Mitigating Biases in Decision-Making Systems: a Control Systems Perspective

Giulia De Pasquale

ETH Zürich

December 11, 2024

Automated Decision Making (ADM)

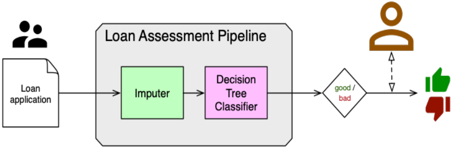

✓ High scalability

✗ Exhacerbate existing biases and even introduce new ones

Applications

employment
health
education
law
. . .

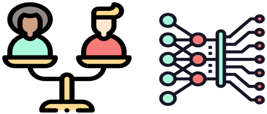

Algorithmic fairness

Enforce group fairness metrics to mitigate biases

✗ solutions are designed for **stationary** systems

Automated Decision Making (ADM)

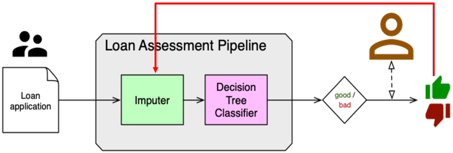

Applications
employment
health
education
law
. . .

✓ High scalability

✗ Exhacerbate existing biases and even introduce new ones

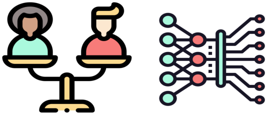

Algorithmic fairness

💡 Enforce group fairness metrics to mitigate biases

✗ solutions are designed for **stationary** systems

[1]

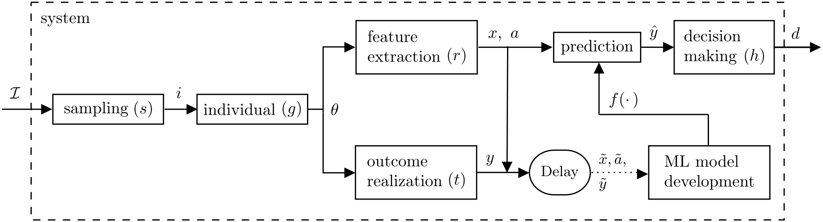

The ML-based decision making pipeline as an open loop system

[1] "A classification of feedback loops and their relation to biases in automated decision-making systems", J. Baumann, N. Pagan, E. Elokda, GDP, S. Bolognani, A. Hannak, Conference on Equity and Access in Algorithms, Mechanisms, and Optimization

1

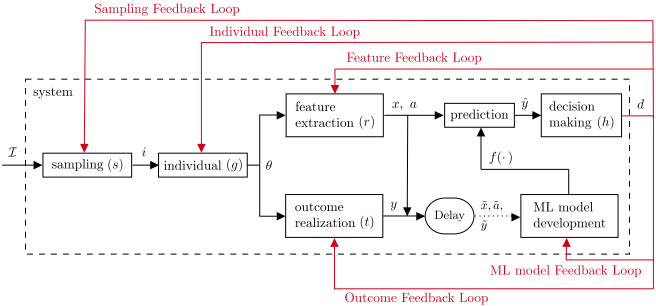

The ML-based decision making pipeline as a closed loop system

[1] "A classification of feedback loops and their relation to biases in automated decision-making systems", J. Baumann, N. Pagan, E. Elokda, GDP, S. Bolognani, A. Hannak, Conference on Equity and Access in Algorithms, Mechanisms, and Optimization

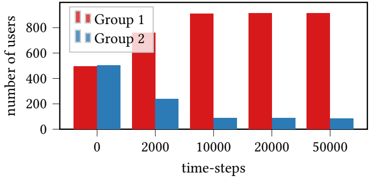

Sampling FL: Representation bias

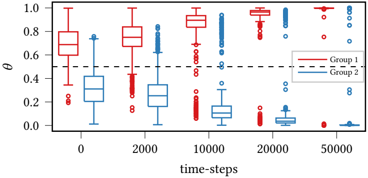

Individual FL: Historical bias

The available data is not representative of the population: **the ML model does not generalize well for the disadvantaged group**, e.g. Amazon's Alexa.

Users with high initial interests get recommended the item: θ increases over time. **Decisions change individual properties**, leads to polarization of interests.

A Solution to Representation Bias[2]

[2] "Fairness in Social Influence Maximization via Optimal Transport", S. Chowshary, GDP*, N. Lanzetti*, A. Stoica, F. Dörfler, NeurIPS 2024

Social Influence Maximization

Suppose you want to sell a product, or make an information spread as much as possible in a social network:

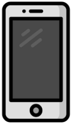

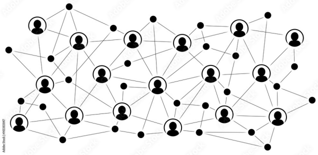

Social Influence Maximization (SIM) is the problem of how to strategically selects seeds that spread information throughout a network in order to **maximize the outreach**.

Social Influence Maximization

Suppose you want to sell a product, or make an information spread as much as possible in a social network:

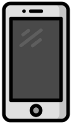

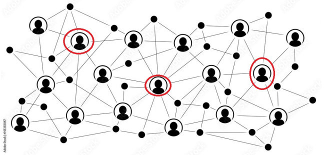

Social Influence Maximization (SIM) is the problem of how to strategically selects seeds that spread information throughout a network in order to **maximize the outreach**.

Suppose you want to spread the news about an open position as Assistant Professor in Control Engineering:

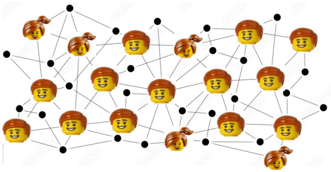

Fairness in SIM: solve SIM by ensuring **balanced outreach** among different communities, e.g. demographic groups.

Spreading mechanism: Independent cascade model

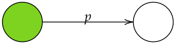

Fairness in SIM

Suppose you want to spread the news about an open position as Assistant Professor in Control Engineering:

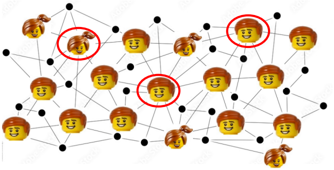

Fairness in SIM: solve SIM by ensuring **balanced outreach** among different communities, e.g. demographic groups.

