T-RECS: A General Simulation Tool to Study the Impact of Recommendation Systems

PRINCETON UNIVERSITY

CENTER FOR INFORMATION TECHNOLOGY POLICY

Elena Lucherini, Matthew Sun, and Arvind Narayanan

Princeton University

t-recs

Simulation has proved to be a valuable tool in assessing the impact of recommendation systems on the content users consume and on society. We show that questions about filter bubbles, political polarization, information diffusion, and more can be modeled with the same unified simulation framework.

To enable more research on the implications of recommendation systems on society, we present **T-RECS (Tool for RECommender Simulation)**, an open-source tool that allows researchers to model the actions of millions of agents and study the effects of these actions at a macro-level.

SIMULATION IN ALGORITHMIC SYSTEMS

Researchers in computer science and the social sciences have used simulation to study filter bubbles [1, 2, 3, 4], political radicalization [5, 6], and the spread of (mis)information on social media [7, 8].
However, often simulations rely on ad-hoc software that is challenging to develop and reproduce. This has led to efforts to provide a unified simulation environment [9, 10].

REFERENCES

[1] Chaney, A.J., Stewart, B.M. and Engelhardt, B.E., 2018, September. How algorithmic confounding in recommendation systems increases homogeneity and decreases utility. In Proceedings of the 12th ACM Conference on Recommender Systems (pp. 224-232).

[2] Ciampaglia, G.L., Nematzadeh, A., Menczer, F. and Flammini, A., 2018. How algorithmic popularity bias hinders or promotes quality. Scientific reports, 8(1), pp.1-7.

[3] Möller, J., Trilling, D., Helberger, N. and van Es, B., 2018. Do not blame it on the algorithm: an empirical assessment of multiple recommender systems and their impact on content diversity. Information, Communication & Society, 21(7), pp.959-977.

[4] Geschke, D., Lorenz, J. and Holtz, P., 2019. The triple-filter bubble: Using agent-based modelling to test a meta-theoretical framework for the emergence of filter bubbles and echo chambers. British Journal of Social Psychology, 58(1), pp.129-149.

[5] Perra, N. and Rocha, L.E., 2019. Modelling opinion dynamics in the age of algorithmic personalisation. Scientific reports, 9(1), pp.1-11.

T-RECS: DESIGN

T-RECS is designed to study emerging complex phenomena caused by millions of individual actions and interactions with a simulation technique called **agent-based modeling**. Agent-based modeling has been widely used in the social sciences, dating back to the economist Thomas Schelling, who studied models of segregation [11].

T-RECS models agents (e.g., users) interacting with items (e.g., posts or news pieces) over time in a shared environment mediated by an algorithmic system.

T-RECS: BENEFITS

Benefits for researchers seeking to understand the impact of algorithms on society:

1. T-RECS drastically **reduces the engineering effort** needed to develop a simulator and the chances of bugs.
2. T-RECS enables a much **greater volume of simulation studies**.
3. T-RECS allows researchers to **shift the focus from the mechanics** of the simulations **to the assumptions** behind them.

Scientific benefits of applying the same tool to different problems:

1. It **promotes reproducibility** by allowing researchers to easily share their simulations.
2. It **provides a common language** to describe problems in the literature.
3. It **fosters discovery** of principles that apply across seemingly different problems.

[6] Schweighofer, S., Garcia D., and Schweitzer, F., 2020. An agent-based model of multi-dimensional opinion dynamics and opinion alignment. Chaos: An Interdisciplinary Journal of Nonlinear Science, 30(9).

[7] Goel, S., Anderson, A., Hofman, J. and Watts, D.J., 2016. The structural virality of online diffusion. Management Science, 62(1), pp.180-196.

[8] Tambuscio, M., Oliveira, D.F., Ciampaglia, G.L. and Ruffo, G., 2018. Network segregation in a model of misinformation and fact-checking. Journal of Computational Social Science, 1(2), pp.261-275.

[9] Bountouridis, D., Harambam, J., Makhortykh, M., Marrero, M., Tintarev, N. and Hauff, C., 2019, January. Siren: A simulation framework for understanding the effects of recommender systems in online news environments. In Proceedings of the conference on fairness, accountability, and transparency (pp. 150-159).

REPLICATION STUDIES

To demonstrate the effectiveness of T-RECS, we reproduced results from two fundamentally distinct studies. We replicated both studies by adding an average of 50 lines of code to define the core behavior of each system.

