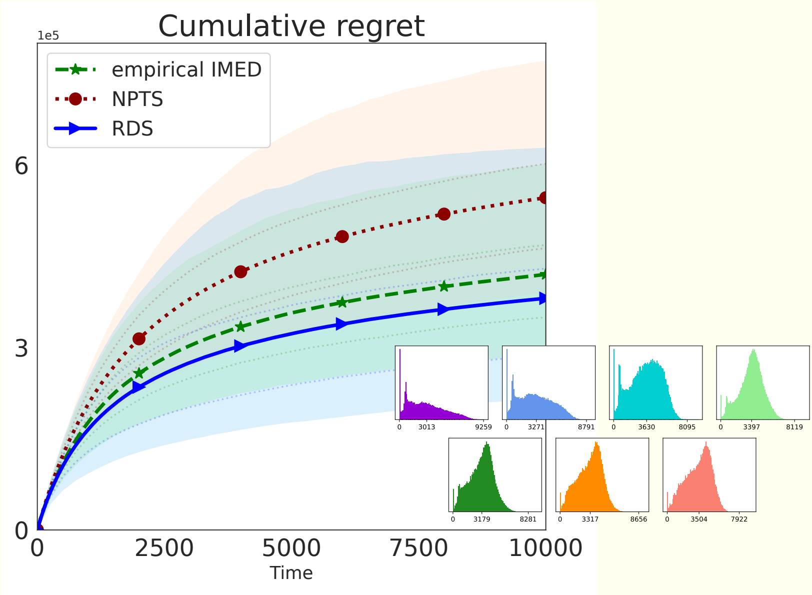

Baudry, **PS** et al., NeurIPS 2021.

Challenges of healthcare recommendations

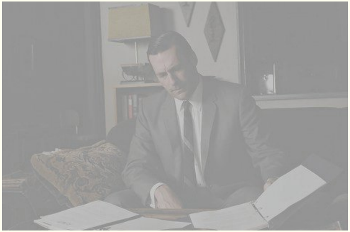

Online advertising

- ✔ Huge volume of data.
- ✔ Easy to model (Bernoulli, logistic...).
- ✔ Limited risks ($\rho = \mathbb{E}$).

Healthcare

- ✔ Slow and scarce data.
- ✔ Nonparametric models.
- ✖ Risky.

III. From risk-neutral to risk-aware

- Non-contextual bandits:

  - Maillard. Robust risk-averse stochastic multi-armed bandits.
    ALT, 2013.

  - Baudry, Gautron, et al. Optimal Thompson sampling strategies for support-aware CVAR bandits.
    ICML, 2021a.

- **Contextual bandits?**

III. From risk-neutral to risk-aware

- Classical linear bandit $Y = \langle \theta^\star, X \rangle +$ zero-mean noise:

$$\widehat{\theta}_{t+1} = \min_{\theta \in \Theta} \sum_{s=1}^{t} \left(Y_s - \langle \theta, X_s \rangle\right)^2 \quad + \quad \text{exploration}\,.$$

Estimator of
$\rho(\nu) = \mathbb{E}_\nu \left[Y \mid X \right]$

Abbasi-Yadkori, Pál, and Szepesvári. Improved algorithms for linear stochastic bandits. NIPS, 2011.

III. From risk-neutral to risk-aware

- Classical linear bandit $Y = \langle \theta^\star, X \rangle +$ zero-mean noise:

$$\widehat{\theta}_{t+1} = \min_{\theta \in \Theta} \sum_{s=1}^{t} \mathcal{L}\left(Y_s, \langle \theta, X_s \rangle\right) \quad + \quad \text{exploration}\,.$$

Estimator of

$$\rho(\nu) = \min_{\theta} \mathbb{E}_\nu \left[\mathcal{L}\left(Y, \langle \theta, X \rangle\right) \mid X\right]$$

(M-estimator)

$$\mathcal{L}(y, \xi) = (y - \xi)^2$$

Abbasi-Yadkori, Pál, and Szepesvári. Improved algorithms for linear stochastic bandits. NIPS, 2011.

III. From risk-neutral to risk-aware

- Elicitable risk measure linear bandit $Y \approx \langle \theta^\star, X \rangle +$ zero ρ-noise:

$$\widehat{\theta}_{t+1} = \min_{\theta \in \Theta} \sum_{s=1}^{t} \mathcal{L}\left(Y_s, \langle \theta, X_s \rangle\right) \quad + \quad \text{exploration}\,.$$

Estimator of
$\rho(\nu) = \min_{\theta} \mathbb{E}_\nu \left[\mathcal{L}\left(Y, \langle \theta, X \rangle\right) \mid X\right]$
(M-estimator)

$\mathcal{L}$ (strongly) convex mapping

Abbasi-Yadkori, Pál, and Szepesvári. Improved algorithms for linear stochastic bandits. NIPS, 2011.