Spreading mechanism: Independent cascade model

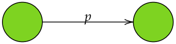

Fairness Metrics in SIM

Given the groups $C_1, \ldots, C_m$, a configuration is said to be

Equal, if the SIM algorithm chooses a seed set S such that

$$\frac{\mathbb{E}[|v \in S | v \in C_i|]}{|C_i|} = \frac{\mathbb{E}[|v \in S | v \in C_j|]}{|C_j|} \quad \forall i, j.$$

Equitable, if the SIM algorithm chooses a seed set S such that

$$\frac{\mathbb{E}[|v \text{ reached} | v \in C_i|]}{|C_i|} = \frac{\mathbb{E}[|v \text{ reached} | v \in C_j|]}{|C_j|} \quad \forall i, j.$$

Max-Min Fair, if the SIM algorithm chooses a seed set S such that

$$\min_{i \in [m]} \frac{\mathbb{E}[|v \text{ reached} | v \in C_i|]}{|C_i|}$$

is maximized.

Given the groups $C_1, \ldots, C_m$, a configuration is said to be

Equal, if the SIM algorithm chooses a seed set S such that

$$\frac{\mathbb{E}[|v \in S | v \in C_i|]}{|C_i|} = \frac{\mathbb{E}[|v \in S | v \in C_j|]}{|C_j|} \quad \forall i, j.$$

Equitable, if the S ch that

$\mathbb{E}[|v$ rea] $\forall i, j.$

Max-Min Fair, if S such that

is maximized

What's wrong with the Expectation?

Consider the outcome: "In 50% if the cases, no one in group 1 gets the information and everyone in group 2 does, and in the other 50 % it is the opposite."

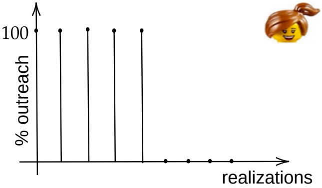

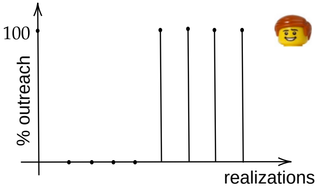

$$\frac{\mathbb{E}[|v \text{ reached}|v \in C_1|]}{|C_1|} = \frac{\mathbb{E}[|v \text{ reached}|v \in C_2|]}{|C_2|}$$

The outcome is classified as **equitable**, however it is **highly unfair**.
Note: this also happens in experimental settings!

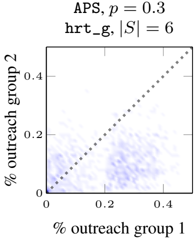

We want to answer questions such as as:

i) When group 1 receives the information, will group 2 also receive it?
ii) Even if the two groups have the same marginal outreach probability distributions, will the final configurations always be **fair**?

Motivating Example

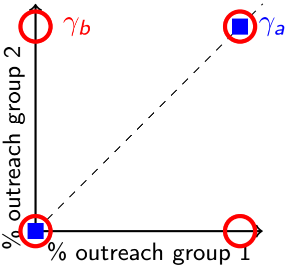

Figure: Illustration of the (γ_a,γ_b) example.

Marginals: $\mu_i = \frac{1}{2}\delta_0 + \frac{1}{2}\delta_1, i \in \{1, 2\}$
Distributions:

$$\gamma_a = 0.5 \cdot \delta_{(0,0)} + 0.5 \cdot \delta_{(1,1)}, \qquad \gamma_b = 0.25 \cdot \delta_{(0,0)} + 0.25 \cdot \delta_{(1,1)} + 0.25 \cdot \delta_{(0,1)} + 0.25 \cdot \delta_{(1,0)}.$$

Use the **joint** outreach probability distribution to capture the correlation between the two groups!

A New Fairness Metric

Quantify fairness by computing the distance of the probability distribution γ from an ideal reference distribution γ^* along the diagonal.

Optimal Transport Problem: quantifies the minimum transportation cost to morph γ into γ^* when transporting a unit of mass from (x_1, x_2) to (y_1, y_2) costs $c((x_1, x_2), (y_1, y_2))$.

$$W_c(\gamma, \gamma^*) = \min_{\pi \in \Pi(\gamma, \gamma^*)} \mathbb{E}_{(x_1,x_2),(y_1,y_2)\sim\pi}, [c((x_1, x_2), (y_1, y_2))]$$

Ingredients:

i) transportation cost;

ii) reference distribution.

Transportation Cost:

- moving mass **along** the diagonal costs 0, as it does not affect fairness
- moving mass **orthogonally** towards the diagonal comes at a price. We quantify the price as the **Euclidean distance**.

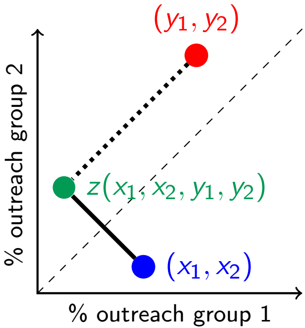

$$c((x_1, x_2), (y_1, y_2)) = \|z(x_1, x_2, y_1, y_2) - (x_1, x_2)\| = \frac{\sqrt{2}}{2}|(x_2 - x_1) - (y_2 - y_1)|,$$

Definition (Mutual Fairness)

Given a network with communities $(C_i)_{i\in[2]}$, a SIM algorithm is said to be *mutually fair* if the algorithm propagation is such that it maximizes

$$\text{FAIRNESS}(\gamma) := 1 - \sqrt{2} W_c(\gamma, \gamma^*),$$

$W_c(\gamma,\gamma^*) = \min_{\pi\in\Pi(\gamma,\gamma^*)} \mathbb{E}_{(x_1,x_2),(y_1,y_2)\sim\gamma}, [c((x_1,x_2),(y_1,y_2))]$ and $\gamma^* = \delta_{(1,1)}$.

Observations:

- min $\text{FAIRNESS}(\gamma) = 0$; argmin $= \gamma = \delta_{(0,1)}$;
- max $\text{FAIRNESS}(\gamma) = 1$; argmax $= \gamma^*$.
- since γ^* is a delta distribution, we can solve the OT problem in closed form and $\text{FAIRNESS}(\gamma) = 1 - \frac{1}{N}\sum_{i=1}^{N} |x_{1,i} - x_{2,i}|$

Back to the Motivating Example

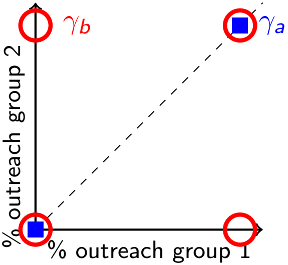

Figure: Illustration of the (γ_a,γ_b) example.

$\text{FAIRNESS}(\gamma_a) = 1$
$\text{FAIRNESS}(\gamma_b) = 0.5.$

Are the outcomes fair?

