Doctoral defense in Genetics

Combining gene regulatory networks and sequential learning for drug repurposing

Clémence Réda

Supervised by Andrée Delahaye-Duriez & Émilie Kaufmann
Université Paris Cité – ED BioSPC 562

September 9^{th}, 2022

Drug development is expensive[1], time-consuming[2] & prone to failure[3]

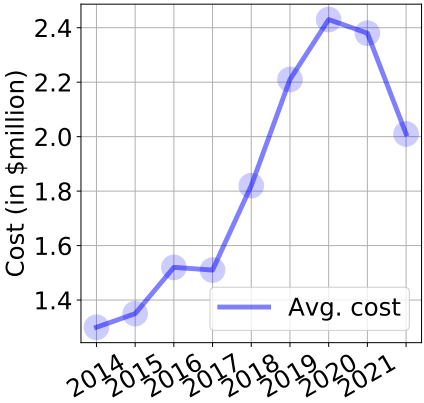

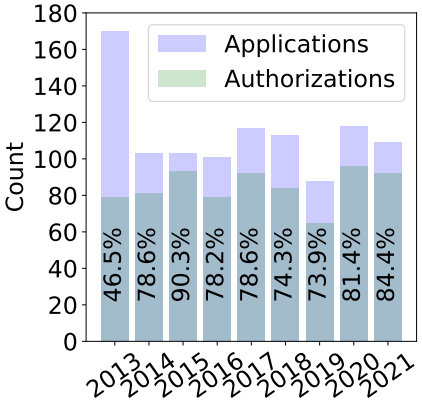

[1] *Deloitte Centre for Health Solutions (2022). https://www2.deloitte.com/. Accessed: [May 5, 2022].*

[2] *Sun et al. (2022). Acta Pharmaceutica Sinica B.*

[3] *European Medecine Agency (2021). https://www.ema.europa.eu/en. Accessed: [July 27, 2022].*

Drug repurposing

Screen a library of known molecules to find new therapeutic indications

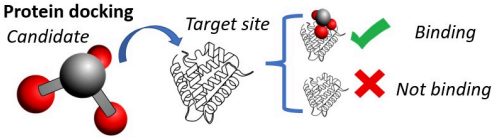

Pros[4]

- Known ADME-Tox profile (safer)
- Lower production cost (cheaper)

[4] *Réda and Kaufmann and Delahaye-Duriez (2020).* Comp. & struct. biotech. journ. (CSBJ) *18, pp. 241–252*

Key idea in signature reversion

Does the treatment "revert the gene level effects of the disease"?

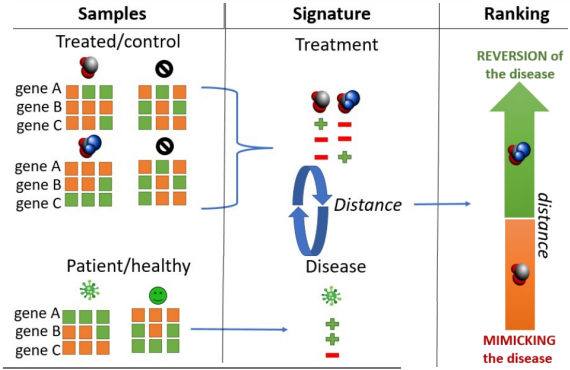

Examples

- CMap[5]
- L1000 CDS[26]
- Application to encephalopathy of prematurity (EoP)[7]

[5] *Lamb et al. (2006). Science 313.5795, pp. 1929–1935*

[6] *Duan et al. (2016). NPJ Systems Biology & Applications 2.1, pp. 1–12*

[7] *Bokobza and Réda and Guenoun et al. (In preparation)*

Signatures $\sim$ signed "fold changes"
(differential expression analysis)

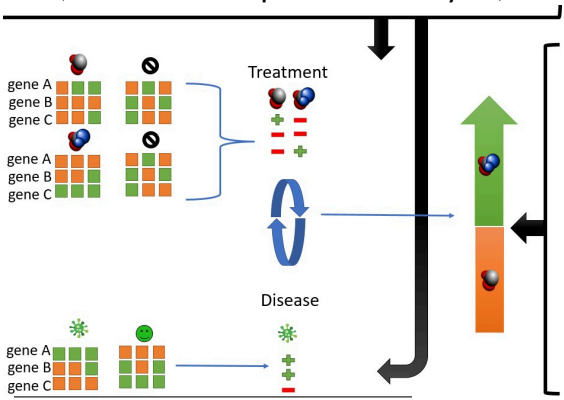

cosine score (L1000 CDS[2] [8])

$$\cos(v, w) := 1 - \frac{v \cdot w}{\|v\|_2 \|w\|_2},$$

many other measures[9]...

[8] *Duan et al. (2016). NPJ Systems Biology & Applications 2.1, pp. 1–12*

[9] *Musa et al. (2018). Briefings in bioinformatics 19.3, pp. 506–523*

(A) Appropriate evaluation of drug efficacy
regulatory interplay between genes, cell line-effect

(B) Design of a reproducible and generic approach
analysis of transcriptomic data, adapting to other diseases

(C) Control of the error in recommendation
error rate in the recommendations of candidates

Objectives | Signature reversion approach

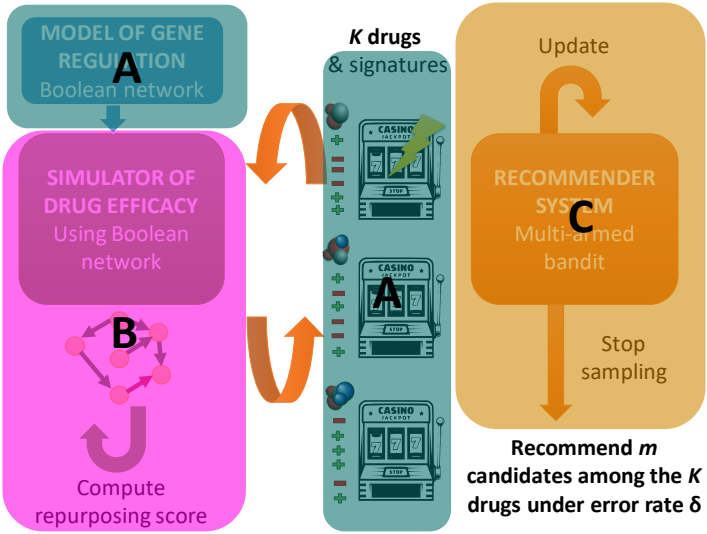

Epileptic crises: Transitory abnormal neuron electric discharge

- Linked to the hippocampus (brain region)[10]
- Affects ≈ 50 million people worldwide[11]
- $> 25\%$ of cases are refractory to known treatments[12]

[10] *Kuruba, Hattiangady, and Shetty (2009). Epilepsy & Behavior 14.1, pp. 65–73; Ogren et al. (2009). Annals of Neurology 66.6, pp. 783–791.*

[11] *World Health Organization (WHO) (2022). who.int/news-room/fact-sheets/detail/epilepsy.*

[12] *González et al. (2015). Neurología (Eng. ed.) 30.7, pp. 439–446.*

I. Analysis of a disease-specific regulatory network

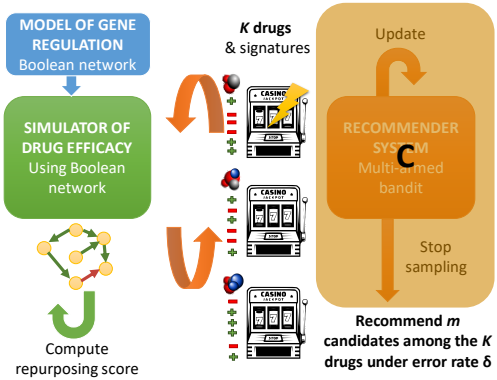

Gene regulatory networks

Graph: `vertices`=genes, `edges`=regulatory interactions

Boolean networks[13]

- Expression level $\in \{0, 1\}$
- Update of expression level: `regulatory functions`=logics

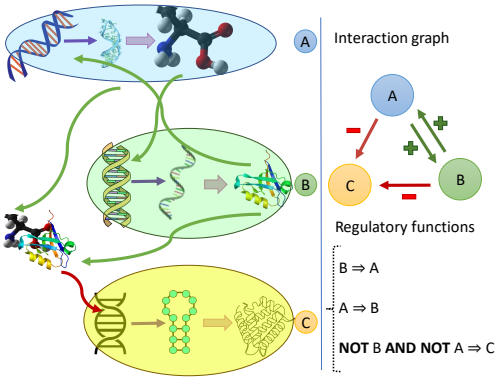

[13] *Kauffman (1969). Journal of theoretical biology 22.3, pp. 437–467; Thomas (1973). Journal of theoretical biology 42.3, pp. 563–585*

Predict the profile(s) resulting from gene regulation under the perturbation of the network (under a fixed type of *update*[14])

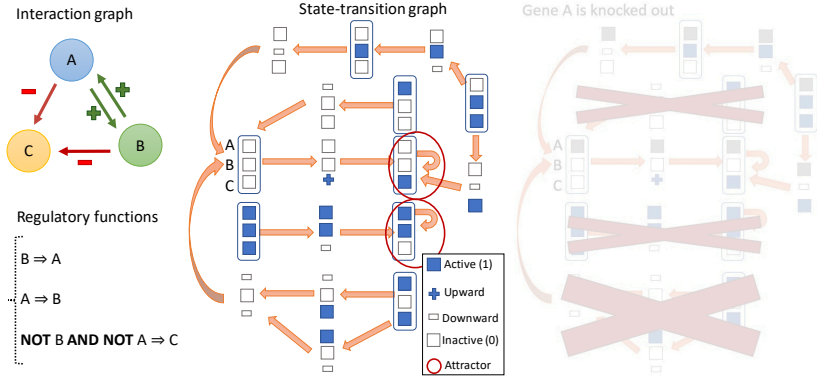

[14] *Paulevé et al. (2020). Nature communications 11.1, pp. 1–7*

Predict the profile(s) resulting from gene regulation under the perturbation of the network (under a fixed type of *update*[14])

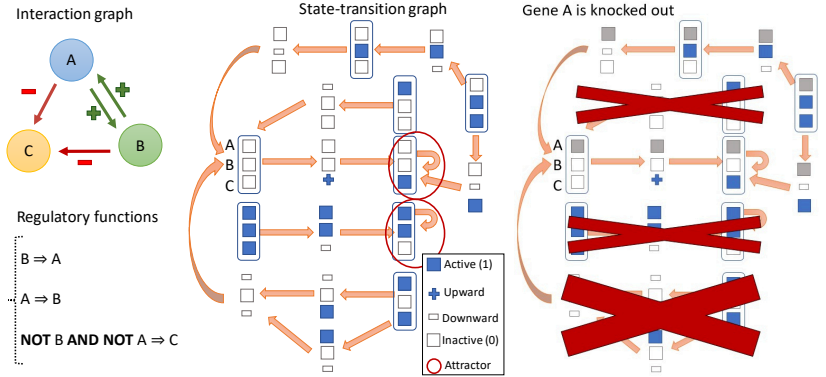

[14] *Paulevé et al. (2020). Nature communications 11.1, pp. 1–7*

Guided network inference | State-of-the-art

- **Building from maps** Regulatory map $\mapsto$ Boolean network (CasQ[15])

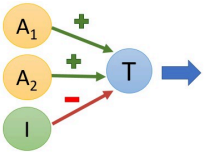

A_1 **AND** A_2 **AND NOT** $I \Rightarrow T$

⚠ How to get maps?

