The closed loop between opinion formation and personalised recommendations

Wilbert Samuel Rossi[1], Jan Willem Polderman[2], Paolo Frasca[3]

[2]University of Twente (Netherlands)

[1]University of Groningen (Netherlands)

[3]CNRS, Grenoble (France)

Outline

1. Why this research
2. Dynamical model: interconnecting users and recommenders
   - User model
   - Recommender model
3. Results on the closed-loop system
   - Types of trajectories
   - Simulations and analytical results
4. Conclusion

Social dynamics in the XXI century

Some basic observations:

1. Nowadays, much social dynamics takes place on **online social media**
2. Online activities influence offline behaviours [Aral (2012)]
3. Online dynamics depends on how digital platforms distribute information between the users

Social dynamics in the XXI century

Some basic observations:

1. Nowadays, much social dynamics takes place on **online social media**
2. Online activities influence offline behaviours [Aral (2012)]
3. Online dynamics depends on how digital platforms distribute information between the users

Actually, online platforms manage huge amounts of information:

recommender systems are indispensable, but also blamed for producing "information disorders":

- the formation of **filter bubbles** [Pariser (2011)]
- the viral spreading of **fake news** [Venturini (2019)]

because platforms want to maximize user engagement

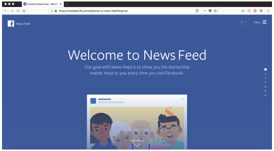

Dynamical model

A model of feedback interconnection

Case study: a **news aggregator** that recommends news articles to **readers**

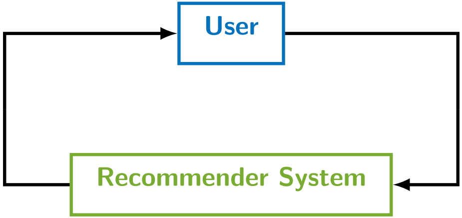

User model: Opinion dynamics

User has a time-dependent *opinion* $o_{usr}(t) \in [-1, 1]$ about an issue

At time t,

- user receives an article that has *position* $p_{art}(t) \in \{-1, 1\}$
- user updates her opinion by

$$o_{usr}(t+1) = \alpha\, o_{usr}^{0} + \beta\, o_{usr}(t) + \gamma\, p_{art}(t) \qquad t \in \mathbb{N}_0$$

where

$o_{usr}^{0} \in [-1, 1]$ is a *prejudice* that coincides with initial opinion (i.e. $o_{usr}(0) = o_{usr}^{0}$)

$\alpha, \beta, \gamma \geq 0$ and $\alpha + \beta + \gamma = 1$ are weights that describe the relative importance of prejudice, memory, and new information

Influence model supported by Chaiken (1987); Friedkin and Johnsen (1990)

User Model: Click Model

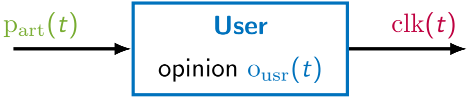

At time t, user also decides whether to *read the recommended article or not*

The user is subject to a **confirmation bias** [Nickerson (1998)]: she prefers contents that are consistent with her opinion o_{usr}

The *click decision* $clk \in \{0, 1\}$ is **stochastic** [Dandekar et al. (2013)]:

$$clk(t) = \begin{cases} 1 & \text{with probability } \frac{1}{2} + \frac{1}{2}\, o_{usr}(t)\, p_{art}(t) \\ 0 & \text{with probability } \frac{1}{2} - \frac{1}{2}\, o_{usr}(t)\, p_{art}(t) \end{cases}$$

Recommender model: solving a bandit problem

The recommender system has the **purpose of maximizing clicks**

(measured by *click-through rate*: $\mathrm{ctr}(t) = \frac{1}{t}\sum_{s=0}^{t-1}\mathrm{clk}(s)$)

Recommender model: solving a bandit problem

The recommender system has the **purpose of maximizing clicks**

(measured by *click-through rate*: $\mathrm{ctr}(t) = \frac{1}{t} \sum_{s=0}^{t-1} \mathrm{clk}(s)$)

The recommender system, therefore, sees
the user as a **"two-armed bandit":**

Recommender model: solving a bandit problem

The recommender system has the **purpose of maximizing clicks**

(measured by *click-through rate*: $\mathrm{ctr}(t) = \frac{1}{t} \sum_{s=0}^{t-1} \mathrm{clk}(s)$)

The recommender system, therefore, sees the user as a **"two-armed bandit":**

This is a one-armed bandit...

