



Delft University of Technology

SimuRec

Workshop on synthetic data and simulation methods for recommender systems research

Ekstrand, Michael D.; Chaney, Allison; Castells, Pablo; Burke, Robin; Rohde, David; Slokom, Manel

DOI

[10.1145/3460231.3470938](https://doi.org/10.1145/3460231.3470938)

Publication date

2021

Document Version

Final published version

Published in

RecSys 2021 - 15th ACM Conference on Recommender Systems

Citation (APA)

Ekstrand, M. D., Chaney, A., Castells, P., Burke, R., Rohde, D., & Slokom, M. (2021). SimuRec: Workshop on synthetic data and simulation methods for recommender systems research. In *RecSys 2021 - 15th ACM Conference on Recommender Systems* (pp. 803-805). (RecSys 2021 - 15th ACM Conference on Recommender Systems). Association for Computing Machinery (ACM).
<https://doi.org/10.1145/3460231.3470938>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

SimuRec: Workshop on Synthetic Data and Simulation Methods for Recommender Systems Research

Michael D. Ekstrand
People and Information Research
Team, Boise State University
Boise, Idaho, USA
michael.ekstrand@boisestate.edu

Allison Chaney
Duke University
Durham, North Carolina, USA
ajb.chaney@duke.edu

Pablo Castells
Universidad Autónoma de Madrid
Madrid, Spain
pablo.castells@uam.es

Robin Burke
That Recommender Systems Lab,
University of Colorado at Boulder
Boulder, Colorado, USA
robin.burke@colorado.edu

David Rohde
Criteo
Paris, France
d.rohde@criteo.com

Manel Slokom
TU Delft
Delft, The Netherlands
M.Slokom@tudelft.nl

ABSTRACT

There is significant interest lately in using synthetic data and simulation infrastructures for various types of recommender systems research. However, there are not currently any clear best practices around how best to apply these methods. We proposed a workshop to bring together researchers and practitioners interested in simulating recommender systems and their data to discuss the state of the art of such research and the pressing open methodological questions. The workshop resulted in a report authored by the participants that documents currently-known best practices on which the group has consensus and lays out an agenda for further research over the next 3–5 years to fill in places where we currently lack the information needed to make methodological recommendations.

KEYWORDS

synthetic data, evaluation, simulation

ACM Reference Format:

Michael D. Ekstrand, Allison Chaney, Pablo Castells, Robin Burke, David Rohde, and Manel Slokom. 2021. SimuRec: Workshop on Synthetic Data and Simulation Methods for Recommender Systems Research. In *Fifteenth ACM Conference on Recommender Systems (RecSys '21)*, September 27–October 1, 2021, Amsterdam, Netherlands. ACM, New York, NY, USA, 3 pages. <https://doi.org/10.1145/3460231.3470938>

1 WORKSHOP DESCRIPTION

We proposed a workshop on the appropriate and rigorous use of simulation and synthetic data for studying recommender system effectiveness and behavior. Modeled after the design of SWIRL (Strategic Workshop on Information Retrieval in Lorne), this was a highly interactive workshop focused on discussion between participants to produce a joint report on current knowledge and open

research directions to advance and refine simulation and data synthesis as robust tools for RecSys researchers, practitioners, and students to apply to a range of situations.

1.1 Rationale and Relevance

A growing body of literature on recommender systems is making use of synthetic data and/or simulation methods in order to understand the behavior of recommenders. There are many different uses for such synthetic data: to preserve privacy in underlying data set [5], to test algorithms over a range of data conditions [8], to synthesize unobservable attributes for algorithmic experimentation [1], to study experimental method behavior under controlled conditions [2, 7], and to assess reinforcement learning algorithms [4], among other purposes. Our search of Google Scholar indicates that of articles relevant to the ACM Conference on Recommendation Systems published after 2017, around 27% (13,120¹ out of 11,700²) use or discuss simulations.

Despite the recent surge in use of simulation methods in both academic and industry settings, the assumptions, implementations, and application of these methods vary vastly. Further, there has been little research on the methods themselves to identify the merits of different simulation designs, learn how to properly validate simulations for recommender systems research, or identify best practices in carrying out and reporting on experiments using such methods. This lack of methodological study makes it more difficult to rigorously and robustly apply simulation methods to either research or practice.

This workshop was intended to catalyze that discussion and necessary research. By bringing together a community of researchers and practitioners interested in simulation and data synthesis as a method of studying recommender systems, we identified what is currently known about the methods that could inform best practices, and the open lines of research needed in order to advance simulation as a robust, reproducible, and useful experimental method. This expands the range of research questions that can be studied by the recommender systems community, both by enabling simulation as

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

RecSys '21, September 27–October 1, 2021, Amsterdam, Netherlands

© 2021 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-8458-2/21/09.

<https://doi.org/10.1145/3460231.3470938>

¹https://scholar.google.com/scholar?hl=en&as_sdt=0,34&as_ylo=2017&q=simulation+%22ACM+Conference+on+Recommender+Systems%22

²https://scholar.google.com/scholar?hl=en&as_sdt=0%2C34&as_ylo=2017&q=%22ACM+Conference+on+Recommender+Systems%22&btnG=

a tool for academic research on topics usually limited to industrial application with direct access to user data, and by unlocking new questions that can only be studied with simulation.

