

NORMALize

The First Workshop on Normative Design and Evaluation of Recommender Systems

Vrijenhoek, Sanne; Michiels, Lien; Kruse, Johannes; Starke, Alain; Viader Guerrero, Jordi; Tintarev, Nava

DOI

[10.1145/3604915.3608757](https://doi.org/10.1145/3604915.3608757)

Publication date

2023

Document Version

Final published version

Published in

Proceedings of the 17th ACM Conference on Recommender Systems, RecSys 2023

Citation (APA)

Vrijenhoek, S., Michiels, L., Kruse, J., Starke, A., Viader Guerrero, J., & Tintarev, N. (2023). NORMALize: The First Workshop on Normative Design and Evaluation of Recommender Systems. In *Proceedings of the 17th ACM Conference on Recommender Systems, RecSys 2023* (pp. 1252-1254). (Proceedings of the 17th ACM Conference on Recommender Systems, RecSys 2023). Association for Computing Machinery (ACM). <https://doi.org/10.1145/3604915.3608757>

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.



NORMALize: The First Workshop on Normative Design and Evaluation of Recommender Systems

Sanne Vrijenhoek
s.vrijenhoek@uva.nl
University of Amsterdam
Amsterdam, The Netherlands

Alain D. Starke
a.d.starke@uva.nl
University of Amsterdam
Amsterdam, The Netherlands
University of Bergen
Bergen, Norway

Lien Michiels
lien.michiels@uantwerpen.be
University of Antwerp
Antwerpen, Belgium

Jordi Viader Guerrero
j.viaderguerrero@tudelft.nl
TU Delft
Delft, The Netherlands

Johannes Kruse
johannes.kruse@eb.dk
Ekstra Bladet
Copenhagen, Denmark

Nava Tintarev
n.tintarev@maastrichtuniversity.nl
Maastricht University
Maastricht, The Netherlands

Abstract

Recommender systems are among the most widely used applications of artificial intelligence. Since they are so widely used, it is important that we, as practitioners and researchers, think about the impact these systems may have on users, society, and other stakeholders. To that effect, the NORMalize workshop seeks to introduce *normative thinking*, to consider the norms and values that underpin recommender systems in the recommender systems community. The objective of NORMalize is to bring together a growing community of researchers and practitioners across disciplines who want to think about the norms and values that should be considered in the design and evaluation of recommender systems; and further educate them on how to reflect on, prioritise, and operationalise such norms and values. NORMalize offers a comprehensive program designed to cater to both the *norm-curious* and the *norm-active*. The morning session is on-site and features a lecture on normative thinking and an interactive workshop. The afternoon is a hybrid program focused on the dissemination of results. NORMalize publishes proceedings, as well as a technical report that summarises the outcomes of the interactive morning session.

CCS Concepts

• **Information systems** → **Recommender systems**; • **Social and professional topics** → **Systems analysis and design**.

Keywords

normative thinking, normative design, recommender systems, norms, values, value-sensitive design

ACM Reference Format:

Sanne Vrijenhoek, Lien Michiels, Johannes Kruse, Alain D. Starke, Jordi Viader Guerrero, and Nava Tintarev. 2023. NORMalize: The First Workshop on Normative Design and Evaluation of Recommender Systems. In

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 '23, September 18–22, 2023, Singapore, Singapore

© 2023 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0241-9/23/09.

<https://doi.org/10.1145/3604915.3608757>

Seventeenth ACM Conference on Recommender Systems (RecSys '23), September 18–22, 2023, Singapore, Singapore. ACM, New York, NY, USA, 3 pages.
<https://doi.org/10.1145/3604915.3608757>

1 Introduction

Users and developers of recommender systems are becoming increasingly aware of the possible societal impact of their systems [5]. As 'beyond-accuracy' metrics are becoming more common in recommender research, much attention has been given to methods related to notions of fairness, such as statistical parity or equality of opportunity in the design or evaluation of recommender systems [7, 8]. However, many values could be considered in the development and goal of a recommender systems, of which fairness towards the end-users of the system is but one example [14].

Identifying and balancing these values requires so-called *normative thinking* and decision-making [1, 2, 12]. Normative thinking requires us to reflect on how or what the system *should be*, rather than focusing on what the current state of the system (output) *is*. Besides identifying relevant values, this includes determining how these values would be expressed in what is recommended by a system, how different values may be conflicting, and justifying how certain values in such cases should be prioritised over others [10].