Joint outreach probability distribution for different real datasets, each with a chosen demographic partitioning the population in two groups.

Four qualitative outcomes:

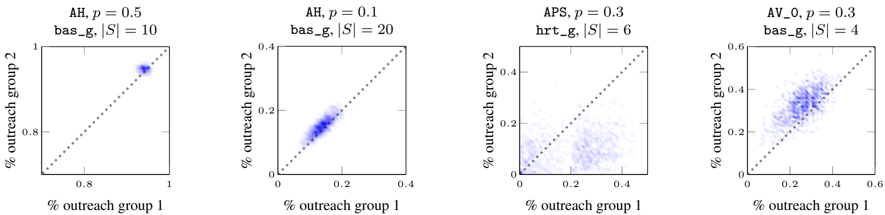

Trading-off Fairness and Efficiency

For both $\gamma = \delta_{(0,0)}$ and $\gamma^* = \delta_{(1,1)}$ the fairness score is maximal:
We need a fairness-efficiency trade-off!
We can define the transportation cost as a weighted sum:

$$c_\beta((x_1, x_2), (y_1, y_2)) =$$
$$\beta \| z(x_1, x_2, y_1, y_2) - (x_1, x_2) \| + (1-\beta) \| z(x_1, x_2, y_1, y_2) - (y_1, y_2) \| =$$
$$\beta \frac{\sqrt{2}}{2} |(x_2 - x_1) - (y_2 - y_1)| + (1-\beta) \frac{\sqrt{2}}{2} |(x_1 + x_2) - (y_1 + y_2)|.$$

Heatmap of c_β:

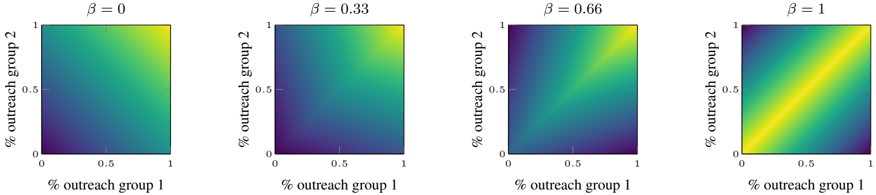

Definition (β-Fairness)

Consider a network with groups C_1, C_2, a SIM algorithm is said to be β-*fair* if the algorithm propagation is such that it maximizes

$$\beta - \text{FAIRNESS}(\gamma) := 1 - \frac{\sqrt{2}}{\max\{1, 2-2\beta\}} W_{c_\beta}(\gamma, \gamma^*),$$

The OT problem can be solved in closed form

$$\beta - \text{FAIRNESS}(\gamma) = \mathbb{E}_{(x_1,x_2)\sim\gamma}\left[1 - \frac{\beta|x_1 - x_2| + (1-\beta)|x_1 + x_2 - 2|}{\max\{1, 2-2\beta\}}\right]$$

In particular, for $\beta = 1$, we recover the mutual fairness FAIRNESS(γ) and for $\beta = 0$ we obtain the efficiency metric $\mathbb{E}_{(x_1,x_2)\sim\gamma}[1 - \frac{x_1+x_2-2}{2}]$.

Fairness-promoting Seed Selection Algorithm

Algorithm 1 Stochastic Seedset Selection Descent

Input: Social Graph $G(V_G, E_G)$, initial seed set S_0, β fairness weight, ϵ-tolerance
Output: Optimal seedset S^*

$\mathcal{S} \leftarrow \{\}$, $S \leftarrow S_0$ ▷ initial collection of candidates, running seedset
for k iterations **do** ▷ configurable k
    $V_S \leftarrow$ nodes reachable from S via cascade, using SEEDSET_REACH routine
    $S' \leftarrow \{\}$
    for $|S|$ iterations **do** ▷ searching nearby states, $V_{S'}$, to get S' (Appendix E.3)
        $S' \leftarrow S' \cup \{v\} \mid v \sim V_S$
        $V_{S'} \leftarrow$ nodes reachable from S' in a fixed horizon, using SEEDSET_REACH 1
        $V_S \leftarrow V_S \setminus V_{S'}$
    $E_S \leftarrow -\text{BETA\_FAIRNESS}(S, \beta)$ ▷ expected potential energy defined on β-fairness
    $E_{S'} \leftarrow -\text{BETA\_FAIRNESS}(S', \beta)$
    $p_{\text{accept}} \leftarrow \min\{1, e^{E_S - E_{S'}}\}$ ▷ S' acceptance on energy minimization 2
    if $x \sim \mathcal{B}(p_{\text{accept}})$ **then** ▷ Metropolis sampling
        $S^+ \leftarrow S'$ ▷ get a better seedset
    else
        if $x \sim \mathcal{B}(\epsilon)$ **then** ▷ for some small constant ϵ 3
            $S^+ \leftarrow \{v_i\}_{i=1}^{|S|} \overset{|S|}{\sim} V_G$ ▷ random seedset
        else
            $S^+ \leftarrow S$ ▷ retain existing choice
    $\mathcal{S} \leftarrow \mathcal{S} \cup \{S^+\}$
    $S \leftarrow S^+$ ▷ for next iteration
$S^* \leftarrow S \in \mathcal{S} \mid \text{BETA\_FAIRNESS}(S, \beta)$ is maximum ▷ via S3D_ITERATE
return S^*

Are the outcomes more fair?

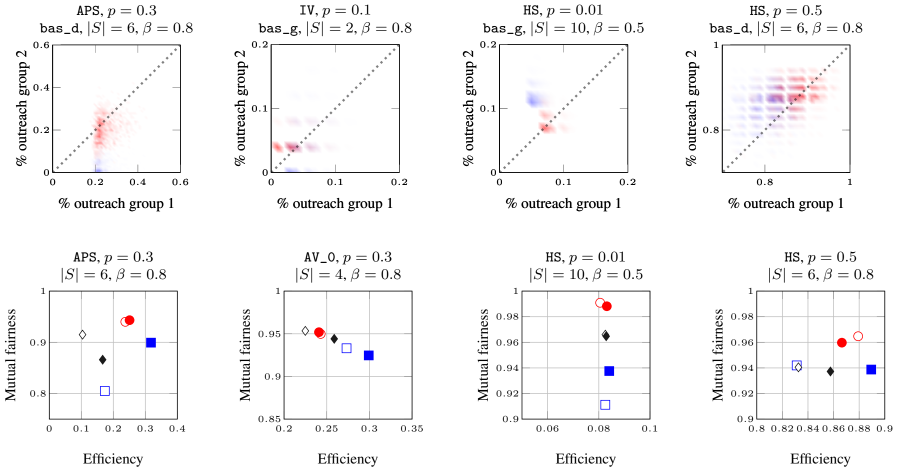

Greedy-based algorithms: ■ = bas_g, ● = S3D_g, and ♦ = hrt_g.