1.2 Format

This workshop takes advantage of virtual participation for a novel (to RecSys) working-group format. Instead of a half- or full-day meeting during the conference, participants worked asynchronously through the month *preceding* the conference. We created a Slack channel for participants to discuss the workshop topics, and combined this with three shorter Zoom calls for synchronous discussion: one at the beginning to start the conversation, a midpoint call to check in, and a wrap-up call right before the conference to conclude. If RecSys is able to retain an in-person component, we will arrange a meetup during the conference for participants who are attending in-person. Section 3 documents in more detail how we promoted engagement through this format.

1.3 Submissions

As detailed in Section 4, we solicited short position papers from prospective participants to seed the discussion.

2 ORGANIZERS

Michael D. Ekstrand (michaalekstrand@boisestate.edu) is an Assistant Professor of Computer Science at Boise State University, where he co-directs the People & Information Research Team studying recommender systems and information retrieval from a human-centered perspective. He has organized several workshops on fairness and related topics (FAccTRec @ RecSys 2017, 2018, and 2020; FACTS-IR @ SIGIR 2019; and FairUMAP @ UMAP 2018–2020); served as General Co-chair for RecSys 2018 itself; has served on program committees for RecSys and its workshops; and is on the ACM FAccT executive committee and program committee. He has also published work using simulations for studying recommender systems evaluation [7] and has an active grant that will make further use of simulation methods.

Allison Chaney (ajb.chaney@duke.edu) is an Assistant Professor of Business Administration (Marketing) and Computer Science at Duke University. Her research focuses on developing scalable and interpretable recommendation systems and understanding the impacts of these methods on individuals and society when they are deployed in real-world markets. For example, she has used simulations to explore the impacts of feedback loops in recommendation systems [3]. She was the Program Chair of the 2014 Women in Machine Learning (WiML) Workshop and served on the WiML Board 2016–2019.

Pablo Castells (pablo.castells@uam.es) is an Associate Professor at Universidad Autónoma de Madrid, where he leads the Information Retrieval group, and an Amazon Scholar. In his research he is interested in different problems involved in the design of evaluation experiments, beyond-relevance evaluation, algorithmic and evaluation bias, and interactive recommendation. He has co-organized five workshops at RecSys. He was a program co-chair of RecSys 2016, workshops co-chair in 2020, and DS co-chair in 2021. He served as general co-chair of ECIR 2020; he is program co-chair of SIGIR 2021, and general co-chair of SIGIR 2022.

Robin Burke (robin.burke@colorado.edu) is Professor and Chair at the Dept. of Information Science, University of Colorado where he directs That Recommender Systems Lab and does research on multistakeholder and fairness properties of recommender systems. He is a member of the RecSys Steering Committee and SC chair from 2017–2020. He was the general co-chair of RecSys 2011, and program co-chair of RecSys 2009 and program co-chair of UMAP 2020. He has also organized a number of workshops at RecSys and other venues including VAMS 2017, RMSE 2019, and FairUMAP 2018 and 2019.

David Rohde (d.rohde@criteo.com) is a research scientist at Criteo. His recent research focuses around improving state of the art algorithms using simulation environments. He is one of the original developers of the RecoGym simulator for evaluating recommender systems algorithms. In 2019 he was one of the organisers of the RecoGym Challenge. He has also published novel algorithms and demonstrated their performance using the RecoGym simulator. At UMAP 2020 and RecSys 2020 he delivered tutorials on policy and value based recommendation and used RecoGym as a teaching aid.

Manel Slokom (M.Slokom@tudelft.nl) is a PhD student at Delft University of Technology, The Netherlands. Her research focuses on purpose-aware privacy-preserving data for recommender systems. She works on generating synthetic data that protect users' information while maintaining the quality of recommendation. At the beginning of her PhD, she worked on generating partially synthetic data [5, 6] and now she is exploring and testing fully synthetic data for recommender systems. She is also interested in fairness. She served as a student volunteer at RecSys for three years (2018, 2019 and 2020). In 2021, she is a student volunteer co-chair.

3 DURATION AND ACTIVITIES

As documented in Section 1.2, this workshop is an asynchronous virtual workshop taking place the month prior to the conference. We began on August 27 and concluded on September 24.

The activities for the workshop were:

- Reading position papers submitted by other participants.
- Three one-hour Zoom calls for synchronous discussion of workshop topics.
- Asynchronous discussion via Slack or e-mail of workshop topics.
- Drafting a report on currently-known best practices and open research questions for using simulation and synthetic data for recommender systems research.

In order to promote continued engagement in the workshop, the organizers arranged the work into weekly topics and objectives, and each Monday will posted to the Slack channel to remind the group of the week's topic and start the discussion. Organizers drafted a schedule of topics and discussions based on the position papers, and one of the primary outcomes for the first Zoom call was to discuss the proposed agenda and revise based on feedback from participants. Unlike most RecSys workshops, participation in this workshop's primary activity (the working group) is limited to authors of accepted position papers, to ensure an engaged and committed group; broader community members were welcome to

attend the public session where the working group reports on its findings.