Recommender model: solving a bandit problem

The recommender system has the **purpose of maximizing clicks**

(measured by *click-through rate*: $\mathrm{ctr}(t) = \frac{1}{t}\sum_{s=0}^{t-1} \mathrm{clk}(s)$)

The recommender system, therefore, sees the user as a **"two-armed bandit":**

This is a one-armed bandit...

The recommender faces the **exploration-exploitation dilemma**
of sequential decision problems that arises between staying with the most successful option so far (i.e. exploitation) and testing the other option (i.e. exploration), which might become better in the future [Bubeck and Cesa-Bianchi (2012); Li et al. (2010)]

Recommender model: solving a bandit problem

The recommender system has the **purpose of maximizing clicks**

(measured by *click-through rate*: $\mathrm{ctr}(t) = \frac{1}{t}\sum_{s=0}^{t-1} \mathrm{clk}(s)$)

This is a one-armed bandit...

The recommender system, therefore, sees the user as a **"two-armed bandit":**

The recommender faces the **exploration-exploitation dilemma** of sequential decision problems that arises between staying with the most successful option so far (i.e. exploitation) and testing the other option (i.e. exploration), which might become better in the future [Bubeck and Cesa-Bianchi (2012); Li et al. (2010)]

- We model *balancing exploration and exploitation by an* **ϵ-greedy algorithm**

$$\mathrm{p_{art}}(t) = \begin{cases} \textit{exploitation} & \text{with probability } 1-\epsilon \\ \textit{exploration} & \text{with probability } \epsilon \end{cases}$$

Recommender model: Details of the ϵ-greedy algorithm

The recommender needs to compute the **most successful "arm"**

Define counters that track

- recommendations $\mathrm{r}_+(t)$, $\mathrm{r}_-(t)$

$$T_+(t) = \{s : 0 \le s \le t-1 \text{ and } \mathrm{p}_{\mathrm{art}}(s) = +1\} \qquad \mathrm{r}_+(t) = \# T_+$$
$$T_-(t) = \{s : 0 \le s \le t-1 \text{ and } \mathrm{p}_{\mathrm{art}}(s) = -1\} \qquad \mathrm{r}_-(t) = \# T_-$$

- and 'successes' $\mathrm{a}_+(t)$, $\mathrm{a}_-(t)$: $\quad \mathrm{a}_+(t) = \sum_{s \in T_+(t)} \mathrm{clk}(s)\,, \qquad \mathrm{a}_-(t) = \sum_{s \in T_-(t)} \mathrm{clk}(s)$

Apply the randomized decision rule (with small $\epsilon > 0$):

$$\begin{cases} \text{if } \frac{\mathrm{a}_+(t)}{\mathrm{r}_+(t)} > \frac{\mathrm{a}_-(t)}{\mathrm{r}_-(t)} \text{ then} & \mathbb{P}(\mathrm{p}_{\mathrm{art}}(t) = 1) = 1 - \epsilon, \quad \mathbb{P}(\mathrm{p}_{\mathrm{art}}(t) = -1) = \epsilon \\ \text{if } \frac{\mathrm{a}_+(t)}{\mathrm{r}_+(t)} = \frac{\mathrm{a}_-(t)}{\mathrm{r}_-(t)} \text{ then} & \mathbb{P}(\mathrm{p}_{\mathrm{art}}(t) = 1) = 0.5, \quad \mathbb{P}(\mathrm{p}_{\mathrm{art}}(t) = -1) = 0.5 \\ \text{if } \frac{\mathrm{a}_+(t)}{\mathrm{r}_+(t)} < \frac{\mathrm{a}_-(t)}{\mathrm{r}_-(t)} \text{ then} & \mathbb{P}(\mathrm{p}_{\mathrm{art}}(t) = 1) = \epsilon, \quad \mathbb{P}(\mathrm{p}_{\mathrm{art}}(t) = -1) = 1 - \epsilon \end{cases}$$