Degree-based algorithms: □ = bas_d, ○ = S3D_d, and ◊ = hrt_d.

Recap

- New fairness metric for SIM that captures new fairness-related aspects;
- We leverage β-fairness to design a new seed selection strategy that tradeoffs fairness and efficiency;
- We show superior fairness performance with minor decrease in efficiency.

Note: Mutual fairness is applicable whenever you have empirical distributions associated with groups.

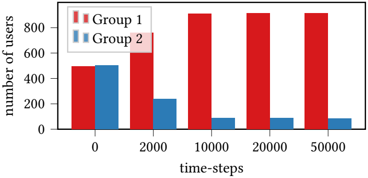

Sampling FL: Representation bias

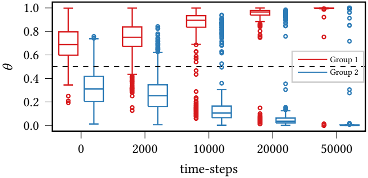

Individual FL: Historical bias

The available data is not representative of the population: **the ML model does not generalize well for the disadvantaged group**, e.g. Amazon's Alexa

Users with high initial interests get recommended the item: θ increases over time. **Decisions change individual properties**, leads to polarization of interests.

A Solution to Historical Bias[3]

[3] S. Chandrasekaran, GDP, G. Belgioioso, F. Dörfler, "Mitigating Polarization in Recommender Systems via Network-aware Feedback Optimization", submitted.

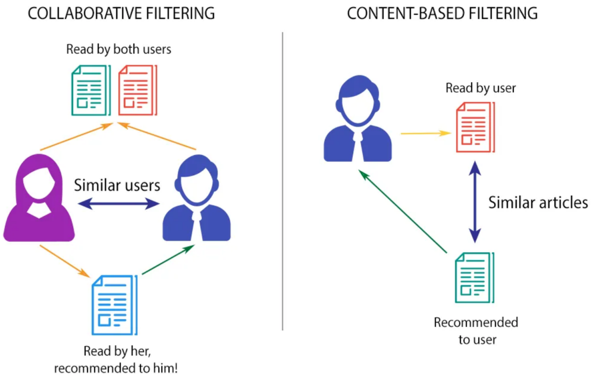
COLLABORATIVE FILTERING
Read by both users
Similar users
Read by her,
recommended to him!
CONTENT-BASED FILTERING
Read by user
Similar articles
Recommended
to user

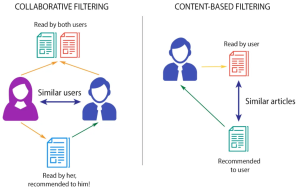

Make the feedback loop explicit, to understand

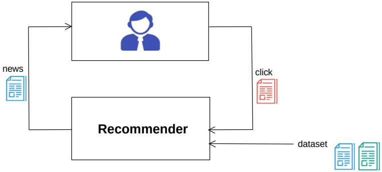

i) the impact of recommendation on users opinions;
ii) how recommender systems should depart from engagement maximization to mitigate polarization.

We leverage on **online feedback optimization** to design a RS as a dynamic feedback controller that mitigates polarization by providing user personalized content, using only **implicit feedback**.

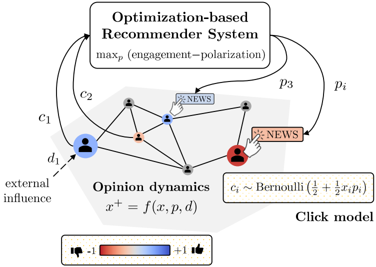

Assumption: one single topic of discussion

Assumption: The dynamics is exponentially stable and admits a unique steady-state map

$$h(p, d) = f(h(p, d), p, d)$$

with $h(p, d)$ continuously-differentiable and L-lipschitz wrt p.

Problem Formulation

$$\min_{p,x} \varphi^{\mathrm{clk}}(p,x) + \gamma \varphi^{\mathrm{pol}}(x)$$

$$\text{s.t. } x = h(p,d)$$

$$p \in [-1,1]^n$$

$$\varphi^{\mathrm{clk}} = -\sum_{i\in[n]} \mathbb{E}_{c_i \sim \mathcal{B}(g_i(x_i,p_i))}[c_i]$$

$$\varphi^{\mathrm{pol}}(x) = \|x\|^2$$

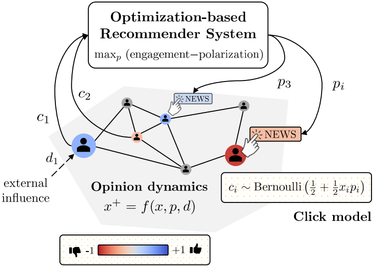

Challenges:

- only clicks are available: opinions, opinion dynamics, network topology, clicking behaviour, external influence unknown $\to$ the problem must be solved online
- non-convex problem

The recommender system only relies on clicks:

$$\frac{\#\mathrm{clk}}{\#\mathrm{news}} \approx \mathbb{E}[\mathcal{B}(g(p,x))] = g(p,x).$$

Recommender System Design

The recommender system dynamically generates recommendation via projected gradient descent

$$p^{+} = \mathrm{proj}_{[-1,1]}[p - \eta \underbrace{(\nabla_p \varphi(p, x) + \nabla_p h(p, d)^{\top} \nabla_x \varphi(p, x))}_{\nabla \varphi}]$$

$\varphi = \varphi^{\mathrm{clk}} + \varphi^{\mathrm{pol}}$.

Challenges

Evaluating $\nabla \varphi$ requires access to:

i) Online opinions x

ii) Sensitivity mapping $\nabla_p h(p, d)$

iii) Gradients $\nabla_p \varphi(p, x)$, $\nabla_x \varphi(p, x)$

None of these information is available online!