The agenda was designed with topics and outcomes to produce, by the end of the month, a rough draft of the project report (an outline with content from various participants in place; report objectives are detailed in Section 5). After the workshop, the workshop organizers edited the report and circulated to participants for comment and approval.

4 SUBMISSIONS AND SELECTION

Participants were invited through an open call for position papers, broadly distributed (as described in 5). People interested in contributing to the workshop were asked to submit position papers of up to 5 pages in ACM Manuscript format, describing their goals and use cases for simulation methods, and their perspective, experience, or open questions on how to use them effectively and rigorously to advance the state of knowledge in the field. The CFP solicited papers on topics including, but not limited to, the following topics:

- What kinds of research questions and problem settings are simulation methods uniquely suited for?
- What are advantages and possibilities of simulation methods compared to other research approaches?
- What are limitations and pitfalls researchers should be aware of when using simulation?
- What have you found particularly promising or difficult in your own application of simulation to research and/or system development?
- What should the field be studying to improve the rigor and usefulness of simulation?
- What results so far shed light on the effective and appropriate use of simulation?

Authors were specifically asked *not* to include significant empirical results in their position papers, but rather cite their published work or separate preprints with the details of empirical findings (the position paper should, of course, summarize the findings with their relevance to supporting the authors' arguments).

The workshop organizers reviewed submitted position papers, and select final papers (and therefore workshop participants) on the basis of:

- Clarity and rigor of arguments presented
- Ensuring a broad and diverse representation of the different topics, subproblems, methods, and goals discussed

5 OUTCOMES AND DISSEMINATION

The primary outcome of this workshop is a report, jointly authored by the organizers and participants, documenting the group's consensus on the following topics:

- What are use cases where simulation methods are particularly or uniquely useful for promoting research?
- What are use cases where simulation methods are ill-suited?
- What is currently known about how to effectively use simulations and synthetic data for recommender systems research, and what can be promoted as a current best practice?

- How should RecSys research using synthetic data or simulations be evaluated?
- What open questions need further research in order to identify good practices, evaluation criteria, etc. to improve the robustness, validity, and usefulness of simulation-based research methods?

Selected position papers are *not* published by the workshop, and will only be distributed to workshop participants unless the authors themselves distribute elsewhere (e.g. arXiv). Public dissemination of participant positions is through the jointly-authored workshop report.

ACKNOWLEDGMENTS

Michael Ekstrand's participation is supported by the National Science Foundation under grant IIS 17-51278. Robin Burke's participation is supported by the National Science Foundation under grant IIS 19-11025.

REFERENCES

- [1] Robin Burke, Jackson Kontny, and Nasim Sonboli. 2018. Synthetic Attribute Data for Evaluating Consumer-side Fairness. (Sept. 2018). arXiv:1809.04199 [cs.CY]
- [2] Rocío Cañameres and Pablo Castells. 2018. Should I Follow the Crowd?: A Probabilistic Analysis of the Effectiveness of Popularity in Recommender Systems. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval* (Ann Arbor, MI, USA) (SIGIR '18). ACM, New York, NY, USA, 415–424. <https://doi.org/10.1145/3209978.3210014>
- [3] Allison J.B. Chaney, Brandon M. Stewart, and Barbara E. Engelhardt. 2018. How Algorithmic Confounding in Recommendation Systems Increases Homogeneity and Decreases Utility. In *Proceedings of the 12th ACM Conference on Recommender Systems* (RecSys '18). 224–232. <https://doi.org/10.1145/3240323.3240370>
- [4] David Rohde, Stephen Bonner, Travis Dunlop, Flavian Vasile, and Alexandros Karatzoglou. 2018. RecoGym: A Reinforcement Learning Environment for the problem of Product Recommendation in Online Advertising. (Aug. 2018). arXiv:1808.00720 [cs.IR]
- [5] Manel Slokom. 2018. Comparing Recommender Systems Using Synthetic Data. In *Proceedings of the 12th ACM Conference on Recommender Systems* (RecSys '18). Association for Computing Machinery, 548–552. <https://doi.org/10.1145/3240323.3240325>
- [6] Manel Slokom, Martha Larson, and Alan Hanjalic. 2019. Data Masking for Recommender Systems: Prediction Performance and Rating Hiding. *Late breaking results, in conjunction with the 13th ACM Conference on Recommender Systems* (RecSys'19) (2019).
- [7] Mucun Tian and Michael D Ekstrand. 2020. Estimating Error and Bias in Offline Evaluation Results. In *Proceedings of the 2020 Conference on Human Information Interaction and Retrieval* (Vancouver BC, Canada) (CHIIR '20). Association for Computing Machinery, New York, NY, USA, 392–396. <https://doi.org/10.1145/3343413.3378004>
- [8] Karen H L Tso and Lars Schmidt-Thieme. 2006. Empirical analysis of attribute-aware recommender system algorithms using synthetic data. *Journal of computers* 1, 4 (July 2006). <https://doi.org/10.4304/jcp.1.4.18-29>