Normative thinking has always been a pillar of humanities and social science research, but it is finding more and more interest in the exact sciences as well. For example, Springer recently launched the 'AI and Ethics' journal, as well as the AAAI/ACM Conference on Artificial Intelligence, Ethics and Society. Also, within the RecSys community, we see an increase in interest in normative thinking [5–8, 11, 13, 14]. However, leveraging the wealth of knowledge on normative thinking built up within the humanities and social sciences over the years within the exact sciences has proven difficult, as there are currently few opportunities for interdisciplinary collaboration. The goal of this workshop is, therefore, to introduce normative theory to RecSys' predominantly technical audience. By bridging the gap between the humanities and social sciences and the exact sciences, the workshop aims to enhance the normative design and evaluation of recommender systems.

During the morning of the workshop, we will host a lecture from a normative scholar, explaining the core principles of normative thinking. This lecture is followed by breakout groups in which we

work with the participants on defining the goals and values of their systems. The afternoon session aims to connect people working on this topic through presentations and panel discussions of the work submitted by the participants. Our target audience is a mix of academia and industry, and also of both exact and social sciences.

Key Outcomes. The goal of NORMalize is to foster a space for interdisciplinary discussion between the humanities, social sciences, and exact sciences about the norms and values underlying recommender systems. The morning session is intended for the norm-curious and will allow for knowledge sharing and the creation of a shared vocabulary that can facilitate research collaborations. In the afternoon session, we invite the norm-active to share and discuss their work. By approaching the workshop in this way, we facilitate a deep dive with explicit sharing of knowledge and experiences with social scientists and others working on similar problems, which is not usually possible during the main conference. In general, we hope to communicate the importance of normative thinking and to provide recommender system developers with practical tools for the justification of normative design choices.

2 Organiser Biographies

NORMalize is organised by an interdisciplinary team of norm-active and norm-curious researchers and practitioners:

Sanne Vrijenhoek is a PhD Candidate at the University of Amsterdam's Institute of Information Law with a background in Artificial Intelligence. She works in an interdisciplinary project on assessing diversity in news recommendations. An important part of this project is translating normative notions of diversity into concrete concepts that can be used to inform recommender system design. Her work was awarded Best Paper Runner Up at RecSys'22 [14].

Lien Michiels is a PhD Candidate in the ADReM Data Lab at the University of Antwerp, Belgium. She is the lead researcher on the FWO SBO funded 'Serendipity Engine' project for the ADReM Data Lab. As part of this project, she applies normative design principles to urban and news recommender systems leading to more diverse and serendipitous experiences for users. Previously, she combined her PhD research with her work as a Machine Learning Engineer at Froomle where she led the design of its recommendation platform.

Johannes Kruse is an industrial PhD Candidate at the Technical University of Denmark's Department of Applied Mathematics and Computer Science in collaboration with the Danish news publisher Ekstra Bladet. He is in charge of developing and maintaining the core recommendation systems at EkstraBladet.dk, which serve millions of users. He focuses on creating algorithms that provide personalised recommendations while balancing relevance and diversity.

Alain Starke is an assistant professor in persuasive communication for a digital society, at the University of Amsterdam, Netherlands. He is also an adjunct associate professor in recommender systems at the SFI MediaFutures research centre for responsible media technology, which is part of the University of Bergen, Norway. His research aims to develop recommender systems that can support preference shifts and behavioural change in domains of self-actualisation, such as energy conservation, healthy eating, and news recommendation.

Jordi Viader Guerrero (MX/ES) is a practice-based researcher on philosophy of technology and media, and a PhD student at TU Delft on the politics and aesthetics of social media. His research and practice chiefly focus on using audiovisual media and theory to articulate digital technologies within wider cultural, political, and aesthetic logics.

Nava Tintarev is a full professor in explainable AI in the Department of Advanced Computing Sciences at Maastricht University, Netherlands. Her research looks at how to improve transparency in, and decision support for, recommender systems. She is a Co-Investigator in the ROBUST consortium carrying out long-term (10-years) research into trustworthy artificial intelligence. She is also a co-lab director of the TAIM lab, working on trustworthy media, in collaboration with UvA and RTL. Her recent work on a.o., diversification of news and social media items has received four best paper awards in the last 3 years [3, 4, 9, 15].