Outline

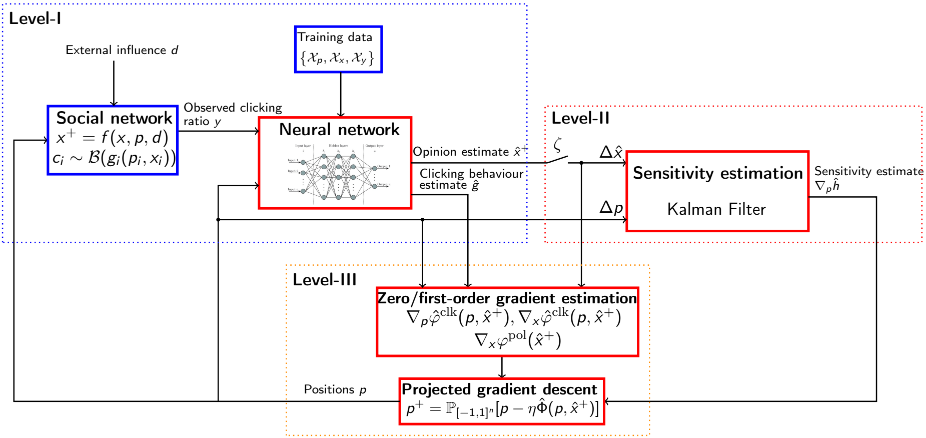
Level-I
External influence d
Training data $\{\mathcal{X}_p, \mathcal{X}_x, \mathcal{X}_y\}$
Social network
$x^+ = f(x, p, d)$
$c_i \sim \mathcal{B}(g_i(p_i, x_i))$
Observed clicking ratio y
Neural network
Opinion estimate $\hat{x}^+$
Clicking behaviour estimate $\hat{g}$
Level-II
ζ
$\Delta\hat{x}$
Δp
Sensitivity estimation
Kalman Filter
Sensitivity estimate $\nabla_p \hat{h}$
Level-III
Zero/first-order gradient estimation
$\nabla_p \hat{\varphi}^{\mathrm{clk}}(p, \hat{x}^+), \nabla_x \hat{\varphi}^{\mathrm{clk}}(p, \hat{x}^+)$
$\nabla_x \varphi^{\mathrm{pol}}(\hat{x}^+)$
Positions p
Projected gradient descent
$p^+ = \mathbb{P}_{[-1,1]^n}[p - \eta\hat{\Phi}(p, \hat{x}^+)]$

Outline

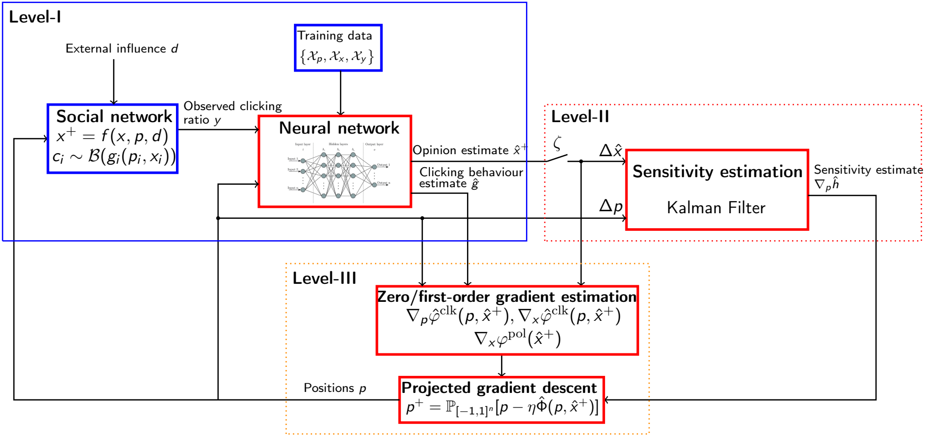
Level-I
External influence d
Training data $\{\mathcal{X}_p, \mathcal{X}_x, \mathcal{X}_y\}$
Social network
$x^+ = f(x, p, d)$
$c_i \sim \mathcal{B}(g_i(p_i, x_i))$
Observed clicking ratio y
Neural network
Opinion estimate $\hat{x}^+$
Clicking behaviour estimate $\hat{g}$
Level-II
ζ
$\Delta\hat{x}$
Δp
Sensitivity estimation
Kalman Filter
Sensitivity estimate $\nabla_p \hat{h}$
Level-III
Zero/first-order gradient estimation
$\nabla_p \hat{\varphi}^{\mathrm{clk}}(p, \hat{x}^+), \nabla_x \hat{\varphi}^{\mathrm{clk}}(p, \hat{x}^+)$
$\nabla_x \varphi^{\mathrm{pol}}(\hat{x}^+)$
Positions p
Projected gradient descent
$p^+ = \mathbb{P}_{[-1,1]^n}[p - \eta\hat{\Phi}(p, \hat{x}^+)]$

Level 1: Opinions & Clicking Behaviour Estimation

The recommender system dynamically generates recommendation via projected gradient descent

$$p^+ = \mathrm{proj}_{[-1,1]}[p - \eta \underbrace{(\nabla_p \varphi(p, x) + \nabla_p h(p, d)^\top \nabla_x \varphi(p, x))}_{\nabla \varphi}]$$

Training data collection

Repeat #training times:

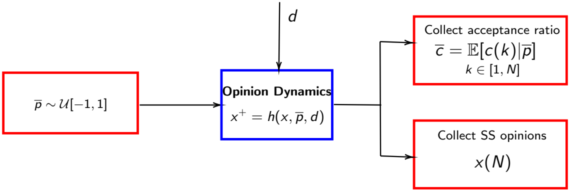

Level 1: Opinions Estimation

Assumption:

i) There exists a continuous mapping $\beta(\bar{c}, p) = x + \theta(x)$, $\|\theta(x)\| \leq \theta$

ii) $g(x, p)$ is Lipshitz and globally smooth.

There exists α s.t. $g(p, \beta(\bar{c}, p)) = \bar{c} + \nabla_x g(p, x)^\top \theta(x) + \alpha(\bar{c})$, $\|\alpha(\bar{c})\| \leq \alpha$

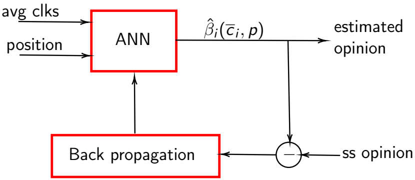

$$\begin{cases} \hat{x}_i^+ = \hat{\beta}_i(\bar{c}_i, p) \\ p \text{ via OFO} \end{cases}$$

Opinion estimation error

$$\|\overbrace{h(p, d) - \hat{\beta}}^{\epsilon_x}\| \leq \sqrt{n}(\sup_{\bar{c},p} \|\beta - \hat{\beta}\|_\infty + \text{ ANN bias})$$

Training is carried out distributedly

[3]Tabuada, Charesifard, "Universal approximation power of deep residual neural networks through the lens of control", TAC, 2023

Level 1: Clicking Behaviour Estimation

Assumption:

i) There exists a continuous mapping $\beta(\bar{c}, p) = x + \theta(x)$, $\|\theta(x)\| \leq \theta$

ii) $g(x, p)$ is Lipshitz and globally smooth.