- **Guess from observations** Interaction graph & regulatory fcts matching observational data (Re:In, BoNeSiS[16])

⚠ How to select observational data?

[15] *Aghamiri et al. (2020). Bioinformatics 36.16, pp. 4473–4482.*

[16] *Dunn et al. (2014). Science 344.6188, pp. 1156–1160; Chevalier et al. (2019). ICTAI. IEEE, pp. 34–41.*

Guided network inference | Identification

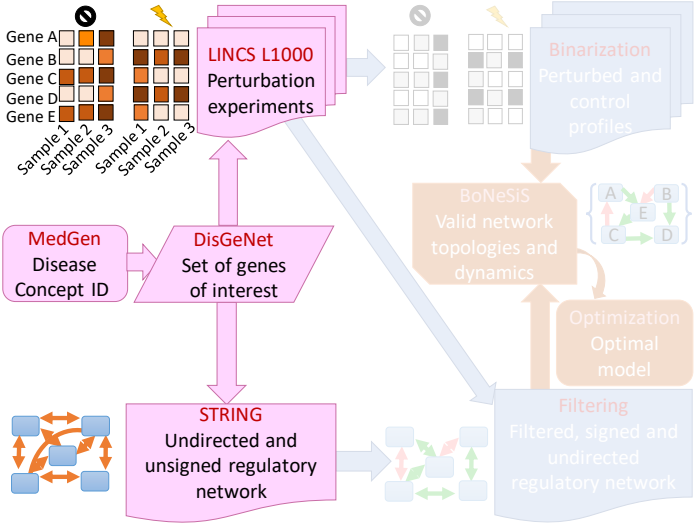

Method

Cell line-specific network from databases & methods[17]

[17] *Réda and Delahaye-Duriez (2022). Int. Conf. on Comp. Meth. in Sys. Bio. (CMSB), pp. 89–121*

Guided network inference | Identification

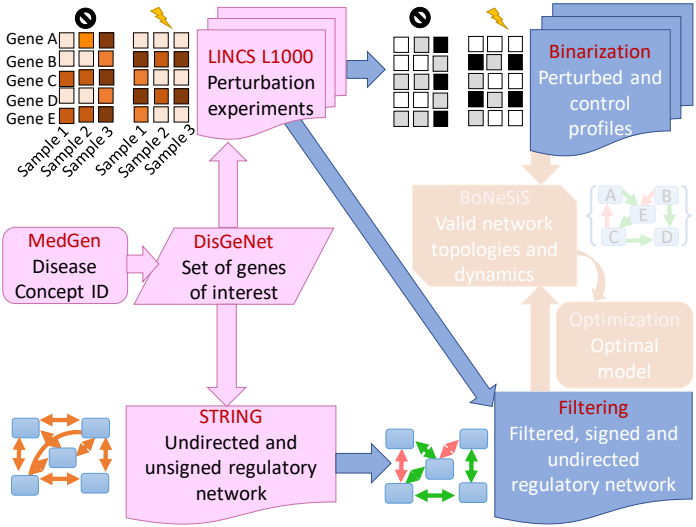

Method

Cell line-specific network from databases & methods[17]

[17] *Réda and Delahaye-Duriez (2022). Int. Conf. on Comp. Meth. in Sys. Bio. (CMSB), pp. 89–121*

Guided network inference | Identification

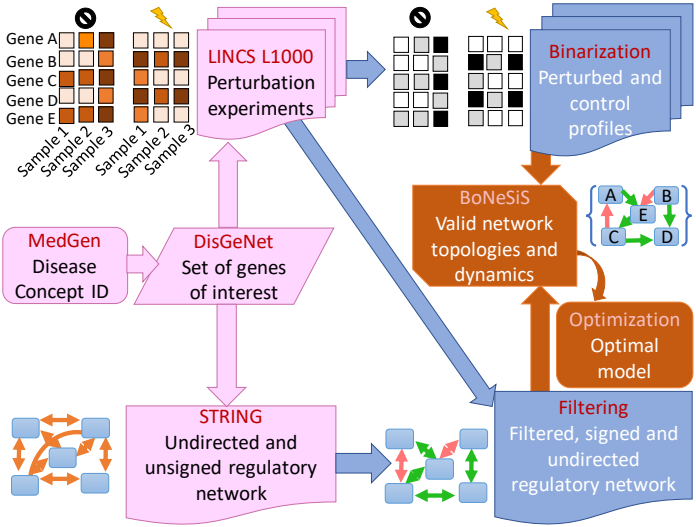

Method

Cell line-specific network from databases & methods[17]

[17] *Réda and Delahaye-Duriez (2022). Int. Conf. on Comp. Meth. in Sys. Bio. (CMSB), pp. 89–121*

Guided network inference | Application: Epilepsy

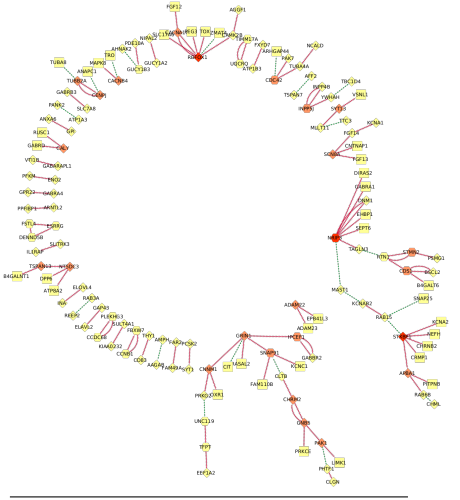

M30 gene module[18]
Identification of a brain cell network & key genes[19]

[18] *Delahaye-Duriez et al. (2016). Genome biology 17.1, pp. 1–18*

[19] *Réda and Delahaye-Duriez (2022). Int. Conf. on Comp. Meth. in Sys. Bio. (CMSB), pp. 89–121*

Simulation of efficacy | State-of-the-art

Compare reachable attractor states
from drug-treated patients

Signatures $\sim$ signed "fold changes"

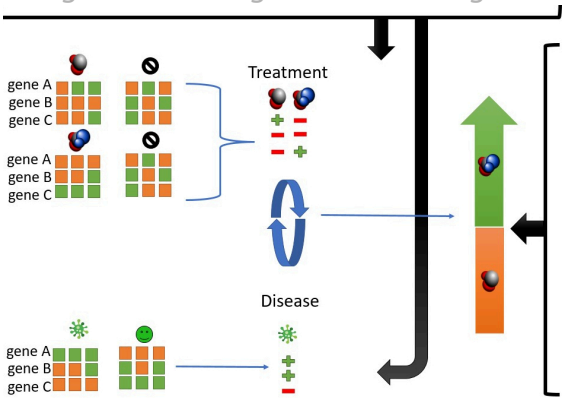

and healthy profiles on
well-chosen features

cosine score (L1000 CDS^2)

$$\cos(v, w) := 1 - \frac{v \cdot w}{\|v\|_2 \|w\|_2},$$

many other measures...

Signature $\approx$ drug effect on gene targets $\sim$ predict gene perturbations

Similar clusterings

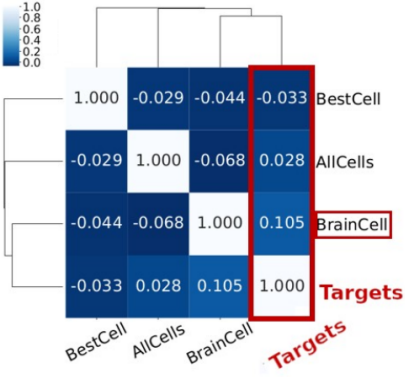

Similar gene targets

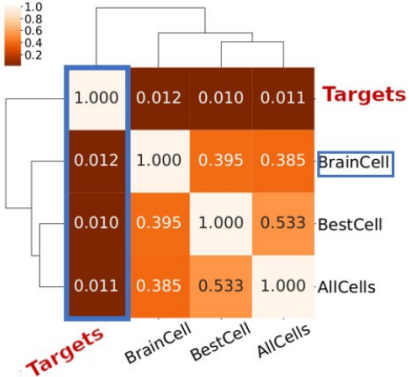

Signatures on brain cell lines (“BrainCell”)
closest to reported gene targets (“Targets”)

Simulation of efficacy | Scoring with the GRN

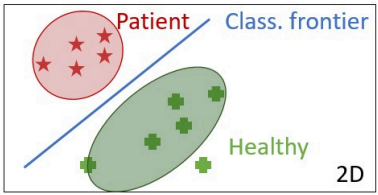

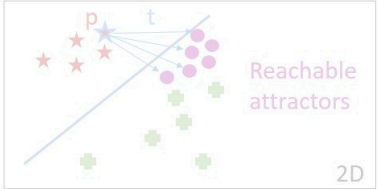

At fixed patient p and drug (signature) t

Key idea: Appropriate space to compute the distance from frontier to attractor

Aggregate distances for the final score

Simulation of efficacy | Scoring with the GRN

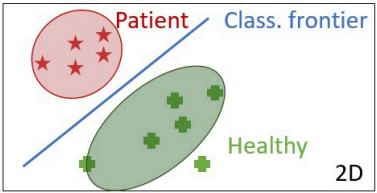

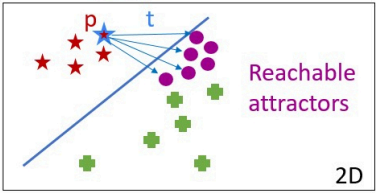

At fixed patient p and drug (signature) t

Key idea: Appropriate space to compute the distance from frontier to attractor

Aggregate distances for the final score

Simulation of efficacy | Scoring with the GRN

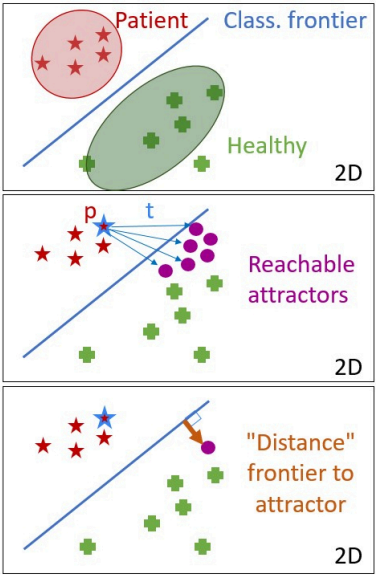

At fixed patient p and drug (signature) t

Key idea: Appropriate space to compute the distance from frontier to attractor

Aggregate distances for the final score

Baseline is L1000 CDS^2

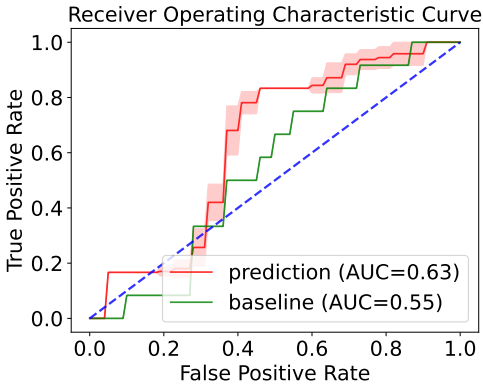

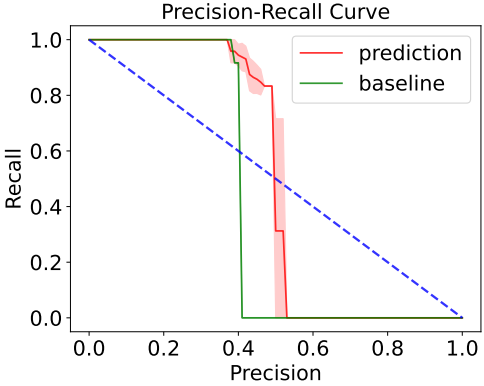

Simulation of efficacy | Sample-efficient testing

- Hardly tractable for hundreds or millions of drugs across patients
- Non-negligible variability in score across patient profiles

$$\text{std}(\textbf{AUC}) \approx 0.01$$

$\Longrightarrow$ design an adaptive sample-efficient procedure to test candidates