3 Workshop Format

3.1 Morning

In the on-site morning session, participants are first introduced to the principles and practices of normative thinking. After this lecture, participants are split into breakout groups. In these groups, they discuss a specific use case of a recommender system, e.g., YouTube or Spotify. First, they identify when, where and how the system is used and what it recommends. Then, they identify relevant stakeholders and the norms and values that matter to them. Next, they consider relationships between values and their possible (negative) consequences. For instance, *are diversity and a user's right to relevant content at odds with each other? Or, if we value freedom of speech, could that lead to hate speech and misinformation?* Subsequently, each group is allocated a total of one hundred points, to be divided amongst various values. Each member within the group is given the responsibility to represent a stakeholder of the recommender system and to champion their respective values. The group work concludes with a discussion of what a recommender system that prioritises values and stakeholders in such a way would look like. Finally, each group presents the outcomes of their discussion to all workshop participants and organisers.

3.2 Afternoon

In the hybrid afternoon session, authors of accepted contributions - whether full papers, short papers, or extended abstracts - have the opportunity to present their work. In addition, an expert in the design of normative AI systems delivers a keynote on how to create cutting-edge socio-technical solutions while being mindful of norms and values.

4 Contributions

We accept contributions in the form of full papers (12 pages), short papers (6 pages) and extended abstracts (3 pages). The topics of interest include, but are not limited to:

- **Normative/Value-Sensitive Algorithm Design:** How can different norms and values be operationalised? How can we design algorithms that optimise for these norms and values?

How can we balance these multiple objectives and multiple stakeholders?

- **Metrics & Evaluation Methods:** How should norms and values be measured? What data representations are required to measure a norm or value? How do metrics for norms and values behave in different domains? Do they generalise? How should we design experiments that measure norms and values?
- **Datasets:** How can norms and values be operationalised in available public datasets? What is the influence of the data representation on the metric? New datasets that contain annotations and metadata that can be used to compute metrics of norms and values.
- **Case Studies:** What norms and values are of interest in practice? What issues do you face when operationalising norms and values? Empirical studies of how recommender systems behave with regards to norms and values.
- **Philosophical & Conceptual Work:** Which norms and values are of interest in a specific domain and why? How should these norms and values be balanced?

5 Program Committee

The Program Committee concurs with the interdisciplinary setup of the workshop, mixing scholars and practitioners from computer sciences, social sciences, and humanities. We are grateful for their invaluable contribution to NORMalize:

- **Computer Sciences:** Dietmar Jannach (University of Klagenfurt, AT), Toine Bogers (IT University of Copenhagen, DK), Antske Fokkens and Myrthe Reuver (Vrije Universiteit Amsterdam, NL), Jes Frellsen (Technical University of Denmark, DK), Andres Ferraro (Pandora, US), Meike Zehlike (Zalando, DE), Olivier Jeunen (ShareChat, UK), Savvina Daniil (CWI, NL), Lucien Heitz (University of Zürich, CH)
- **Social Sciences:** Nicholas Diakopoulos (Northwestern University, USA), Kasper Lindskow (Ekstra Bladet), Judith Möller (Hans-Bredow-Institut, DE), Annelien Smets (Imec-SMIT VUB, BE), Rupert Kiddle, Damian Trilling, Valeria Resendez, Sophie Morosoli, Hannes Cools and Mathias Felipe (All University of Amsterdam, NL), Laura Jansen (Wageningen University & Research, NL), Tomás Dodds Rojas (Leiden University, NL)
- **Humanities & Law:** Natali Helberger, Marijn Sax, Max van Drunen, Laurens Naudts, and Naomi Appelman (All University of Amsterdam, NL)

Acknowledgments

We would like to thank the participants and authors of accepted contributions for their valuable inputs to the workshop, as well as the RecSys '23 organisers for their support in the organisation of NORMalize. Finally, we would like to thank our employers and funding bodies. Sanne Vrijenhoek's contribution to this research is supported by the AI, Media and Democracy Lab. Lien Michiels' contribution to this research was supported by the Research Foundation Flanders (FWO) under grant number S006323N. Johannes Kruse's contribution to this research is supported by the Innovation