There exists α s.t. $g(p, \beta(\bar{c}, p)) = \bar{c} + \nabla_x g(p, x)^\top \theta(x) + \alpha(\bar{c})$, $\|\alpha(\bar{c})\| \leq \alpha$

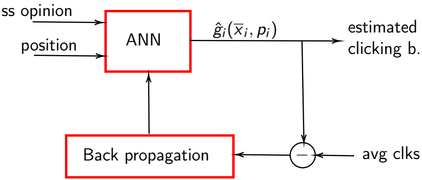

$$\begin{cases} \hat{c}_i^+ = \hat{g}_i(\hat{x}_i^+, p_i) \\ p \text{ via OFO} \end{cases}$$

clicking behaviour estimation error

$$\|\overbrace{\hat{g}(p, \hat{x}) - g(p, h(p, d))}^{\epsilon_g}\| \leq \sqrt{n}(\sup_{p,x} \|g(p, x) - \hat{g}(p, x)\|_\infty + \text{ANN bias} + f(\theta, \alpha))$$

[3] Tabuada, Charesifard, "Universal approximation power of deep residual neural networks through the lens of control", TAC, 2023

Outline

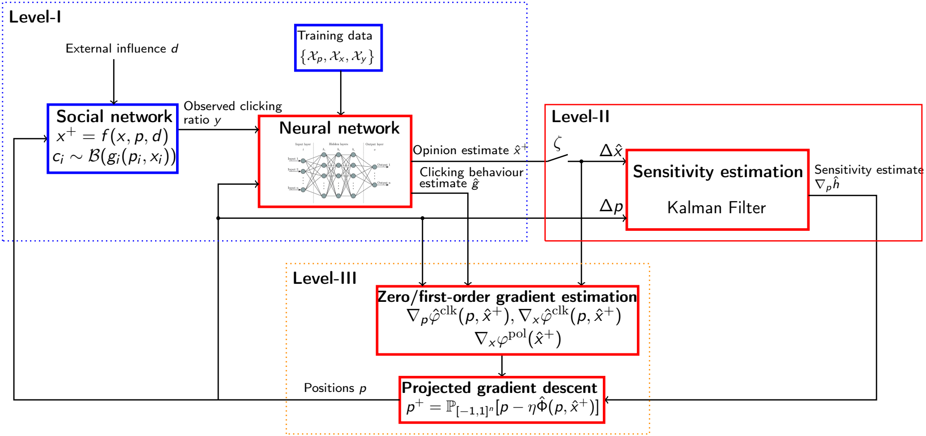
Level-I
External influence d
Training data
$\{\mathcal{X}_p, \mathcal{X}_x, \mathcal{X}_y\}$
Social network
$x^+ = f(x, p, d)$
$c_i \sim \mathcal{B}(g_i(p_i, x_i))$
Observed clicking ratio y
Neural network
Opinion estimate $\hat{x}^+$
Clicking behaviour estimate $\hat{g}$
Level-II
ζ
$\Delta\hat{x}$
Δp
Sensitivity estimation
Kalman Filter
Sensitivity estimate $\nabla_p\hat{h}$
Level-III
Zero/first-order gradient estimation
$\nabla_p\hat{\varphi}^{\mathrm{clk}}(p, \hat{x}^+), \nabla_x\hat{\varphi}^{\mathrm{clk}}(p, \hat{x}^+)$
$\nabla_x\varphi^{\mathrm{pol}}(\hat{x}^+)$
Positions p
Projected gradient descent
$p^+ = \mathbb{P}_{[-1,1]^n}[p - \eta\hat{\Phi}(p, \hat{x}^+)]$

Level 2: Online Sensitivity Learning

The recommender system dynamically generates recommendation via projected gradient descent

$$p^+ = \mathrm{proj}_{[-1,1]}[p - \eta \underbrace{(\nabla_p \varphi(p,x) + \nabla_p h(p,d)^\top \nabla_x \varphi(p,x))}_{\nabla \varphi}]$$

To estimate the sensitivity online we rely on **Kalman filter**.
Note: $\nabla_p h_{ij}(p,d) \neq 0 \rightarrow j$ and i are connected

To ensure the sensitivity estimate is accurate:
Assumption: The inputs p are persistently exciting.

Outline

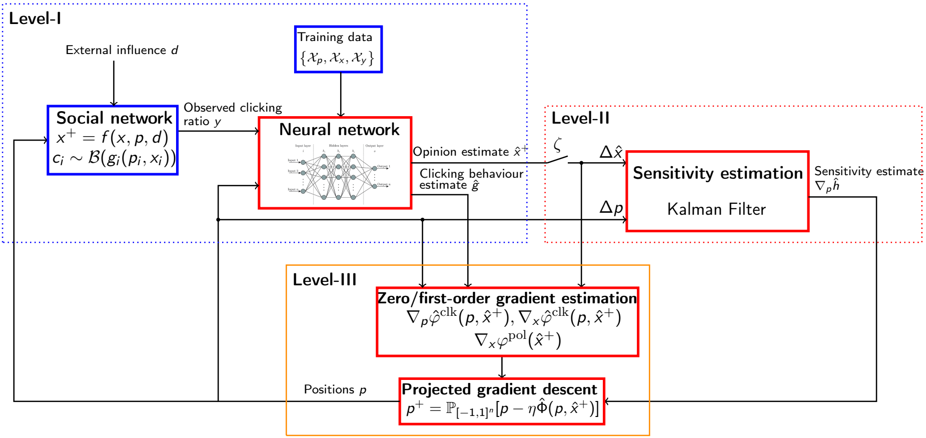
Level-I
External influence d
Training data
$\{\mathcal{X}_p, \mathcal{X}_x, \mathcal{X}_y\}$
Social network
$x^+ = f(x, p, d)$
$c_i \sim \mathcal{B}(g_i(p_i, x_i))$
Observed clicking ratio y
Neural network
Opinion estimate $\hat{x}^+$
Clicking behaviour estimate $\hat{g}$
Level-II
ζ
$\Delta\hat{x}$
Δp
Sensitivity estimation
Kalman Filter
Sensitivity estimate $\nabla_p \hat{h}$
Level-III
Zero/first-order gradient estimation
$\nabla_p \hat{\varphi}^{\mathrm{clk}}(p, \hat{x}^+), \nabla_x \hat{\varphi}^{\mathrm{clk}}(p, \hat{x}^+)$
$\nabla_x \varphi^{\mathrm{pol}}(\hat{x}^+)$
Positions p
Projected gradient descent
$p^+ = \mathbb{P}_{[-1,1]^n}[p - \eta\hat{\Phi}(p, \hat{x}^+)]$