Detailed feedback interconnection

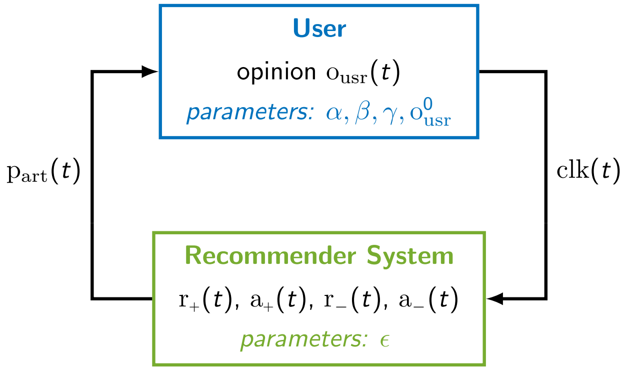

Results: behavior of the interconnection

Example of trajectories: Random recommendations

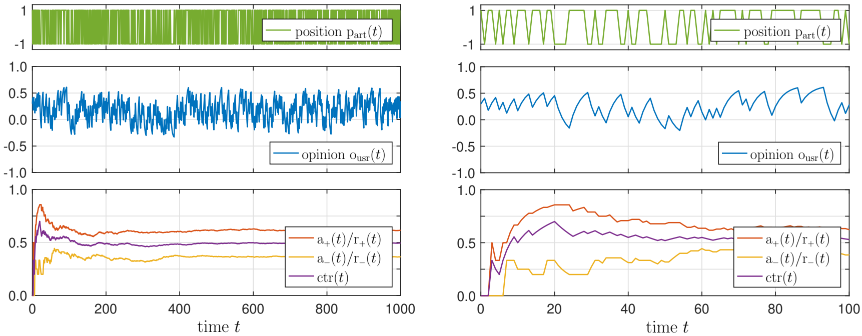

Parameters: $\alpha = 0.15$, $\beta = 0.70$, $\gamma = 0.15$, $o^0_{usr} = 0.30$ and $\epsilon = 0.50$

Left: up to time $t_{max} = 1000$. Right: zooming into the first 100 steps.

Non-random recommendations

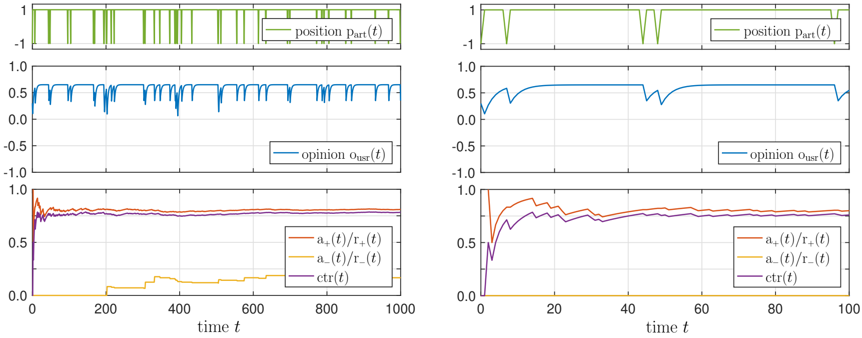

Parameters: $\alpha = 0.15$, $\beta = 0.70$, $\gamma = 0.15$, $o_{usr}^0 = 0.30$ and $\epsilon = 0.05$

Left: up to time $t_{max} = 1000$. Right: zooming into the first 100 steps.