II. Recommendation through bandit algorithms

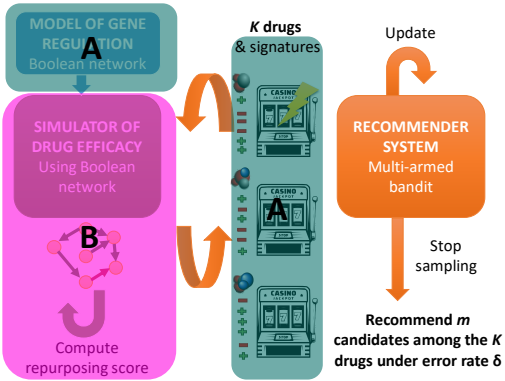

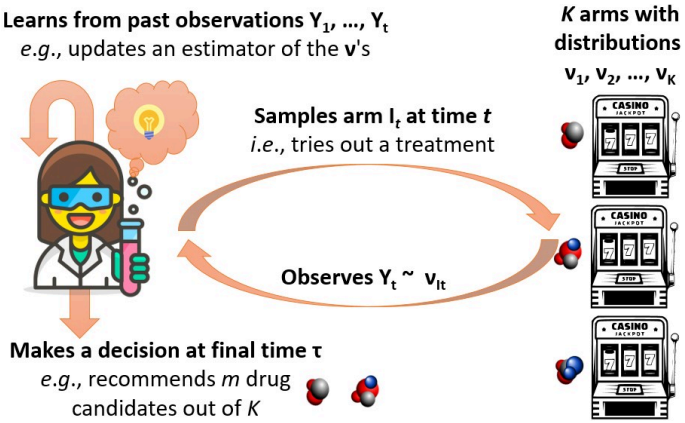

Observations $(Y_t)_t$ are *noisy*, of the form $Y_t = \underbrace{\mathbb{E}_{\{Y \sim \nu_{I_t}\}}[Y]}_{\text{mean of } \nu_{I_t}} + \underbrace{\phi_t}_{\text{1-subgaussian noise}}$

Leverage the info from $\underbrace{\text{signatures}}_{\text{on } d \text{ genes}}$ ("structure") to infer scores *faster*

$\mathcal{M}$-structured bandits

$\mathcal{M} \subseteq (\mathbb{R}^d \to \mathbb{R})$: class of functions,
$X = [x_1, \dots, x_K]^\top$ feature matrix where $x_i \in \mathbb{R}^d$ feature vector for arm i

$(\nu_1, \nu_2, \dots, \nu_K)$ is $\mathcal{M}$-structured if

$$\underbrace{\exists f \in \mathcal{M}}_{\text{shared across arms!}} \qquad \forall a \in \{1, \dots, K\}, \mathbb{E}_{\{Y \sim \nu_a\}}[Y] = f(x_a) \ .$$

Note: if $\mathcal{M} = (\mathbb{R}^d \to \mathbb{R})$: no assumption on the model (unstructured)

Bandit algorithms | FC-Top-m identification

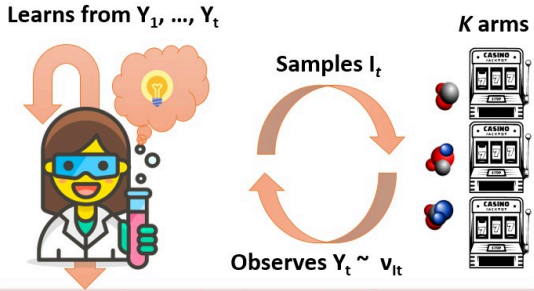

Design an algorithm $\mathfrak{A}$ that recommends $\widehat{S}$ of size m *s.t.*

$\mathfrak{A}$ is

δ-correct

For any model function $f \in \mathcal{M}$

$$\underbrace{\mathbb{P}_{\{\mathfrak{A},f\}}\left(\widehat{S} \not\subseteq \mathcal{S}^\star(f)\right) \leq \delta}_{\text{prob. of error is less than } \delta} .$$

Top-m arms $\mathcal{S}^\star(f) := \underbrace{\underset{k\in\{1,\ldots,K\}}{\overset{[m]}{\arg\max}}\ f(x_k)}_{m\text{-maximizers}}$

using only few samples: small τ

Linear models

$$\mathcal{M}_{\text{lin}} := \{x \mapsto \theta^\top x = \textstyle\sum_{i \leq d} \theta[i]x[i] : \theta \in \mathbb{R}^d\}$$

$\theta[i]$ quantifies the dependency of observations in feature $\sim$ gene i

$m = 1$: adaptive sampling [20], elimination-based [21]

adaptive sampling: sequentially samples & re-estimates

elimination-based: iteratively eliminates arms from a candidate set

$m > 1$: unstructured models only[22]

⚠ A more general principle for structured Top-m?

[20] *Xu et al. (2018). AISTATS. vol. 84.*

[21] *Fiez et al. (2019). NeurIPS 32.*

[22] *Kalyanakrishnan et al. (2012). ICML, pp. 655–662; Chen, Li, and Qiao (2017). AISTATS. vol. 54.*

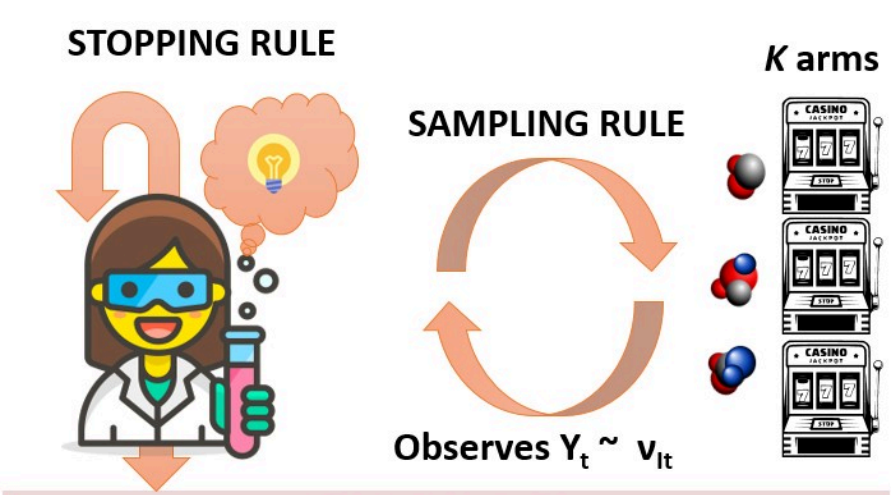
STOPPING RULE
SAMPLING RULE
K arms
CASINO
JACKPOT
STOP
Observes $Y_t \sim \nu_{I_t}$
DECISION RULE

Gap-Index Focused Algorithms (GIFA)[23]

1. *Partially specified* algorithms for structured Top-m
2. Any algorithm in GIFA is δ-correct
3. Unified *sample complexity* $\sim$ upper bound on #samples τ
   (for **linear** models)

[23] *Réda and Kaufmann and Delahaye-Duriez (2021). AISTATS. vol. 130. PMLR*

Gap Index. For $f \in \mathcal{M}$, pair of arms (i, j),

$$\mathcal{B}_{i,j}(\cdot) : \mathbb{N}^* \to \mathbb{R}^+ \text{ high prob. bound on gap } f(x_i) - f(x_j)$$

Good gap indices

$\mathcal{B}_{\cdot,\cdot}(\cdot)$ is a good gap index if $\mathbb{P}(\mathcal{E}) \geq 1 - \delta$, where

$$\mathcal{E} := \underbrace{\bigcap_{t \geq 0}}_{\text{at all times}} \underbrace{\bigcap_{k \in \mathcal{S}^\star(f)}}_{\text{any top-}m\text{ arm}} \underbrace{\bigcap_{j \notin \mathcal{S}^\star(f)}}_{\text{any non top-}m\text{ arm}} (\mathcal{B}_{k,j}(t) \geq f(x_k) - f(x_j))$$

Linear bandit models | 1. GIFA

$\hat{\mu}_a(t)$ is the empirical average of rewards from arm a up to t

⚠ **Any** SAMPLING & good gap index $\mathcal{B}_{\cdot,\cdot}(\cdot)$! $\Leftarrow$ embedding of structure f

LUCB-GIFA

repeat

$\hat{\alpha}_a(t) := \max_{b \notin J(t)} \mathcal{B}_{b,a}(t)$

$J(t) \leftarrow m$ m-maximizers of $\hat{\mu}_\cdot(t)$

$b_t \leftarrow \arg\max_{j \in J(t)} \hat{\alpha}_j(t)$

$c_t \leftarrow \arg\max_{c \notin J(t)} \mathcal{B}_{c,b_t}(t)$

$l_t \leftarrow \mathrm{SAMPLING}(b_t, c_t)$

Observe, update, $t \leftarrow t+1$

until $\hat{\alpha}_{b_t}(t) \leq 0$ # STOPPING

return J(t) # DECISION

Gap-GIFA

repeat

$\hat{\omega}_a(t) := \max^{m}_{b \neq a} \mathcal{B}_{b,a}(t)$

$J(t) \leftarrow m$ m-minimizers of $\hat{\omega}_a(t)$

$b_t \leftarrow \arg\max_{j \in J(t)} \hat{\omega}_j(t)$

$c_t \leftarrow \arg\max_{c \notin J(t)} \mathcal{B}_{c,b_t}(t)$

$l_t \leftarrow \mathrm{SAMPLING}(b_t, c_t)$

Observe, update, $t \leftarrow t+1$

until $\hat{\omega}_{b_t}(t) \leq 0$ # STOPPING

return J(t) # DECISION

Linear bandit models | 1. GIFA

$\hat{\mu}_a(t)$ is the empirical average of rewards from arm a up to t

⚠ **Any** SAMPLING & good gap index $\mathcal{B}_{\cdot,\cdot}(\cdot)$! $\Leftarrow$ embedding of structure f