III. From risk-neutral to risk-aware

Elicitable risk measure linear bandit $Y \approx \langle \theta^\star, X \rangle +$ zero ρ-noise:

$$\widehat{\theta}_{t+1} = \min_{\theta \in \Theta} \sum_{s=1}^{t} \mathcal{L}\left(Y_s, \langle \theta, X_s \rangle\right) \quad + \quad \text{exploration}\,.$$

Estimator of
$\rho(\nu) = \min_\theta \mathbb{E}_\nu \left[\mathcal{L}\left(Y, \langle \theta, X \rangle\right) \mid X\right]$
(M-estimator)

$\mathcal{L}$ (strongly) convex mapping

Expectile: $\mathcal{L}(y, \xi) = \left|p - \mathbb{1}_{y<\xi}\right| (y - \xi)^2$
($\neq$ loss for negative and positive rewards).

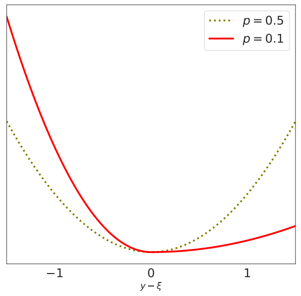

Abbasi-Yadkori, Pál, and Szepesvári. Improved algorithms for linear stochastic bandits. NIPS, 2011.

III. Risk-aware policies

- Deterministic exploration:

$$\pi_{t+1} = \operatorname{argmax}_x \left\langle \widehat{\theta}_{t+1}, x \right\rangle + \mathrm{UCB}(x),$$

Worst case: $\mathcal{R}_T = \widetilde{\mathcal{O}}\left(\sqrt{T}\right)$.

- Randomised Dirichlet exploration:

$$\widetilde{\theta}_{t+1} = \min_{\theta \in \Theta} \sum_{s=1}^{t} w_s^t \mathcal{L}\left(Y_s, \langle \theta, X_s \rangle\right) + \widetilde{w}_{t+1}^t \mathcal{L}\left(\widetilde{Y}, \langle \theta, \widetilde{X} \rangle\right),$$

$$\pi_{t+1} = \operatorname*{argmax}_x \left\langle \widetilde{\theta}_{t+1}, x \right\rangle,$$

Worst case: $\mathcal{R}_T \overset{?}{=} \widetilde{\mathcal{O}}\left(\sqrt{T}\right)$.

Linear Gaussian Thompson sampling?
(strong approximation of weighted bootstrap)

PS, Maillard, EWRL 2022; **PS**, Maillard, AISTATS 2023.

III. Risk-aware policies

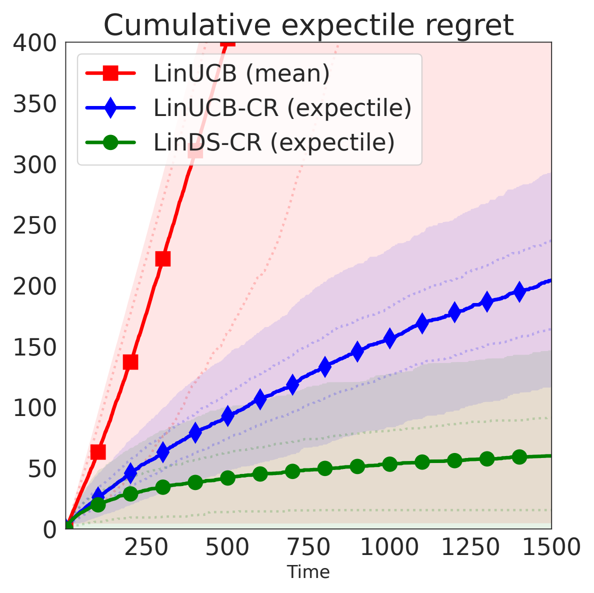

PS, Maillard, EWRL 2022; **PS**, Maillard, AISTATS 2023.

Challenges of healthcare recommendations

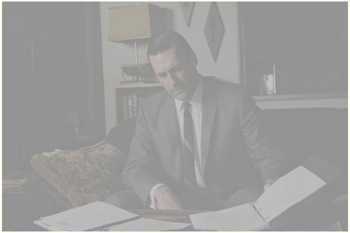

Online advertising

- Huge volume of data.
- Easy to model (Bernoulli, logistic...).
- Limited risks ($\rho = \mathbb{E}$).

Healthcare

- Slow and scarce data.
- Nonparametric models.
- Risky.