Note: Here the most recommended position is $+1$

$+1$-majority and -1-majority trajectories

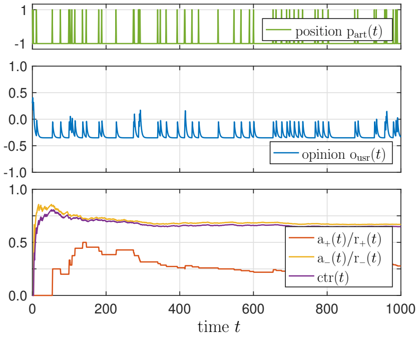

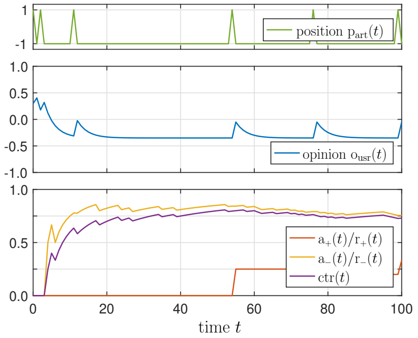

Parameters: $\alpha = 0.15$, $\beta = 0.70$, $\gamma = 0.15$, $o^0_{usr} = 0.30$ and $\epsilon = 0.05$

Left: up to time $t_{max} = 1000$. Right: zooming into the first 100 steps.

Note: Here the most recommended position is -1

Analysis of the closed-loop system

State vector $\mathbf{x}(t) = [\mathrm{r}_+(t), \mathrm{r}_-(t), \mathrm{a}_+(t), \mathrm{a}_-(t), \mathrm{o}_{\mathrm{usr}}(t)]^\top$ has closed dynamics from initial condition $\mathbf{x}(0) = [0, 0, 0, 0, \mathrm{o}_{\mathrm{usr}}^0]^\top$

We could study $\mathbb{E}[\mathbf{x}(t)]$, but...

- the dynamics of $\mathbb{E}[\mathbf{x}(t)]$ is impractical to write due to the **nonlinearities** and **dependences** between the variables
- since there are two kinds of trajectories, an average would be a poor description of either

Analysis of the closed-loop system

State vector $\mathbf{x}(t) = [\mathrm{r}_+(t), \mathrm{r}_-(t), \mathrm{a}_+(t), \mathrm{a}_-(t), \mathrm{o}_{\mathrm{usr}}(t)]^\top$ has closed dynamics from initial condition $\mathbf{x}(0) = [0, 0, 0, 0, \mathrm{o}_{\mathrm{usr}}^0]^\top$

We could study $\mathbb{E}[\mathbf{x}(t)]$, but...

- the dynamics of $\mathbb{E}[\mathbf{x}(t)]$ is impractical to write due to the **nonlinearities** and **dependences** between the variables
- since there are two kinds of trajectories, an average would be a poor description of either

Our approach:

1. We **condition** on the type of trajectory:
   $\mathbb{E}^+[\mathbf{x}(t)] := \mathbb{E}[\mathbf{x}(t) \,|\, +1 \text{ is more likely}]$
   $\mathbb{E}^-[\mathbf{x}(t)] := \mathbb{E}[\mathbf{x}(t) \,|\, -1 \text{ is more likely}]$

Analysis of the closed-loop system

State vector $\mathbf{x}(t) = [\mathrm{r}_+(t), \mathrm{r}_-(t), \mathrm{a}_+(t), \mathrm{a}_-(t), \mathrm{o}_{\mathrm{usr}}(t)]^\top$ has closed dynamics from initial condition $\mathbf{x}(0) = [0, 0, 0, 0, \mathrm{o}^0_{\mathrm{usr}}]^\top$

We could study $\mathbb{E}[\mathbf{x}(t)]$, but...

- the dynamics of $\mathbb{E}[\mathbf{x}(t)]$ is impractical to write due to the **nonlinearities** and **dependences** between the variables
- since there are two kinds of trajectories, an average would be a poor description of either

Our approach:

1. We **condition** on the type of trajectory:
   $\mathbb{E}^+[\mathbf{x}(t)] := \mathbb{E}[\mathbf{x}(t) \mid +1 \text{ is more likely}]$
   $\mathbb{E}^-[\mathbf{x}(t)] := \mathbb{E}[\mathbf{x}(t) \mid -1 \text{ is more likely}]$
2. We write and solve the **linear** dynamics for $\mathbb{E}^\pm[\mathbf{x}(t)]$

Analysis of the closed-loop system

State vector $\mathbf{x}(t) = [\mathrm{r}_+(t), \mathrm{r}_-(t), \mathrm{a}_+(t), \mathrm{a}_-(t), \mathrm{o}_{\mathrm{usr}}(t)]^\top$ has closed dynamics from initial condition $\mathbf{x}(0) = [0, 0, 0, 0, \mathrm{o}^0_{\mathrm{usr}}]^\top$

We could study $\mathbb{E}[\mathbf{x}(t)]$, but...