LUCB-GIFA

repeat

$\hat{\alpha}_a(t) := \max_{b \notin J(t)} \mathcal{B}_{b,a}(t)$

$J(t) \leftarrow m$ m-maximizers of $\hat{\mu}_\cdot(t)$

$b_t \leftarrow \arg\max_{j \in J(t)} \hat{\alpha}_j(t)$

$c_t \leftarrow \arg\max_{c \notin J(t)} \mathcal{B}_{c,b_t}(t)$

$l_t \leftarrow \mathrm{SAMPLING}(b_t, c_t)$

Observe, update, $t \leftarrow t+1$

until $\hat{\alpha}_{b_t}(t) \leq 0$ # STOPPING

return J(t) # DECISION

Gap-GIFA

repeat

$\hat{\omega}_a(t) := \max^{m}_{b \neq a} \mathcal{B}_{b,a}(t)$

$J(t) \leftarrow m$ m-minimizers of $\hat{\omega}_a(t)$

$b_t \leftarrow \arg\max_{j \in J(t)} \hat{\omega}_j(t)$

$c_t \leftarrow \arg\max_{c \notin J(t)} \mathcal{B}_{c,b_t}(t)$

$l_t \leftarrow \mathrm{SAMPLING}(b_t, c_t)$

Observe, update, $t \leftarrow t+1$

until $\hat{\omega}_{b_t}(t) \leq 0$ # STOPPING

return J(t) # DECISION

Linear bandit models | 1. GIFA

$\hat{\mu}_a(t)$ is the empirical average of rewards from arm a up to t

⚠ **Any** SAMPLING & good gap index $\mathcal{B}_{\cdot,\cdot}(\cdot)$! $\Leftarrow$ embedding of structure f

LUCB-GIFA

repeat

$\hat{\alpha}_a(t) := \max_{b \notin J(t)} \mathcal{B}_{b,a}(t)$

$J(t) \leftarrow m$ m-maximizers of $\hat{\mu}_\cdot(t)$

$b_t \leftarrow \arg\max_{j \in J(t)} \hat{\alpha}_j(t)$

$c_t \leftarrow \arg\max_{c \notin J(t)} \mathcal{B}_{c,b_t}(t)$

$I_t \leftarrow$ SAMPLING(b_t, c_t)

Observe, update, $t \leftarrow t+1$

until $\hat{\alpha}_{b_t}(t) \leq 0$ # STOPPING

return J(t) # DECISION

Gap-GIFA

repeat

$\hat{\omega}_a(t) := \max^{m}_{b \neq a} \mathcal{B}_{b,a}(t)$

$J(t) \leftarrow m$ m-minimizers of $\hat{\omega}_a(t)$

$b_t \leftarrow \arg\max_{j \in J(t)} \hat{\omega}_j(t)$

$c_t \leftarrow \arg\max_{c \notin J(t)} \mathcal{B}_{c,b_t}(t)$

$I_t \leftarrow$ SAMPLING(b_t, c_t)

Observe, update, $t \leftarrow t+1$

until $\hat{\omega}_{b_t}(t) \leq 0$ # STOPPING

return J(t) # DECISION

LUCB-GIFA

repeat

$\hat{\alpha}_a(t) := \max_{b \notin J(t)} \mathcal{B}_{b,a}(t)$

$J(t) \leftarrow m$ m-maximizers of $\hat{\mu}_\cdot(t)$

$b_t \leftarrow \arg\max_{j \in J(t)} \hat{\alpha}_j(t)$

$c_t \leftarrow \arg\max_{c \notin J(t)} \mathcal{B}_{c,b_t}(t)$

$I_t \leftarrow$ SAMPLING(b_t, c_t)

Observe, update, $t \leftarrow t+1$

until $\hat{\alpha}_{b_t}(t) \leq 0$ # STOPPING

return J(t) # DECISION

Why does LUCB-GIFA work?

- $J(\cdot)$: set of empirical best arms
- $\hat{\alpha}_{b_t}(t) \approx$ min. gap between empirical worst & best arms

$$\max_{c \notin J(t)} \mathcal{B}_{c,b_t}(t) \geq \max_{c \notin J(t)} f(x_c) - f(x_{b_t})$$

$$= \left(\max_{c \notin J(t)} f(x_c) \right) - \left(\min_{b \in J(t)} f(x_b) \right)$$

Why does Gap-GIFA work?

- $\hat{\omega}_{b_t}(t) \approx$ min. gap between true & empirical best arms

$$\max_{c \neq b_t}^{m} \mathcal{B}_{c,b_t}(t) \geq \max_{k \in \{1,\dots,K\}}^{m} f(x_k) - f(x_{b_t})$$

$$= \left(\min_{c \in \mathcal{S}^\star(f)} f(x_c) \right) - \left(\min_{b \in J(t)} f(x_b) \right)$$

Gap-GIFA

repeat

$\hat{\omega}_a(t) := \max_{b \neq a}^{m} \mathcal{B}_{b,a}(t)$

$J(t) \leftarrow m$ m-minimizers of $\hat{\omega}_a(t)$

$b_t \leftarrow \arg\max_{j \in J(t)} \hat{\omega}_j(t)$

$c_t \leftarrow \arg\max_{c \notin J(t)} \mathcal{B}_{c,b_t}(t)$

$l_t \leftarrow$ SAMPLING(b_t, c_t)

Observe, update, $t \leftarrow t+1$

until $\hat{\omega}_{b_t}(t) \leq 0$ # STOPPING

return J(t) # DECISION

Any GIFA is δ-correct

On the "good event" $\mathcal{E}$ of probability $1-\delta$ (*good gap index*),

$$\text{at final time } \tau,\ J(\tau) \subseteq \mathcal{S}^\star(f)\ (\star)$$

Proof (Gap-GIFA). $\neg(\star)$: there exists $b \in J(\tau) \cap (\mathcal{S}^\star(f))^c$
If there were no $c \in \mathcal{S}^\star(f)$ *s.t.* $\mathcal{B}_{c,b}(\tau) \leq 0$, then

$$\underbrace{0 \geq \hat{\omega}_b(\tau)}_{\texttt{STOPPING} \text{ and } b \in J(\tau)} = \max_{a \neq b}^{m} \mathcal{B}_{a,b}(\tau), \underbrace{\{\mathcal{B}_{c,b}(t) : c \in \mathcal{S}^\star(f)\}}_{\geq m \text{ elements } (b \notin \mathcal{S}^\star(f))} > 0 \Rightarrow \underline{c \text{ exists}}$$

That means $\underbrace{\min_{k \in \mathcal{S}^\star(f)} f(x_k) - f(x_b) \leq \mathcal{B}_{c,b}(t)}_{\min_{k \in \mathcal{S}^\star(f)} f(x_k) \leq f(x_c) \text{ and } \mathcal{E} \text{ holds}} \leq 0 \Rightarrow \underline{b \in \mathcal{S}^\star(f)}$

Any GIFA is δ-correct

On the "good event" $\mathcal{E}$ of probability $1-\delta$ (*good gap index*),

$$\text{at final time } \tau,\ J(\tau) \subseteq \mathcal{S}^\star(f)\ (\star)$$

Proof (Gap-GIFA). $\neg(\star)$: there exists $b \in J(\tau) \cap (\mathcal{S}^\star(f))^c$
If there were no $c \in \mathcal{S}^\star(f)$ *s.t.* $\mathcal{B}_{c,b}(\tau) \leq 0$, then

$$\underbrace{0 \geq \hat{\omega}_b(\tau)}_{\texttt{STOPPING} \text{ and } b \in J(\tau)} = \max_{a \neq b}^{m} \mathcal{B}_{a,b}(\tau), \underbrace{\{\mathcal{B}_{c,b}(t) : c \in \mathcal{S}^\star(f)\}}_{\geq m \text{ elements } (b \notin \mathcal{S}^\star(f))} > 0 \Rightarrow \underline{c \text{ exists}}$$

That means $\underbrace{\min_{k \in \mathcal{S}^\star(f)} f(x_k) - f(x_b) \leq \mathcal{B}_{c,b}(t)}_{\min_{k \in \mathcal{S}^\star(f)} f(x_k) \leq f(x_c) \text{ and } \mathcal{E} \text{ holds}} \leq 0 \Rightarrow \underline{b \in \mathcal{S}^\star(f)}$

Any GIFA is δ-correct

On the "good event" $\mathcal{E}$ of probability $1-\delta$ (*good gap index*),

$$\text{at final time } \tau,\ J(\tau) \subseteq \mathcal{S}^\star(f)\ (\star)$$

Proof (Gap-GIFA). $\neg(\star)$: there exists $b \in J(\tau) \cap (\mathcal{S}^\star(f))^c$
If there were no $c \in \mathcal{S}^\star(f)$ *s.t.* $\mathcal{B}_{c,b}(\tau) \leq 0$, then

$$\underbrace{0 \geq \hat{\omega}_b(\tau)}_{\texttt{STOPPING} \text{ and } b \in J(\tau)} = \max_{a \neq b}^{m} \mathcal{B}_{a,b}(\tau), \underbrace{\{\mathcal{B}_{c,b}(t) : c \in \mathcal{S}^\star(f)\}}_{\geq m \text{ elements } (b \notin \mathcal{S}^\star(f))} > 0 \Rightarrow \underline{c \text{ exists}}$$

That means $\underbrace{\min_{k \in \mathcal{S}^\star(f)} f(x_k) - f(x_b) \leq \mathcal{B}_{c,b}(t) \leq 0}_{\min_{k \in \mathcal{S}^\star(f)} f(x_k) \leq f(x_c) \text{ and } \mathcal{E} \text{ holds}} \Rightarrow \underline{b \in \mathcal{S}^\star(f)}$

Any GIFA is δ-correct

On the "good event" $\mathcal{E}$ of probability $1-\delta$ (*good gap index*),

$$\text{at final time } \tau,\ J(\tau) \subseteq \mathcal{S}^\star(f)\ (\star)$$

Proof (Gap-GIFA). $\neg(\star)$: there exists $b \in J(\tau) \cap (\mathcal{S}^\star(f))^c$
If there were no $c \in \mathcal{S}^\star(f)$ *s.t.* $\mathcal{B}_{c,b}(\tau) \leq 0$, then

$$\underbrace{0 \geq \hat{\omega}_b(\tau)}_{\texttt{STOPPING} \text{ and } b \in J(\tau)} = \max_{a \neq b}^{m} \mathcal{B}_{a,b}(\tau), \underbrace{\{\mathcal{B}_{c,b}(t) : c \in \mathcal{S}^\star(f)\}}_{\geq m \text{ elements } (b \notin \mathcal{S}^\star(f))} > 0 \Rightarrow \underline{c \text{ exists}}$$

That means $\underbrace{\min_{k \in \mathcal{S}^\star(f)} f(x_k) - f(x_b) \leq \mathcal{B}_{c,b}(t)}_{\min_{k \in \mathcal{S}^\star(f)} f(x_k) \leq f(x_c) \text{ and } \mathcal{E} \text{ holds}} \leq 0 \Rightarrow \underline{b \in \mathcal{S}^\star(f)}$

Any GIFA is δ-correct

On the "good event" $\mathcal{E}$ of probability $1-\delta$ (*good gap index*),

$$\text{at final time } \tau,\ J(\tau) \subseteq \mathcal{S}^\star(f)\ (\star)$$

Proof (Gap-GIFA). $\neg(\star)$: there exists $b \in J(\tau) \cap (\mathcal{S}^\star(f))^c$
If there were no $c \in \mathcal{S}^\star(f)$ *s.t.* $\mathcal{B}_{c,b}(\tau) \leq 0$, then

$$\underbrace{0 \geq \hat{\omega}_b(\tau)}_{\texttt{STOPPING} \text{ and } b \in J(\tau)} = \max_{a \neq b}^{m} \mathcal{B}_{a,b}(\tau), \underbrace{\{\mathcal{B}_{c,b}(t) : c \in \mathcal{S}^\star(f)\}}_{\geq m \text{ elements } (b \notin \mathcal{S}^\star(f))} > 0 \Rightarrow \underline{c \text{ exists}}$$