The recommender system dynamically generates recommendation via projected gradient descent

$$p^+ = \mathrm{proj}_{[-1,1]}[p - \eta \underbrace{(\nabla_p \varphi(p,x) + \nabla_p h(p,d)^\top \nabla_x \varphi(p,x))}_{\nabla \varphi}]$$

$\varphi = \varphi^{\mathrm{clk}} + \varphi^{\mathrm{pol}}$. Estimation via forward difference method

$$\nabla_x \hat{\varphi}_i^{\mathrm{clk}}(p,x) = \frac{\hat{\varphi}^{\mathrm{clk}}(p, x + \mu e_i) - \hat{\varphi}^{\mathrm{clk}}(p,x)}{\mu}.$$

$$\nabla_p \hat{\varphi}_i^{\mathrm{clk}}(p,x) = \frac{\hat{\varphi}^{\mathrm{clk}}(p + \mu e_i, x) - \hat{\varphi}^{\mathrm{clk}}(p,x)}{\mu},$$

Gradient estimation error

Under the previous regularity assumptions on β, g

$$\|\nabla \hat{\varphi}^{\mathrm{clk}} - \nabla \varphi^{\mathrm{clk}}\| \leq \frac{1}{2} L_x \mu + 2 \frac{\sqrt{n}\epsilon_g}{\mu}; \quad \mu^* = 2n^{1/4}\sqrt{\frac{\epsilon_g}{L}}$$

Smoothing parameter μ, requires fine tooning: small, but not too much!

Recap

We now collected all the ingredients to run gradient descent for the recommender system algorithm:

$$p^{k+1} = \mathrm{proj}\Big[p^k - \eta\zeta^k\big(\nabla_p\hat{\varphi}^{\mathrm{clk}}(p^k, \hat{x}^k) + \nabla_p\hat{h}(p^k, d)^\top \nabla_x\hat{\varphi}^{\mathrm{clk}}(p^k, \hat{x}^k)\big)\Big]$$

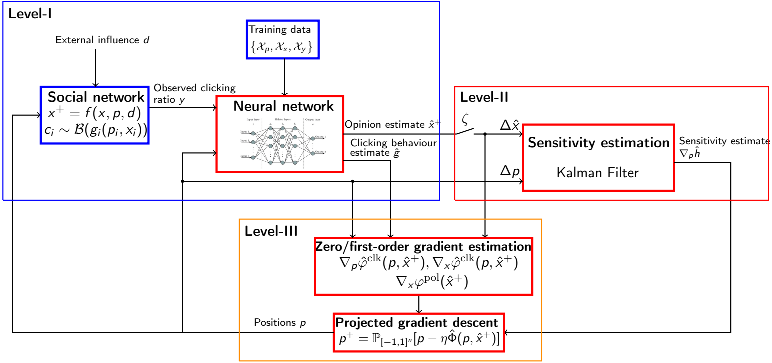

Initialization
Collect data during training
Build opinion and clicking behaviour estimators $(\hat{\beta}, \hat{g})$
Optimization phase
for $k \geq 0$ **do**
  Collect clicks $c_i^k \sim \mathcal{B}\big(g_i(p_i^k, x_i^k)\big)$ from users
  CTR $y^k \leftarrow \frac{\sum_{t=\tau_i}^{k} c^t}{k-\tau_i+1}, \tau_i = (i-1)T < k$
  Estimate opinions $\hat{x}_i^{k+1} \leftarrow \hat{\beta}_i(y_i^k, p^k)$
  if $\zeta^k = 1$ **then**
    $\mathcal{T} \leftarrow \mathrm{append}[k]$
    Estimate sensitivity $\hat{H}^k$
    Estimate gradient
    Update positions p^{k+1}
  else
    $\hat{H}^k \leftarrow \hat{H}^{k-1}$; $p^{k+1} \leftarrow p^k$
  end if
end for

Closed-loop Convergence

We ensure convergence by using the **gradient mapping**

$$\mathcal{G}(p) := \frac{1}{\eta}\Big(p - \mathrm{proj}_{[-1,1]}[p - \eta(\nabla\varphi)]\Big)$$

a common metric to quantify convergence in non convex-regimes.

OFO Convergence

Under all the previous assumptions, for $\eta \in (0, \frac{1}{2(L')})$, $\mu = \mu^*$, the position sequence generated by the projected gradient descent algorithm satisfies

$$\frac{1}{|\mathcal{T}|}\sum_{\substack{l\in\mathcal{T}\\ l\leq k}} \mathbb{E}\left[\|\mathcal{G}(p^l)\|^2\right] \leq K_1, \quad \forall k \geq T$$

$K_1 \propto \varphi(p^0, h(p^0, d)) - \varphi^*, \epsilon_x^2, \epsilon_g^2, L'^2, \frac{1}{\eta^2}$, gradient est. error

Opinion Dynamics and Clicking Behaviour

Extended FJ model

$$x^+ = (I - \Gamma_p - \Gamma_d)Ax + \Gamma_p p + \Gamma_d d$$

Users follow two clicking behaviours

$$c_i \sim \mathcal{B}\Big(\underbrace{\frac{1}{2} + \frac{1}{2}x_i p_i}_{C_a}\Big), \quad c_i \sim \mathcal{B}\Big(\underbrace{\frac{1}{2} + \frac{1}{2}e^{-c(x_i - p_i)^2}}_{C_b}\Big)$$

we perform our algorithm over a network of 15 users, with C_a and C_b randomly distributed. Initial opinion $\sim \mathcal{U}[-1, 1]$, A substochastic, $d^k = x^0 + \text{noise}$, $\Gamma_p \sim \mathcal{U}[10^{-2}, 0.5]$

Training We train the NN for opinion and clicking behaviour with horizon $N = 100$ and collect 75 data points, with trigger period $T = 60$, with the clicks being recorded in the interval $[N - T, N]$. We take $m = 375$ training and 125 testing points.

Online We set $p^0 = 0$ (neutral recommendations). All simulations are conducted for $N = 10^3$ over 50 Monte-Carlo trials.