- the dynamics of $\mathbb{E}[\mathbf{x}(t)]$ is impractical to write due to the **nonlinearities** and **dependences** between the variables
- since there are two kinds of trajectories, an average would be a poor description of either

Our approach:

1. We **condition** on the type of trajectory:
   $\mathbb{E}^+[\mathbf{x}(t)] := \mathbb{E}[\mathbf{x}(t) \mid +1 \text{ is more likely}]$
   $\mathbb{E}^-[\mathbf{x}(t)] := \mathbb{E}[\mathbf{x}(t) \mid -1 \text{ is more likely}]$
2. We write and solve the **linear** dynamics for $\mathbb{E}^\pm[\mathbf{x}(t)]$

We compare analytical $\mathbb{E}^\pm[\mathbf{x}(t)]$ with simulated **time-average** $\bar{\mathbf{x}}(t) = \frac{1}{t}\sum_{s=0}^{t-1} \mathbf{x}(s)$

Results (matching analysis with simulations)

Long-time opinions

Opinions split between $+1$-trajectories and -1-trajectories, concentrating around the conditional expectations

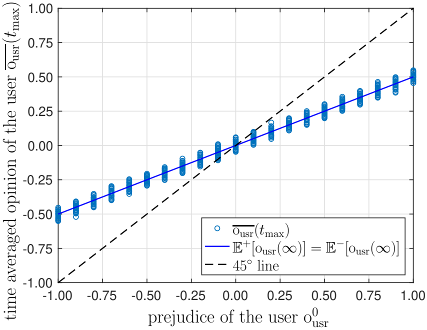

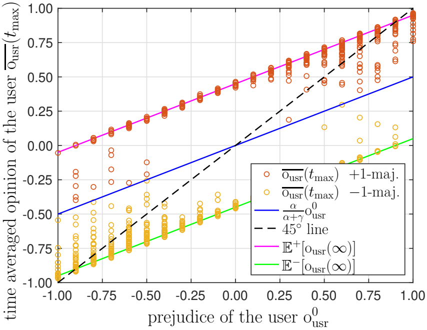

Parameters: $\alpha = 0.15$, $\beta = 0.70$, $\gamma = 0.15$. Left: $\epsilon = 0.50$ (random). Right: $\epsilon = 0.05$

$$\lim_{t\to\infty} \mathbb{E}^{\pm}[o_{usr}(t)] = \frac{\alpha o^0_{usr} \pm \gamma(1-2\epsilon)}{\alpha+\gamma}$$

Prevalence of $+1$ or -1 trajectories

Strong prejudices lead to consistent recommendations

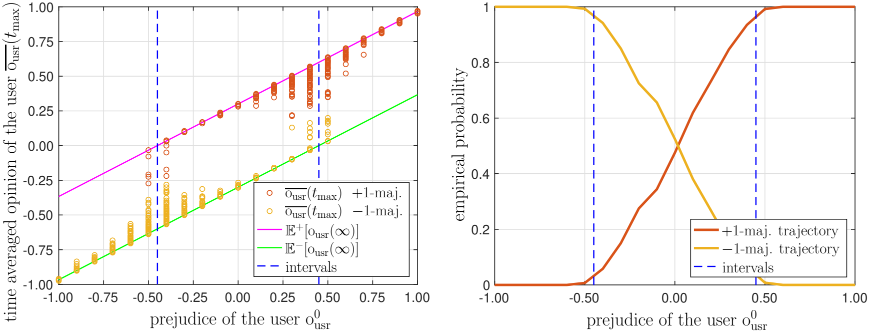

Parameters: $\alpha = 0.20$, $\beta = 0.70$, $\gamma = 0.10$, $\epsilon = 0.05$.