That means $\underbrace{\min_{k \in \mathcal{S}^\star(f)} f(x_k) - f(x_b) \leq \mathcal{B}_{c,b}(t)}_{\min_{k \in \mathcal{S}^\star(f)} f(x_k) \leq f(x_c) \text{ and } \mathcal{E} \text{ holds}} \leq 0 \Rightarrow \underline{b \in \mathcal{S}^\star(f)}$

Good gap indices: $N_a(t)$: # times a is sampled up to t

(unstr) $\mathcal{B}_{i,j}(t) := \underbrace{\hat{\mu}_i(t) - \hat{\mu}_j(t)}_{\text{empirical gap}} + \underbrace{\overbrace{\mathcal{T}_\delta(t)}^{\text{good gap index}} \quad \overbrace{1/\sqrt{N_i(t)} + 1/\sqrt{N_j(t)}}^{\text{past obs + no feature}}}_{\text{quantifies uncertainty without prior}}$

(linear) $\mathcal{B}_{i,j}(t) := \underbrace{\hat{\mu}_i(t) - \hat{\mu}_j(t)}_{\text{empirical gap}} + \underbrace{\overbrace{\mathcal{T}_\delta(t)}^{\text{good gap index}} \quad \overbrace{\|x_i - x_j\|_{(\sum_a N_a(t) x_a x_a^\top)^{-1}}}^{\text{past obs + features}}}_{\text{quantifies uncertainty under linearity}}$

$\mathcal{T}_\delta(t)^2 = \mathcal{O}\left(\ln\left(\frac{Kt^4}{\delta}\right)\right)$ (unstr[24]) $\quad \mathcal{T}_\delta(t)^2 = \mathcal{O}\left(\ln\left(\frac{1}{\delta}\right) + d\ln\left(\frac{t}{d}\right)\right)$ (linear[25])

[24] *Hoeffding (1994). The collected works of Wassily Hoeffding. Springer, pp. 409–426.*

[25] *Abbasi-Yadkori, Pál, and Szepesvári (2011). Advances in neural information processing systems 24.*

⚠ The upper bound on τ depends on SAMPLING (and def. of $\mathcal{B}_{\cdot,\cdot}()$)

SAMPLING: *e.g.*, "highest uncertainty" b/w b_t and challenger c_t

(unstr) $I_t \in \arg\min_{a \in \{b_t, c_t\}} N_a(t)$ (linear) $I_t \in \arg\max_{a \in \{b_t, c_t\}} \|x_a\|_{(\sum_a N_a(t) x_a x_a^\top)^{-1}}$

For linear models: decreasing the variance in direction $x_t := x_{b_t} - x_{c_t}$ [26]

$I_t \in \arg\min_{a \in [K]} \|x_t\|_{(\sum_b (N_b(t) + \mathbb{1}_{b=a}) x_b x_b^\top)^{-1}}$

(greedy)

Compute allocation $\omega(t) \in [0,1]^K$

$I_t \in \arg\max_{a \in [K]} N_a(t) / \omega_a(t)$

(optimized)

[26] *Xu et al. (2018). AISTATS. vol. 84.*

For "highest uncertainty" (unstr, linear) & "optimized" (linear) LUCB-GIFA

Upper bound on sample complexity in (unstr. & linear) GIFA

$$\tau = \mathcal{O}\left(\mathcal{H}(f)\ln\left(\mathcal{H}(f)/\delta\right)\right) \text{ where } \mathcal{H}(f) = \mathcal{O}\left(\sum_{a}(\Delta_{a\text{ bad}} \vee \Delta_{a\text{ top-}m})^{-2}\right),$$

Characteristic gaps

$$\Delta_{a\text{ bad}} := \min_{k\in\mathcal{S}^\star(f)} f(x_k) - f(x_a) \qquad \Delta_{a\text{ top-}m} := f(x_a) - \max_{k\notin\mathcal{S}^\star(f)} f(x_k)$$

⚠ but the scaling might be smaller for linear models

Non-linear bandit models | Misspecification

Why not use linear algorithms on non-linear bandit models?

- Choose $f_{\text{lin}} \in \mathcal{M}_{\text{lin}}$ (here min. gap ≈ 0.28, #arms$= 10$, #dim$= 5$)
- Run any linear GIFA to find the $m = 3$ best arms in

$$f_\psi := x \mapsto f_{\text{lin}}(x) + \psi \times \mathbb{1}(x = x_4)$$

for $\delta = 5\%$ across 100 iterations

ψ	0 (linear)	1	2	2.7	2.8	3	4	5
Empirical error (%)	0	0	0	1	6	28	100	100
δ-correct?	✓	✓	✓	✓	×	×	×	×

A nearly optimal algorithm exists, with a dependency on ψ (MisLid[27])

[27] *Réda and Tirinzoni and Degenne (2021). Advances in Neural Information Processing Systems 34.*

III. Application to epilepsy

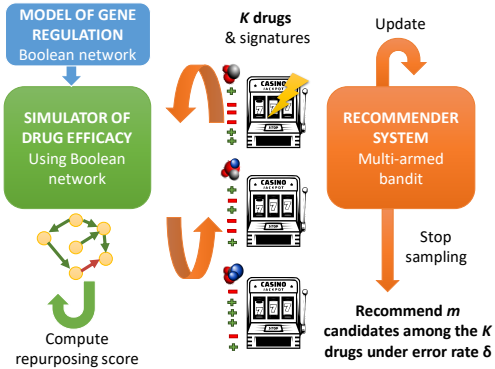

Instance (Part I.)

Scoring[28] on $K = 21$ drugs across #patients $= 24$ using $d = 10$ features*

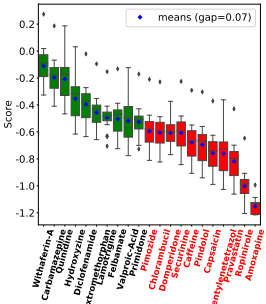

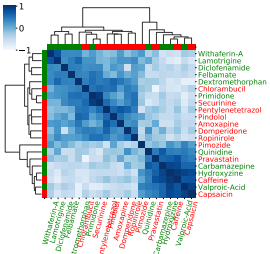

* PCA on drug signatures (196 genes)

Objective (Part II.)

Find the $m = 3$ best treatments under error $\delta = 10\%$

Observes a raw score for a randomly selected patient

[28] *Réda and Delahaye-Duriez (2022). Int. Conf. on Comp. Meth. in Sys. Bio. (CMSB), pp. 89–121*

Antiepileptics & Proconvulsant drugs | Results

for $\delta = 10\%$ across 100 iterations

Algorithm	error $\hat{\delta} \leq \delta$?	sample complexity $\hat{\tau}$
Unstructured LUCB-GIFA	True	$37,425 \pm 271$
Linear LUCB-GIFA	False	157 ± 13
Linear Gap-GIFA	False	91 ± 14
MisLid ($\psi = 1$)	**True**	$16,189 \pm 1,492$

	Correct?	Sample-efficient?
Unstructured	✓	×
Linear	×	✓
Misspecified	✓	✓

Conclusion

Conclusion | Combination of GRNs & bandits

Application to real life dataset for epilepsy

Obj. A

Disease network

Obj. B

Network reversion

Obj. C

Structured Top-m

(linear , misspecified)

Collaborative CTs

Signature reversion for EoP & literature review

Conclusion | Combination of GRNs & bandits

Application to real life dataset for epilepsy

Obj. A

Disease network [29]

Obj. B

Network reversion

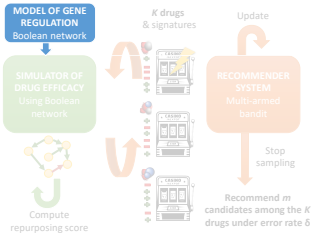

Obj. C

Structured Top-m

(linear , misspecified)

Collaborative CTs

Signature reversion for EoP & literature review

[29] *Réda and Delahaye-Duriez (2022). Int. Conf. on Comp. Meth. in Sys. Bio. (CMSB), pp. 89–121*

Conclusion | Combination of GRNs & bandits

Application to real life dataset for epilepsy

Obj. A

Disease network [29]

Obj. B

Network reversion

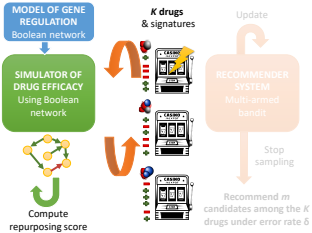

Obj. C

Structured Top-m

(linear , misspecified)

Collaborative CTs

Signature reversion for EoP[30] & literature review[31]

[29] *Réda and Delahaye-Duriez (2022). Int. Conf. on Comp. Meth. in Sys. Bio. (CMSB), pp. 89–121*

[30] *Bokobza and Réda and Guenoun et al. (In preparation)*

[31] *Réda and Kaufmann and Delahaye-Duriez (2020). Comp. & struct. biotech. journ. (CSBJ) 18, pp. 241–252*

Conclusion | Combination of GRNs & bandits

Application to real life dataset for epilepsy

Obj. A

Disease network [29]

Obj. B

Network reversion

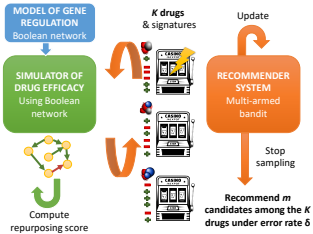

Obj. C

Structured Top-m

(linear [30], misspecified)

Collaborative CTs

Signature reversion for EoP[31] & literature review[32]

---

[29] *Réda and Delahaye-Duriez (2022). Int. Conf. on Comp. Meth. in Sys. Bio. (CMSB), pp. 89–121*

[30] *Réda and Kaufmann and Delahaye-Duriez (2021). AISTATS. vol. 130. PMLR*

[31] *Bokobza and Réda and Guenoun et al. (In preparation)*

[32] *Réda and Kaufmann and Delahaye-Duriez (2020). Comp. & struct. biotech. journ. (CSBJ) 18, pp. 241–252*

Conclusion | Combination of GRNs & bandits

Application to real life dataset for epilepsy

Obj. A

Disease network [29]

Obj. B

Network reversion

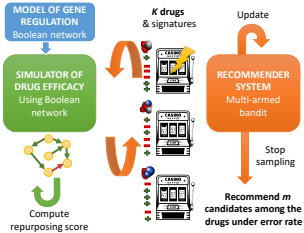

Obj. C

Structured Top-m

(linear [30], misspecified[31])

Collaborative CTs

Signature reversion for EoP[32] & literature review[33]

[29] *Réda and Delahaye-Duriez (2022). Int. Conf. on Comp. Meth. in Sys. Bio. (CMSB), pp. 89–121*

[30] *Réda and Kaufmann and Delahaye-Duriez (2021). AISTATS. vol. 130. PMLR*

[31] *Réda and Tirinzoni and Degenne (2021). Advances in Neural Information Processing Systems 34*