Architecture Comparison

Method	Sensitivity	Opinions	Clicking behaviour
M_1 (Oracle)	✓	✓	✓
M_2	✗	✓	✓
M_3	✗	✗	✓
M_4 (Alg. 1)	✗	✗	✗

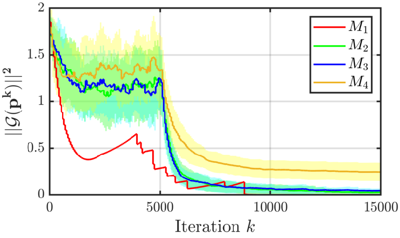

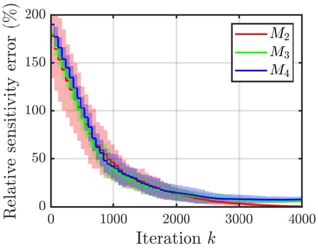

Trading off CTR and Polarization

$$\min_{p,x} \varphi^{\mathrm{clk}}(p,x) + \gamma \varphi^{\mathrm{pol}}(x)$$
$$\text{s.t. } x = h(p,d)$$
$$p \in [-1,1]^n$$

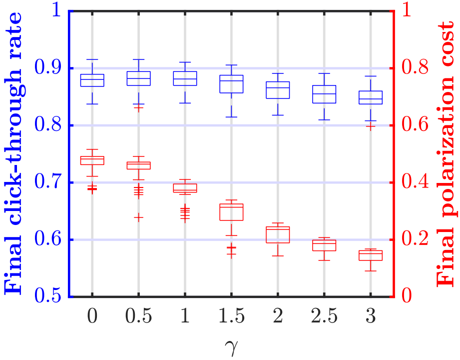

The Impact of the Network

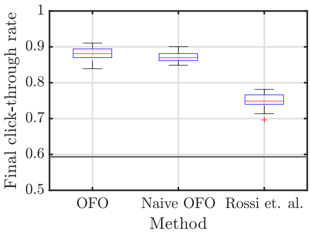

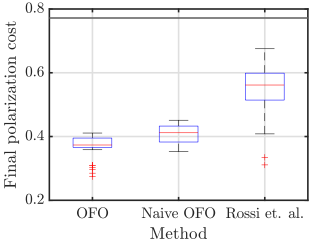

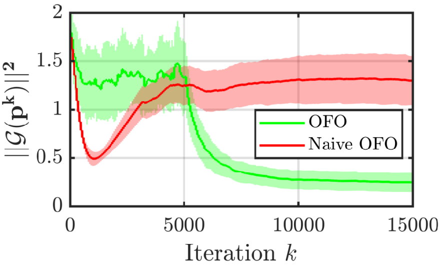

Conclusions

- A Model-free recommender system algorithm that balances engagement maximization and polarization mitigation;
- Theoretical guarantees for CL stability;
- Validation on synthetic data

Future Directions

- Relax smoothness hypothesis on clicking behaviour;
- Consider other interests drivers than confirmation bias, e.g. repulsion.

Acknowledgements

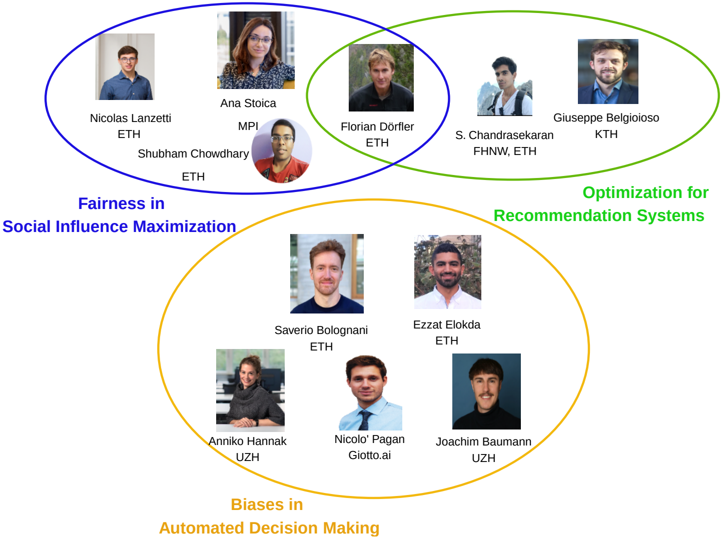

Thanks for your attention

Appendix

Level 2: Online Sensitivity Learning

Sensitivity dynamics as a random process[1]:

$$\mathrm{vec}(\nabla_p h(p,d))^+ = \mathrm{vec}(\nabla_p h(p,d)) + w \qquad \text{Process model}$$
$$\Delta x_{\mathrm{ss}}^+ = \Delta\tilde{p} * \mathrm{vec}(\nabla_p h(p,d)) + v \qquad \text{Measurement model}$$

where

- $\Delta x_{\mathrm{ss}}^+ = h(p^k, d) - h(p^{k-1}, d)$
- $w^k \sim \mathcal{N}(0, Q^k)$
- $v^k \sim \mathcal{N}(0, R^k)$, accounts for the external influence
- $\Delta\tilde{p} = (p^k - p^{k-1})^\top \otimes I_n$

Sensitivity and covariance updates:

$$\mathrm{vec}(\nabla_p h)^k = \mathrm{vec}(\nabla_p h)^{k-1} + \zeta^k (K^{k-1}\Delta\hat{x}^{k+1} - \Delta\tilde{p}^k \mathrm{vec}(\nabla_p h)^{k-1})$$
$$\Sigma^k = \Sigma^{k-1} + \zeta^k (Q^k - K^{k-1}\Delta\tilde{p}^k \Sigma^{k-1})$$

Trigger mecanism: Enforces time-scale separation and ensures that a sufficient number of clicks is collected (clicking ratio accuracy).

[1] Picallo, Ortmann, Bolognani, Dörfler, *Adaptive real time grid operation via online feedback optimization with sensitivity estimation* Electric Power Systems Research, 2022

Note: The CTR is recorded over a time horizon with constant p. The dynamics is exponentially stable: the opinion esitmate is close to the steady state opinion $h(p, d) \to$ we can treat the opinion dynamics as a static map.

CL Convergence

Under all the previous assumptions, the sensitivity estimation error $e^k := vec(h^k) - vec(\hat{h}^k)$ has bias and variance bounded in norm, with

$$\|\mathbb{E}[e^k]\| \leq J_1 \quad \mathbb{E}[\|e^k\|^2] \leq J_2$$

with $J_1, J_2 \propto \epsilon_x, \frac{1}{T}$ and $J_2 \propto \sigma_r^2$