Dashed blue lines have abscissas $-\frac{\gamma}{\alpha}(1-2\epsilon)$ and $\frac{\gamma}{\alpha}(1-2\epsilon)$

Effects on the opinions: Polarization

Most trajectories produce more extreme opinions (polarization)

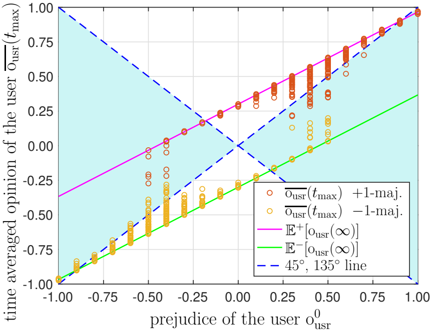

Parameters: $\alpha = 0.20$, $\beta = 0.70$, $\epsilon = 0.05$

In shaded areas, the time averaged opinion $\overline{o_{usr}}(t_{max})$ is *less extreme* than the prejudice o^0_{usr}, i.e. $|\overline{o_{usr}}(t_{max})| \leq |o^0_{usr}|$; in white areas, it is *more extreme*

Combined effects on opinions and click-through rate

Recommendations are more effective when opinions are extreme

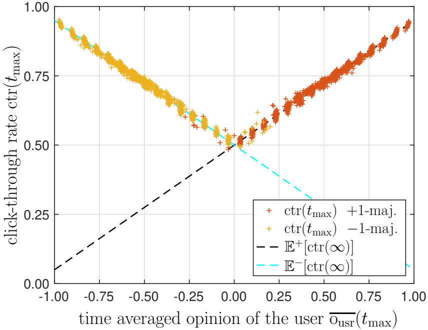

Parameters: $\alpha = 0.20$, $\beta = 0.70$, $\epsilon = 0.05$

$$\mathbb{E}^{\pm}[\mathrm{ctr}(\infty)] = \frac{1}{2} \pm \frac{1}{2}(1 - 2\epsilon)\frac{\alpha o_{\mathrm{usr}}^0 \pm \gamma(1-2\epsilon)}{\alpha+\gamma}$$

Combined effects on opinions and click-through rate II

Effectiveness of recommendations and impact on opinions are positively correlated

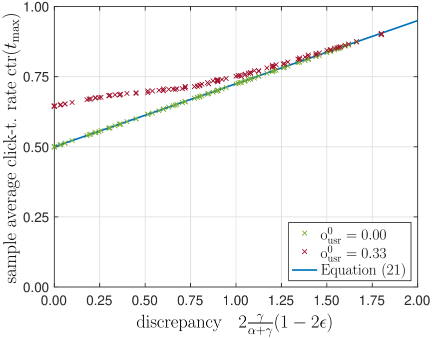

1000 simulations with random parameters α, β, γ and $\epsilon = 0.05$

Discrepancy $\mathbb{E}^+[o_{\mathrm{usr}}(\infty)] - \mathbb{E}^-[o_{\mathrm{usr}}(\infty)] = 2\frac{\gamma}{\alpha+\gamma}(1-2\epsilon)$ measures impact on opinions
Click-through rate measures effectiveness of recommendations

Blue line (21) is $\frac{1}{2}\left(\mathbb{E}^+[\mathrm{ctr}(\infty)] + \mathbb{E}^-[\mathrm{ctr}(\infty)]\right) = \frac{1}{2} + \frac{1}{2}\frac{\gamma}{\alpha+\gamma}(1-2\epsilon)^2$

Combined effects on opinions and click-through rate III

Randomness parameter ϵ controls the trade-off between impact on the opinions and achievable click-through rate