[32] *Bokobza and Réda and Guenoun et al. (In preparation)*

[33] *Réda and Kaufmann and Delahaye-Duriez (2020). Comp. & struct. biotech. journ. (CSBJ) 18, pp. 241–252*

Conclusion | Combination of GRNs & bandits

Application to real life dataset for epilepsy

Obj. A

Disease network [29]

Obj. B

Network reversion

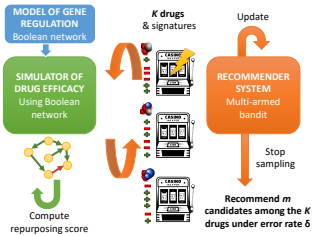

Obj. C

Structured Top-m

(linear [30], misspecified[31])

Collaborative CTs[32]

Signature reversion for EoP[33]& literature review[34]

[29] *Réda and Delahaye-Duriez (2022). Int. Conf. on Comp. Meth. in Sys. Bio. (CMSB), pp. 89–121*

[30] *Réda and Kaufmann and Delahaye-Duriez (2021). AISTATS. vol. 130. PMLR*

[31] *Réda and Tirinzoni and Degenne (2021). Advances in Neural Information Processing Systems 34*

[32] *Réda and Vakili and Kaufmann (Under review)*

[33] *Bokobza and Réda and Guenoun et al. (In preparation)*

[34] *Réda and Kaufmann and Delahaye-Duriez (2020). Comp. & struct. biotech. journ. (CSBJ) 18, pp. 241–252*

Conclusion | Future work

Towards the improvement and extension of the approach

Network & drug efficacy

- Integrating non-coding elements

e.g., enhancers, promoters, miRNAs

- Dose/exposure time effects

Bandit algorithms

- Privacy-preserving collaborative clinical trials

Acknowledgements


Université Paris Cité


Inserm


MINISTÈRE DE L'ENSEIGNEMENT SUPÉRIEUR ET DE LA RECHERCHE

Liberté
Égalité
Fraternité

anr
agence nationale de la recherche
AU SERVICE DE LA SCIENCE

Credits for images under CC license

p. 2, 7	DataBase Center for Life Science (DBCLS)
p. 2	User Titimaster (Wikipedia)
p. 2, 3, 5, 6, 12, 17, 18, 20, 22, 31, 35	User Darkdadaah (Wikipedia)
p. 18, 20, 22	Vincent Le Moign
p. 7	DataBase Center for Life Science (DBCLS)
p. 7	DataBase Center for Life Science (DBCLS)
p. 7	User Vossman (Wikipedia)
p. 7	User Emw (Wikipedia)

1. Gene regulatory networks
   - Robustness
   - Key genes
2. Application of signature reversion
   - Encephalopathy of prematurity (EoP)
3. Multi-armed bandit algorithms
   - Sample complexity lower and upper bounds
   - Misspecified models
   - Collaborative learning
4. Application to epilepsy
5. Future work
   - Integrating non-coding elements
   - Dose/exposure time effects
   - Privacy-preserving collaborative clinical trials

Gene regulatory networks | Robustness

	Min.	25th quantile	Median	Mean	75th quantile	Max.
# RFs	1	1	2	2.635	3	17
OTP	0.794	0.796	0.797	0.797	0.799	0.802

Regulatory function

Criterion for selecting the optimal model (in [0,1])

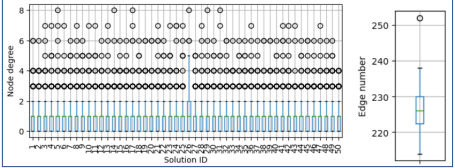

across all 50 solutions

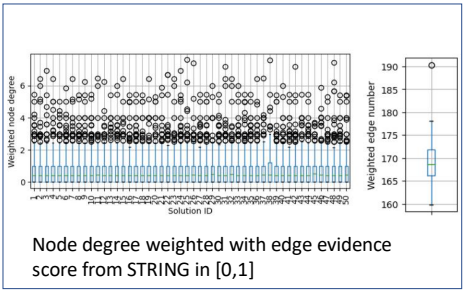

Node degree weighted with edge evidence score from STRING in [0,1]

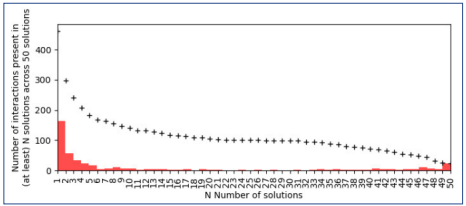

Variability of interactions across solutions

A core set (22 interactions) that are present across all solutions!

Master regulators

Genes at the top of the gene regulation hierarchy[35]

$\sim$ relate to disease onset / act on related pathways

e.g., SESTRIN3[36], CSF1R[37]in epilepsy

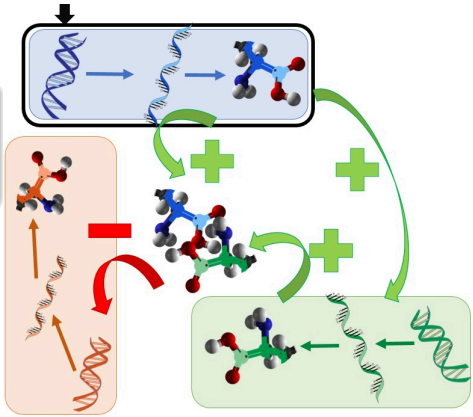

[35] *Mattick, Taft, and Faulkner (2010). Trends in genetics 26.1, pp. 21–28*

[36] *Johnson et al. (2015). Nat. comm. 6.1, pp. 1–11*

[37] *Srivastava et al. (2018). Nat. comm. 9.1, pp. 1–15*

A generic method to detect master regulators in a regulatory network?

Hub nodes

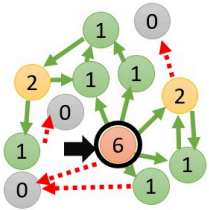

Network controllability

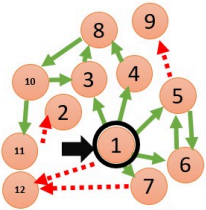

Transcription factor influence (TIF)

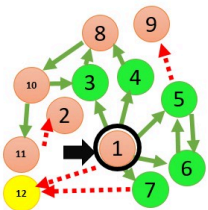

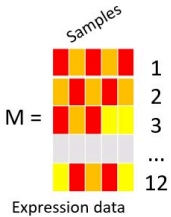

$$\frac{\partial x(t)}{\partial x} = \underbrace{\begin{bmatrix} 00\star\star\star\star\star00000 \\ 00\,0\;\;0\,0\;\;0\;\;0\,0\,0\,0\,0 \\ \cdots \\ 00\,0\;\;0\,0\;\;0\;\;0\,0\,0\,0\,0 \end{bmatrix}}_{\text{Adjacency matrix}} x(t) + \underbrace{\begin{bmatrix} \star \\ 0 \\ \cdots \\ 0 \end{bmatrix}}_{\text{Perturbation}}$$

$$TIF_g = \frac{\text{Average activation} - \text{Average inhibition}}{\text{Variance of the expected change by } g}$$

Outgoing degree | Control Centrality[38] | CoRegNet[39]

[38] *Liu, Slotine, and Barabási (2012). PLOS ONE 7.9, e44459*

[39] *Nicolle, Radvanyi, and Elati (2015). Bioinformatics 31.18, pp. 3066–3068*

Scoring of the spread of the regulatory influence of gene *g*

Compare **reachable attractors** in the absence and under perturbation of gene *g* from a diseased transcriptional context

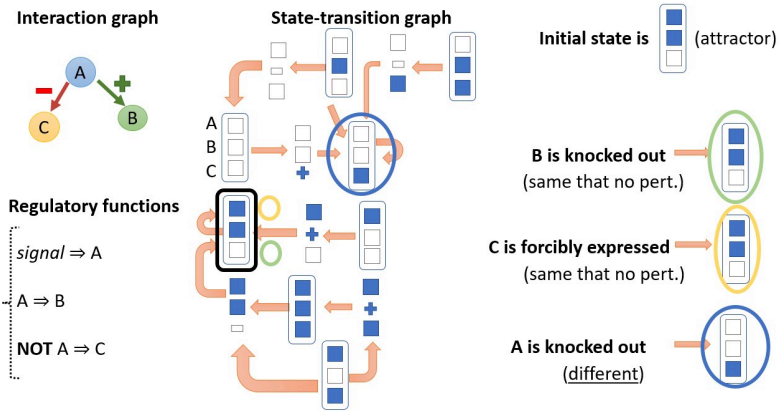

Gene (set) influence in a specific transcriptional profile p

Given a Boolean network $\mathcal{B}(\mathcal{V}, \mathcal{A}_p)$, $\underbrace{\mathcal{A}_p : S(\mathcal{V} \times \{\text{KO}, \text{FE}\}) \to S(\{0,1\}^{|\mathcal{V}|})}_{\text{reachable attractors from profile } p \text{ under gene perts.}}$,

the influence of gene set N in profile $p \in \{0,1\}^{|\mathcal{V}|}$ is

$$\boxed{I_{(\mathcal{B},p)}(N) = 1 - \max\left\{\texttt{sim}\left(a^1, a^2\right) : a^1, a^2 \in \mathcal{A}(p, \emptyset) \times \mathcal{A}(p, P(N))\right\}},$$

where $P(N) := \{(n, \text{KO } \textbf{if } p[n] \textbf{ else } \text{FE}) : n \in N\})\}$.

In particular, $I_{(\mathcal{B},p)}(N) = 0$ iff. $\mathcal{A}(p, \emptyset) \cap \mathcal{A}(p, P(N)) \neq \emptyset$

⚠ Extended to multiple profiles by taking a corrected geometric mean

Adaptation of the greedy algorithm in[40]

1: **Input:** $\mathcal{B}(\mathcal{V}, \{\mathcal{A}_p : p \in \mathcal{P}\})$ Boolean network where $\mathcal{P}$ set of diseased transcriptional profiles, size k

2: Initialize $\mathcal{N} = \emptyset$, $i = 0$

3: **repeat**

4: $\quad i \leftarrow i + 1$

5: $\quad N_i \leftarrow \underset{n \in \mathcal{V} \setminus \mathcal{N}}{\arg\max}\, \mathsf{I}_{(\mathcal{B},\mathcal{P})}(\mathcal{N} \cup \{n\})$ **and** $\mathcal{N} \leftarrow \mathcal{N} \cup N_i$

6: **until** $i = k$ or $\underbrace{\max_{n \in \mathcal{V} \setminus \mathcal{N}} \mathsf{I}_{(\mathcal{B},\mathcal{P})}(\mathcal{N} \cup \{n\}) \leq \mathsf{I}_{(\mathcal{B},\mathcal{P})}(\mathcal{N})}_{\text{do not add redundant or detrimental elements}}$

7: **Output:** $\mathcal{N}$

[40] *Kempe, Kleinberg, and Tardos (2003). Proc. of the 9^{th} ACM SIGKDD ICKDD, pp. 137–146.*

Gene regulatory networks | Key genes 6/7

Choose $k = 1$, Temporal Lobe Epilepsy (TLE) profiles[41] $\sim$ refractory[42]