Homogenization of content in recommendation systems with T-RECS

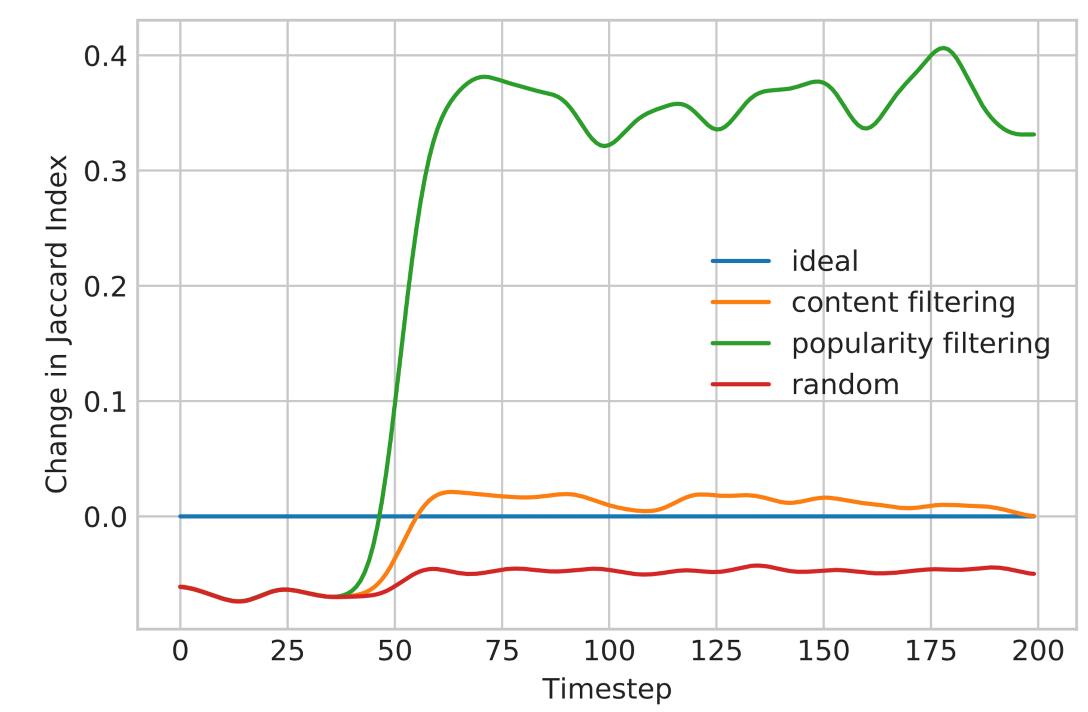

The figure above shows the evolution of homogeneity – that is, how similarly users behave, as defined by Chaney et al. [1] – in four different recommendation system algorithms:

1. Content-based filtering (CF), which recommends content similar to what users have liked in the past.
2. A popularity recommender, which serves the most popular items in the system.
3. A random recommender, which serves random items to users.
4. An “ideal" system, which presents items based on the users' true utility.

In line with the original study, the figure shows that CF and popularity recommendation systems caused users to interact with more similar items than with the ideal system, despite users' differing preferences.

[10] D'Amour, A., Srinivasan, H., Atwood, J., Baljekar, P., Sculley, D. and Halpern, Y., 2020, January. Fairness is not static: deeper understanding of long term fairness via simulation studies. In Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency (pp. 525-534).

Online information diffusion with T-RECS

In our second replication study, we generated scale-free networks of 200, 500, and 1000 users with the parameters that, according to Goel et al. [7], best model the diffusion of viral “events” on Twitter. Observing the spread of tweets from random users, our simulations show that tweets that reach at least 100 users are extremely rare.

CURRENT/FUTURE WORK

Focus on content creators

What happens when we introduce creators with incentives to produce content that maximizes interactions? Does a recommender system like YouTube's incentivize the creation of **increasingly extreme content**? Are creators motivated to publish content that **reinforces viewers' stereotypes**? Do **minority users suffer from decreased utility over time**?

Large-scale replication study

To validate and **expand on the simulation and empirical literature.** Under which conditions **do findings generalize**? Can we **contextualize widely-cited studies yielding contradicting results**, for example on the existence of filter bubbles (e.g., [3, 4])?

The `scikit-learn` of algorithmic simulation

To **democratize algorithmic simulation** by making it more accessible. Can we **develop a standard language** to describe common simulation settings? Can we find shared principles across seemingly distinct problems?

ACKNOWLEDGMENTS

Thanks to Amy Winecoff for the feedback!

[11] Schelling, T.C., 1971. Dynamic models of segregation. Journal of mathematical sociology, 1(2), pp.143-186.

WIP at https://github.com/elucherini/t-recs