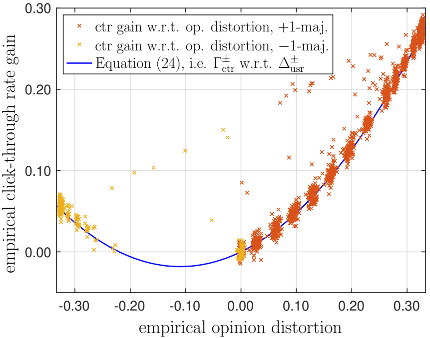

opinion distortion $\Delta^{\pm}_{\mathrm{usr}} := \mathbb{E}^{\pm}[\mathrm{o}_{\mathrm{usr}}(\infty);\epsilon] - \mathbb{E}^{\pm}[\mathrm{o}_{\mathrm{usr}}(\infty);\epsilon = 0.5] = \pm\frac{\gamma}{\alpha+\gamma}(1-2\epsilon)$

click-through rate gain

$\Gamma^{\pm}_{\mathrm{ctr}} := \mathbb{E}^{\pm}[\mathrm{ctr}(\infty);\epsilon] - \mathbb{E}^{\pm}[\mathrm{ctr}(\infty);\epsilon = 0.5] = \pm\frac{1}{2}\frac{\alpha}{\alpha+\gamma}\mathrm{o}^{0}_{\mathrm{usr}}(1-2\epsilon) + \frac{1}{2}\frac{\gamma}{\alpha+\gamma}(1-2\epsilon)^2$

Blue line (24) is $\Gamma^{\pm}_{\mathrm{ctr}} = \frac{1}{2}\frac{\alpha}{\gamma}\mathrm{o}^{0}_{\mathrm{usr}}\Delta^{\pm}_{\mathrm{usr}} + \frac{1}{2}\frac{\alpha+\gamma}{\gamma}(\Delta^{\pm}_{\mathrm{usr}})^2$

Conclusion

Conclusion

Summary

- This was an analytical model of user-recommender interaction (motivated by news aggregators), constructed from "prime principles"
- The connection between personalized recommendations and distorted opinion evolution was made apparent

Conclusion

Summary

- This was an analytical model of user-recommender interaction (motivated by news aggregators), constructed from "prime principles"
- The connection between personalized recommendations and distorted opinion evolution was made apparent

What to do next?

On the sociological/psycological side

- validate the user model (and identify its parameters)
- interpret and validate the recommender model and its tuning

On the machine learning side

- Design optimal recommender algorithms for our closed-loop dynamics

On the modeling side

- Model a network of users
- Refine recommender model (maybe, include collaborative recommendations)

Some references

Aral, S. (2012). Social science: Poked to vote. *Nature*, 489(7415):212–214.

Bubeck, S. and Cesa-Bianchi, N. (2012). Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends® in Machine Learning*, 5(1):1–122.

Chaiken, S. (1987). The heuristic model of persuasion. In *Social influence: the Ontario symposium*, volume 5, pages 3–39. Psychology Press, New York, US.

Dandekar, P., Goel, A., and Lee, D. T. (2013). Biased assimilation, homophily, and the dynamics of polarization. *Proceedings of the National Academy of Sciences*, 110(15):5791–5796.

Friedkin, N. E. and Johnsen, E. C. (1990). Social influence and opinions. *Journal of Mathematical Sociology*, 15(3-4):193–206.

Li, L., Chu, W., Langford, J., and Schapire, R. E. (2010). A contextual-bandit approach to personalized news article recommendation. In *Proceedings of the 19th International Conference on World Wide Web*, WWW '10, pages 661–670, New York, NY, USA. ACM.

Nickerson, R. S. (1998). Confirmation bias: A ubiquitous phenomenon in many guises. *Review of general psychology*, 2(2):175.

Pariser, E. (2011). *The filter bubble: What the Internet is hiding from you*. Penguin UK.

Rossi, W. S., Polderman, J. W., and Frasca, P. (2018). The closed loop between opinion formation and personalised recommendations.

Venturini, T. (2019). From fake to junk news, the data politics of online virality. In Bigo, D., Isin, E., and Ruppert, E., editors, *Data Politics: Worlds, Subjects, Rights*. Routledge, London.