Gene location-based measures

Gene pathogenicity measures

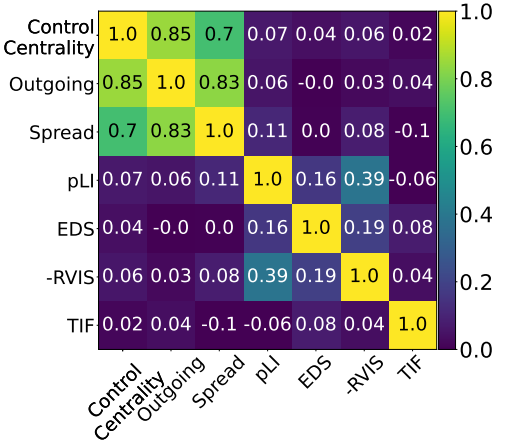

[41] *Mirza et al. (2017). Human molecular genetics 26.9, pp. 1759–1769.*

[42] *Han et al. (2014). Brain research bulletin 109, pp. 13–21.*

Gene regulatory networks | Key genes 7/7

List of prioritized genes (spread > 0.01)

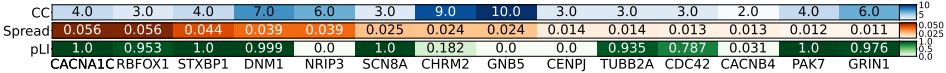

Over-Representation Analysis (ORA) compare significant functional gene enrichments in prioritized genes VS M30 module (FDR $< 20\%$)

DisGeNet[43] functional families

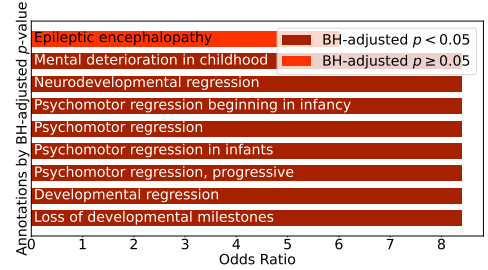

GLAD4U[44] functional families

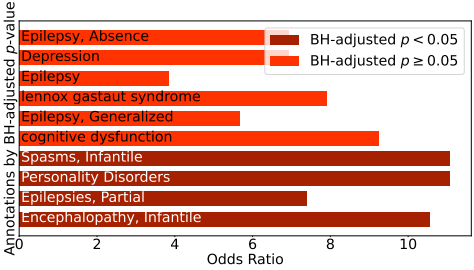

[43] *Piñero et al. (2020). Nucleic acids research 48.D1, pp. D845–D855*

[44] *Jourquin et al. (2012). BMC genomics 13.8, pp. 1–12*

Application of signature reversion | EoP 1/3

Encephalopathy of prematurity in rats

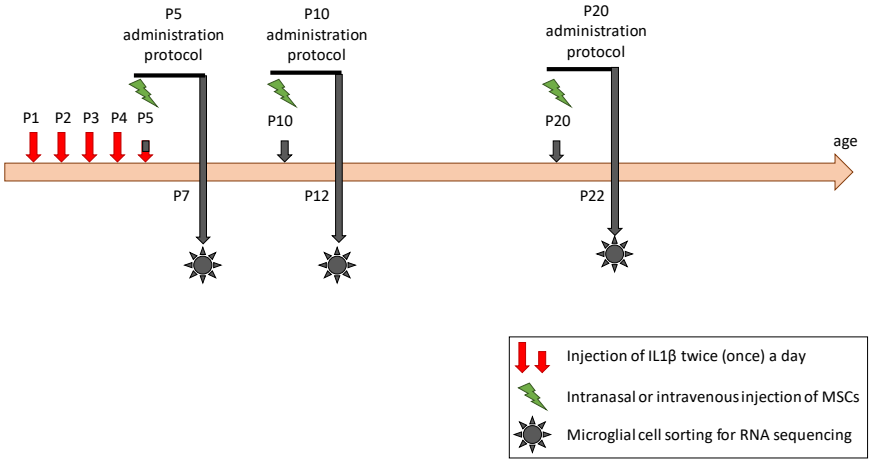

Application of signature reversion | EoP 2/3

Selection of an optimal treatment protocol by stem cell injection

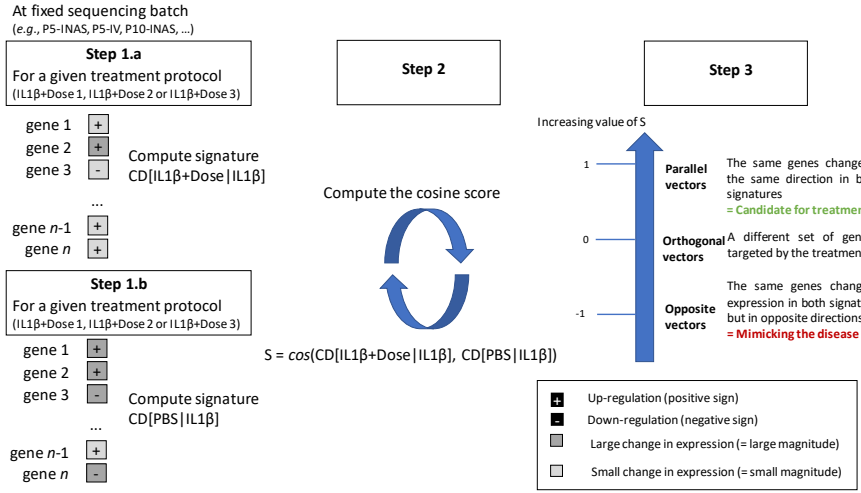

Application of signature reversion | EoP 3/3

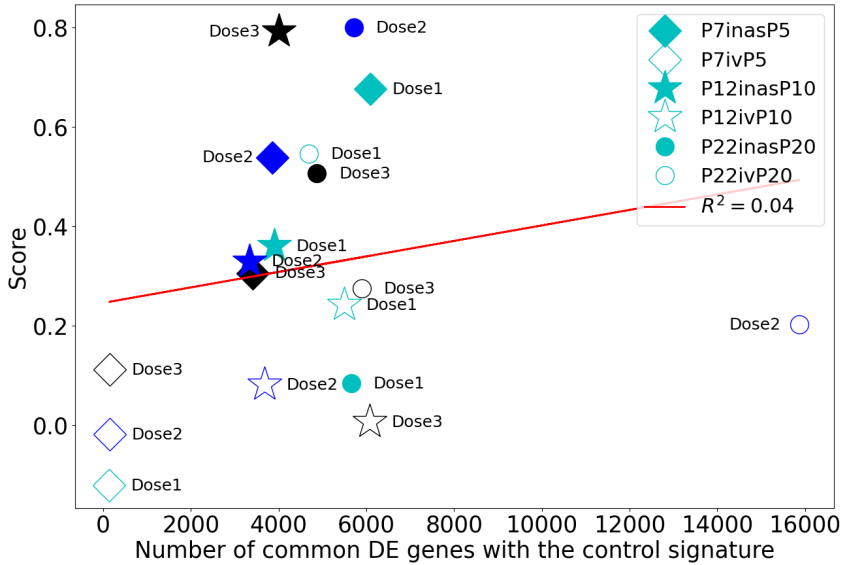

Sample complexity | Lower bound on τ

A δ-*correct* algo samples *at least* an expected $\mathcal{C}^\star(f)^{-1}\mathcal{O}(\ln(\delta^{-1}))$ arms

Lower bound on sample complexity for pure exploration[45]

For any $\delta \leq 0.5$, any 1-Gaussian bandit algorithm $\mathfrak{A}$ δ-correct on $f \in \mathcal{M}$,

$$\mathbb{E}_{\mathfrak{A},f}[\tau] \geq \left(\underbrace{\sup_{\omega \in \triangle_K} \underbrace{\inf_{\substack{g \in \mathcal{M} \\ \mathcal{S}^\star(f) \cap \mathcal{S}^\star(g) = \emptyset}}}_{g \in \mathrm{Alt}(f)} \sum_{k \in [K]} \frac{\omega_k}{2}(f(x_k) - g(x_k))^2}_{= \mathcal{C}^\star(f)} \right)^{-1} \log\left(\frac{1}{2.4\delta}\right) .$$

[45] *Garivier and Kaufmann (2016). COLT. PMLR, pp. 998–1027*

Linear bandit models | Sample complexity

Algorithm	Upper bound
Unstructured	$\inf\{u \in \mathbb{R}^{*+} \mid u > 1 + \mathcal{H}^{\mathsf{uns}}(f)(\mathcal{T}_\delta(u))^2\}$
LUCB	$\mathcal{H}^{\mathsf{uns}}(f) := 2\sum_a \max\left(\varepsilon/2, \Delta_a\right)^{-2}$
UGapE	$\mathcal{H}^{\mathsf{uns}}(f) := 2\sum_a \max\left(\varepsilon, (\varepsilon+\Delta_a)/2\right)^{-2}$
GIFA*	$\inf\left\{u \in \mathbb{R}^{*+} \mid u > 1 + \mathcal{H}^{\mathsf{lin}}(f)(\mathcal{T}_\delta(u))^2\right\}$
Largest var.	$\mathcal{H}^{\mathsf{lin}}(f) := 4\sigma^2 \sum_a \max\left(\varepsilon, (\varepsilon+\Delta_a)/3\right)^{-2}$
Optimized	$\mathcal{H}^{\mathsf{lin}}(f) := \sigma^2 \sum_a \max_{i,j\in[K]^2} \lvert\omega_a^\star(i,j)\rvert \left(\max\left(\varepsilon, (\varepsilon+\Delta_i)/3, (\varepsilon+\Delta_j)/3\right)\right)^{-2}$
Greedy*	$\mathcal{H}^{\mathsf{lin}}(f) := \sigma^2 \sum_a \max_{i,j\in[K]^2} \frac{\lVert X_i - X_j\rVert_2^2}{\lVert X_a\rVert_2^2} \left(\max\left(\varepsilon, (\varepsilon+\Delta_i)/3, (\varepsilon+\Delta_j)/3\right)\right)^{-2}$

Misspecified models | MisLid 1/2

Main ingredients:
- User-provided $\hat{\psi}$
- Online learner $\mathcal{L}$
- Stopping threshold $\{\beta_{t,\delta}\}_t \approx \log(1/\delta) + \min\{K \log t, d \log t + \psi^2 t\}$

⋆ Careful design of $\beta_{t,\delta}$ for unstructured & linear model
⋆ Sampling instead of tracking: reduce terms in K in complexity
⋆ Analysis holds for any no-regret learner

Misspecified Linear IDentification (MisLID)

Compute emp. mean $\widehat{\mu}$, $\tilde{\mu} := \Pi_{\mathcal{M}_{\hat{\psi}}}(\widehat{\mu})$
`# STOPPING`
while $\inf_{\lambda \in \mathsf{Alt}(\tilde{\mu}) \cap \mathcal{M}_{\hat{\psi}}} \|\tilde{\mu} - \lambda\|^2_{N_\cdot(t-1)} \leq 2\beta_{t-1,\delta}$ **do**
  Obtain allocation $\omega(t) \in \triangle_K$ from $\mathcal{L}$
  Compute $\lambda_t \leftarrow \arg\min_{\lambda \in \mathsf{Alt}(\tilde{\mu}) \cap \mathcal{M}_{\hat{\psi}}} \|\tilde{\mu} - \lambda\|^2_{\omega(t)}$
  Update optimistically $\mathcal{L}$
  Pull $k_t \sim \omega(t)$ `# SAMPLING`
  Update $\widehat{\mu}_t$ and projection $\tilde{\mu}$, $t \leftarrow t+1$
end while
return $\arg\max_{k}^{[m]} \widehat{\mu}_k$ `# DECISION`

Misspecified models | MisLid 2/2

Sample complexity upper bound on $f \in \mathcal{M}_\psi$ (if $\hat{\psi} = \psi$)

$$\mathbb{E}_f[\tau] \leq T_\delta + 2 \ (\star) \ ,$$

where

$$T_\delta \leq \mathcal{C}^\star(f)^{-1} \left[\log\left(\frac{1}{\delta}\right) + H(T_\delta; \psi, d, K) o\left(\log\left(\frac{1}{\delta}\right)\right) \right] .$$

$$\psi \approx 0 \implies H(T_\delta; \psi, d, K) = \psi^2 T_\delta + \log K \sqrt{T_\delta} + d\sqrt{T_\delta} \log T_\delta + \sqrt{d \log T_\delta} \ .$$

$$\psi \gg 0 \implies H(T_\delta; \psi, d, K) = \log K \sqrt{T_\delta} + \sqrt{K T_\delta} \log T_\delta + \sqrt{K \log T_\delta} \ .$$

- Asymptotic optimality[46]

$$\delta\text{-correct} \cap (\star) \implies \lim\inf_{\delta \to 0} \frac{\mathbb{E}_f[\tau]}{\log(1/\delta)} = \mathcal{C}^\star(f)^{-1}$$

- No polynomial dependence in K when $\psi = 0$.