Foundation Denmark under grant number 1044-00058B and Platform Intelligence in News under project number 0175-00014B. Alain Starke's contribution was in part supported by the Research Council of Norway with funding to MediaFutures: Research Centre for Responsible Media Technology and Innovation, through the Centre for Research-based Innovation scheme, project number 309339. Jordi Viader Guerrero's contribution is supported by the TU Delft AI Labs Initiative and the AI DeMoS Lab. Nava Tintarev's contribution is supported by the project ROBUST: Trustworthy AI-based Systems for Sustainable Growth with project number KICH3.LTP.20.006, which is (partly) financed by the Dutch Research Council (NWO), RTL, and the Dutch Ministry of Economic Affairs and Climate Policy (EZK) under the program LTP KIC 2020-2023. *All content represents the opinion of the authors, which is not necessarily shared or endorsed by their respective employers and/or sponsors.*

References

- [1] Steve Buckler. 2010. Normative theory. *Theory and methods in political science* 3 (2010), 156–180.
- [2] Theresia Anita Christiani. 2016. Normative and empirical research methods: Their usefulness and relevance in the study of law as an object. *Procedia-Social and Behavioral Sciences* 219 (2016), 201–207.
- [3] Tim Draws, Oana Inel, Nava Tintarev, Christian Baden, and Benjamin Timmermans. 2022. Comprehensive viewpoint representations for a deeper understanding of user interactions with debated topics. In *Proceedings of the 2022 Conference on Human Information Interaction and Retrieval*. 135–145.
- [4] Tim Draws, Alisa Rieger, Oana Inel, Ujwal Gadiraju, and Nava Tintarev. 2021. A checklist to combat cognitive biases in crowdsourcing. In *Proceedings of the AAAI conference on human computation and crowdsourcing*, Vol. 9. 48–59.
- [5] Michael D Ekstrand, Mucun Tian, Mohammed R Imran Kazi, Hoda Mehrpouyan, and Daniel Kluger. 2018. Exploring author gender in book rating and recommendation. In *Proceedings of the 12th ACM conference on recommender systems*. 242–250.
- [6] Bart P Knijnenburg, Saadhika Sivakumar, and Daricia Wilkinson. 2016. Recommender systems for self-actualization. In *Proceedings of the 10th ACM conference on recommender systems*. 11–14.
- [7] Rishabh Mehrotra, James McInerney, Hugues Bouchard, Mounia Lalmas, and Fernando Diaz. 2018. Towards a fair marketplace: Counterfactual evaluation of the trade-off between relevance, fairness & satisfaction in recommendation systems. In *Proceedings of the 27th ACM international conference on information and knowledge management*. 2243–2251.
- [8] Erasmo Purificato, Ludovico Boratto, and Ernesto William De Luca. 2022. Do Graph Neural Networks Build Fair User Models? Assessing Disparate Impact and Mistreatment in Behavioural User Profiling. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*. 4399–4403.
- [9] Alisa Rieger, Tim Draws, Mariët Theune, and Nava Tintarev. 2021. This item might reinforce your opinion: Obfuscation and labeling of search results to mitigate confirmation bias. In *Proceedings of the 32nd ACM conference on hypertext and social media*. 189–199.
- [10] Bernd Carsten Stahl. 2012. Morality, ethics, and reflection: a categorization of normative IS research. *Journal of the association for information systems* 13, 8 (2012), 1.
- [11] Alain Starke, Martijn Willemsen, and Chris Snijders. 2021. Promoting energy-efficient behavior by depicting social norms in a recommender interface. *ACM Transactions on Interactive Intelligent Systems (TiIS)* 11, 3-4 (2021), 1–32.
- [12] Judith Jarvis Thomson. 2010. Normativity.
- [13] Christoph Trattner and David Elweiler. 2019. What online data say about eating habits. *Nature Sustainability* 2, 7 (2019), 545–546.
- [14] Sanne Vrijenhoek, Gabriel Bénédict, Mateo Gutierrez Granada, Daan Odijk, and Maarten De Rijke. 2022. RADIO – Rank-Aware Divergence Metrics to Measure Normative Diversity in News Recommendations. In *Proceedings of the 16th ACM Conference on Recommender Systems (Seattle, WA, USA) (RecSys '22)*. Association for Computing Machinery, New York, NY, USA, 208–219. <https://doi.org/10.1145/3523227.3546780>
- [15] Mireia Yurrita, Tim Draws, Agathe Balayn, Dave Murray-Rust, Nava Tintarev, and Alessandro Bozzon. 2023. Disentangling Fairness Perceptions in Algorithmic Decision-Making: the Effects of Explanations, Human Oversight, and Contestability. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–21.