Conclusion and future works

Preoperative

Bariatric Weight Trajectory Prediction

→

Postperative

Prospective tool to improve follow-up (complications, costs)

Conclusion and future works

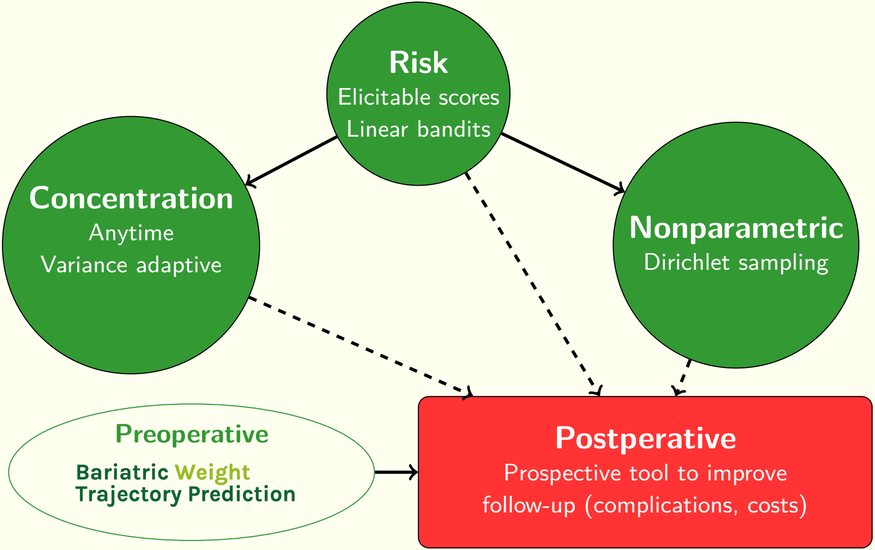

Conclusion and future works

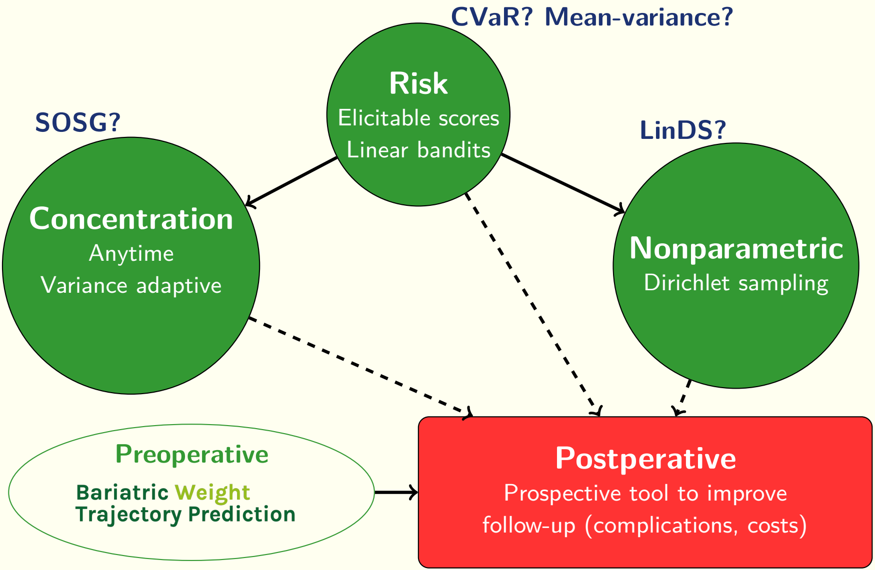

Contributions in medical sciences

Publications in peer-reviewed journals

- **PS**, Pierre Bauvin, Violeta Raverdy, Julien Teigny, Hélène Verkindt, Maxence Debert, Philippe Preux, François Pattou, et al. Development and validation of an interpretable machine learning based calculator for predicting 5 year-weight trajectories after bariatric surgery: a multinational retrospective cohort SOPHIA study.

  The Lancet Digital Health, 2023

- Robert Caiazzo, Pierre Bauvin, Camille Marciniak, **PS**, et al. Impact of robotic assistance on complications in bariatric surgery at expert laparoscopic surgery centers. a retrospective comparative study with propensity score.

  Annals of Surgery, 2023

Contributions in statistics and machine learning

Publications in peer-reviewed international conferences with proceedings

- Sayak Ray Chowdhury, **PS**, Odalric-Ambrym Maillard, and Aditya Gopalan. Bregman deviations of generic exponential families.
  In Conference On Learning Theory, 2023
- **PS** and Odalric-Ambrym Maillard. Risk-aware linear bandits with convex loss.
  In International Conference on Artificial Intelligence and Statistics, 2023
- Dorian Baudry, **PS**, and Odalric-Ambrym Maillard. From optimality to robustness: Adaptive re-sampling strategies in stochastic bandits.
  Advances in Neural Information Processing Systems, 2021b

Workshop presentations in peer-rewiewed international conferences

- **PS** and Odalric-Ambrym Maillard. Risk-aware linear bandits with convex loss.
  In European Workshop on Reinforcement Learning, 2022 (poster)

Ongoing

- **PS**. Empirical chernoff concentration: beyond bounded distributions
- **PS**, Romain Gautron, Odalric-Ambrym Maillard, Marc Corbeels, Chandra A. Madramootoe, and Nitin Joshi. Quantifying the uncertainty of crop management decisions based on crop model simulations

Software contributions

⌨ Bariatric Weight Trajectory Prediction

https://bwtp.univ-lille.fr

⌨ rlberry

https://github.com/rlberry-py/rlberry

⌨ concentration-lib

https://pypi.org/project/concentration-lib

Questions

Question at time $t-1$

Question at time t

Question at time $t+1$

Angry jury

Happy jury

Confused jury