[46] *Rémy Degenne and Wouter M Koolen (2019). arXiv preprint arXiv:1902.03475.*

Weighted Collaborative Model[47]

P patient populations: sampling arm k in population p at time t

$$Y_t = \mathbb{E}_{\{Y \sim \nu_{(k,p)}\}}[Y] + \phi_t \ , \ \phi_t \sim \mathcal{N}(0,1)$$

$W \sim$ similarity matrix on populations, *mixed* reward for arm k in p is

$$\mu'_{k,p} := \sum_{n \in [P]} w_{n,p} \mathbb{E}_{\{Y \sim \nu_{(k,n)}\}}[Y]$$

Goal: find the top-m arms *w.r.t. mixed* rewards

[47] *Réda and Vakili and Kaufmann (Under review)*

Oracle problem

$$\mathcal{P}(\mu') := \arg \min_{t \in (\mathbb{R}^+)^{K \times P}} \sum_{k,p \in [K] \times [P]} t_{k,p} \ s.t.$$

$$\forall p \forall (k,l) \in \mathcal{S}^\star(\mu'_{\cdot,p}) \times \mathcal{S}^\star(\mu'_{\cdot,p})^c, \ \sum_{n \in [P]} \left(\frac{w_{n,p}^2}{t_{k,n}} + \frac{w_{n,p}^2}{t_{l,n}} \right) \leq \frac{(\mu'_{k,p} - \mu'_{l,p})^2}{2}$$

$\mathcal{N}^\star(\mu') := \sum_{k,p} t_{k,p}$ where $t \in \mathcal{P}(\mu')$.

Lower bound on the sample complexity

For any δ-correct 1-Gaussian algorithm $\mathfrak{A}$ where $\delta \leq 0.5$ and $\forall p, W_{p,p} \neq 0$

$$\mathbb{E}_{\mathfrak{A},\mu'}[\tau] \geq \mathcal{N}^\star(\mu') \log \left(\frac{1}{2.4\delta} \right)$$

Why considering *mixed reward* top-m arms instead of local ones?

$\hookrightarrow$ Be more sample-efficient (at the cost of some suboptimality)

ex. Case $P = K = 2$, $m = 1$

$$\mu = \overbrace{\begin{bmatrix} 1 & 0.5 \\ 0 & 0.1 \end{bmatrix}}^{\text{pop 1} \quad \text{pop 2}} \begin{matrix} \} \text{arm 1} \\ \} \text{arm 2} \end{matrix}$$

$$W = \overbrace{\begin{bmatrix} 0.53 & 0.47 \\ 0.47 & 0.53 \end{bmatrix}}^{\text{pop 1 pop 2}} \begin{matrix} \} \text{pop 1} \\ \} \text{pop 2} \end{matrix}$$

For a similarity of 0.90 between populations
$\Rightarrow \mathcal{N}^\star(\mu W) \approx 18$

Solving 2 independent bandit problems
$\Rightarrow \mathcal{N}^\star(\mu \mathsf{Id}_P) \approx 58$

Epilepsy problem ($P = 3$, $K = 21$, $m = 1$)
$\mathcal{N}^\star(\mu W) \approx 2,770$ and $\mathcal{N}^\star(\mu \mathsf{Id}_P) \approx 5,193$ (+54%)

Relaxed oracle

$$\widehat{\mathcal{P}}(\Delta) := \arg\min_{t} \sum_{k,p} t_{k,p}$$

$$s.t., \quad \forall p \, \forall k, \sum_{n} \frac{w_{n,p}^2}{t_{k,n}} \leq \frac{(\Delta'_{k,p})^2}{2}$$

$\Delta'_{k,p}$ characteristic gap for k in $\mu'_{\cdot,p}$

⚠ $\Omega_{k,m}(r,\delta)$ depends on the number of samples from (k,p) up to round r

WCPE-Top-m

$\widehat{\Delta}_{k,p}(0) = 1$, $B_p(0) = \{1,\dots,K\}$

repeat

- $\tau(r) \leftarrow \widehat{\mathcal{P}}(\sqrt{2}\widehat{\Delta}(r))$
- Sample $d_{k,p}(r)$ times from (k,p) *s.t.* $\Omega_{k,p}(r,\delta) \leq \widehat{\Delta}_{k,p}(r)$
- $L_{k,p}(r), U_{k,p}(r) \leftarrow \{\widehat{\mu}'_{k,m} \pm \Omega_{k,p}(r;\delta)\}$
- Remove all arms k from $B_p(r)$ *s.t.* $U_{k,p}(r) < \max^{m}_{j \in B_p(r)} L_{j,p}(r)$
- $r \leftarrow r+1$
- $\widehat{\Delta}_{k,p}(r) \leftarrow \frac{\widehat{\Delta}_{k,p}(r-1)}{2}$ if $k \in B_p(r)$

until $|B_p(r)| = m$ for any p

return $B_p(r)$ for any p

A benchmark for the case $m = 1$ and $W = \alpha \mathrm{Id}_P + \frac{1-\alpha}{P}\mathbb{1}$ with $\alpha \in [0, 1]$

PF-UCB-BAI (inspired by **PF-UCB** [48])

$f(r)$: sampling effort in round r, $B_p(r) = \{1, \ldots, K\}$ for any p

repeat

Sample $d_{k,p}(r)$ times from (k, p) *s.t.*

$d_{k,p}(r) = \lceil (1-\alpha) f(r) \rceil + \lceil \alpha P f(r) \rceil \mathbb{1}(k \in B_p(r))$

$L_{k,p}(r), U_{k,p}(r) \leftarrow \{\widehat{\mu}'_{k,m}(r) \pm \Omega(r; \delta)\}$

Remove all arms k from $B_p(r)$ *s.t.*

$U_{k,p}(r) < \max_{j \in B_p(r)} L_{j,p}(r)$

$r \leftarrow r + 1$

until $|B_p(r)| = 1$ for any p, and **return** $\hat{k}_p \in B_p(r)$ for any p

[48] *Shi, Shen, and Yang (2021). International Conference on Artificial Intelligence and Statistics. PMLR, pp. 2917–2925*

Application to epilepsy | Choice of ψ in MisLid

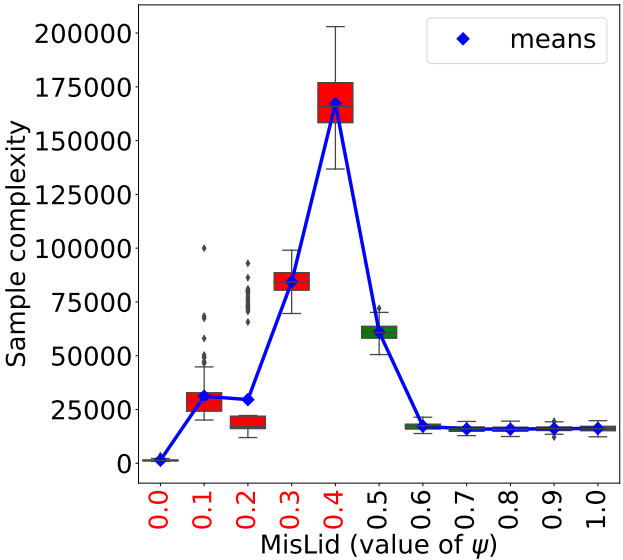

Future work | Integrating non-coding elements

Pre-transcription

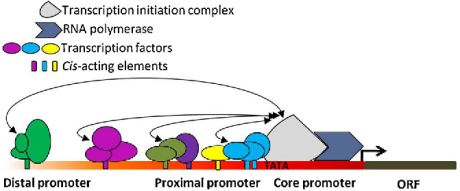

Cis-regulatory modules[49]

Post-transcription

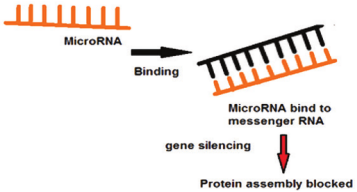

miRNAs[50]

Further constraint the dynamics of the network

Add to the identification pipeline a biologically meaningful decomposition of the set of vertices [51]

[49] *Howard and Davidson (2004). Developmental biology 271.1, pp. 109–118; Hernandez-Garcia and Finer (2014). Plant Sci 217, pp. 109–119*

[50] *Saeidimehr et al. (2016). Cell J. 18.2, p. 117*

[51] *Bernot, Comet, and Khalis (2008). Eu. sim. and mod. conf, pp. 423–432; Réda and Wilczynski (2020). J. of Th. Bio. 486, p. 110091*

Future work | Dose/exposure time effects

Drug efficacy depending on cell-line, *dose & exposure time*

Refining drug signature based on dose and time

Apply the same approach as to cell-line specific signatures

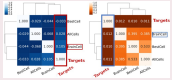

... or change the structure of the bandit model?

One can exploit the dose-response curves available for drugs[52]

[52] *Güvenç Paltun, Kaski, and Mamitsuka (2021). Briefings in Bioinformatics 22.6, bbab293.*

Future work | Privacy-preserving collaborative CTs

Privacy is key in federated learning & data generation
(*e.g.*, security[53], healthcare[54])

Intuition behind ε-differential privacy

On two datasets differing by *a single* sample, the probabilities of any outcome from a ε-differentially private algorithm on those datasets are equal up to a multiplicative constant $\exp(\varepsilon)$

Modify the WCPE-Top-*m* algorithm

Design a differentially private algorithm solving weighted collaborative Top-*m* identification

[53] *Dubey and Pentland (2020). Advances in Neural Information Processing Systems 33, pp. 6003–6014.*

[54] *Jordon, Yoon, and Van Der Schaar (2018). International conference on learning representations.*