Delft University of Technology

Are We Losing Interest in Context-Aware Recommender Systems?

Rook, Laurens; Zanker, Markus; Jannach, Dietmar

DOI
10.1145/3631700.3665190

Publication date
2024

Document Version
Final published version

Published in
UMAP Adjunct '24

Citation (APA)
Rook, L., Zanker, M., & Jannach, D. (2024). Are We Losing Interest in Context-Aware Recommender Systems? In *UMAP Adjunct '24: Adjunct Proceedings of the 32nd ACM Conference on User Modeling, Adaptation and Personalization* (pp. 229-230). Association for Computing Machinery (ACM). https://doi.org/10.1145/3631700.3665190

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.

This work is downloaded from Delft University of Technology.
For technical reasons the number of authors shown on this cover page is limited to a maximum of 10.

Are We Losing Interest in Context-Aware Recommender Systems?

Laurens Rook
Delft University of Technology
Delft, the Netherlands
l.rook@tudelft.nl

Markus Zanker
Free University of Bozen-Bolzano
Bolzano, Italy
University Klagenfurt
Klagenfurt, Austria
mzanker@unibz.it

Dietmar Jannach
University Klagenfurt
Klagenfurt, Austria
dietmar.jannach@aau.at

ABSTRACT

Contextual information is a prerequisite for timely offering of personalized decision support and recommendation. Yet, research on *context-aware* recommender systems (CARS) does not appear to be thriving, and finding public datasets containing context factors is a challenging task. We can make various assumptions about why this drop in research interest happened – be it ethical considerations or the popularity of opaque deep learning models that merely consider context in an implicit way. This is an unwelcome development. We argue that continued effort must be put on the creation of suitable datasets. Furthermore, we see significant opportunities in the development of next-generation CARS in the space of interactive AI assistants powered by Large Language Models.

CCS CONCEPTS

• **Information systems → Recommender systems**.

KEYWORDS

Context, Context-awareness, Personalization, Recommender Systems, User Intent

ACM Reference Format:
Laurens Rook, Markus Zanker, and Dietmar Jannach. 2024. Are We Losing Interest in Context-Aware Recommender Systems?. In *Adjunct Proceedings of the 32nd ACM Conference on User Modeling, Adaptation and Personalization (UMAP Adjunct '24), July 01–04, 2024, Cagliari, Italy*. ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3631700.3665190

1 SITUATION

Context-aware recommender systems (CARS) were put on the research agenda in the early 2000s [1], and refer to decision support systems that offer automated recommendations to users based on contextual information [2]. In this setting, contextual information refers to observational data such as time, geographical location, sentiment, and presence or absence of others – variables that play an important role in human decision making and preference formation. In case of CARS, contextual attributes are typically used to vary recommendations over a larger user set, based on the preferences of the individual (for reviews, see [11, 13, 16]). CARS can handle state-of-the-art personalization techniques like the ubiquity paradigm, in which contextual information is used to proactively recommend to the user items that may not yet be "top of mind". Precisely this aspect of context adaptation renders CARS particularly suited for the offering of highly targeted – i.e., personalized – forms of decision support [21].

The research community seems to address CARS less often in works presented at conferences such as the ACM Recommender Systems (RecSys) and User Modeling, Adaptation and Personalization (UMAP). As already recognized in [12], a lack of publicly available datasets on context features for recommendation seriously hampers research progress in the community. Ilarri et al. [8] identified only six (relatively small) datasets with potential usefulness for CARS. This may be partly due to ethical considerations. For instance, proactive recommendations based on data harvested from the user's personal space tend to be pervasive – with users' recommendation engagement depending on information privacy concerns [17]. Overly intrusive recommendations may cause users to opt out of data collection, and yield datasets that suffer from missing and/or sparse context variables [8]. In addition, context-aware recommendations that exploit personal and intimate data – such as with whom a user watches a movie or goes out for dinner – can be potentially harmful and be easily at odds with legal frameworks for data management such as the European Union's General Data Protection Regulation (GDPR) [19].

A second possible explanation for the seeming decline in research interest in CARS may be the growing reliance on deep learning models. Such models efficiently take on unstructured, multimodal and multilevel panel data. In contrast to traditional ML models, where ML engineers select features such as variables representing context explicitly, DL models decide on features implicitly as part of the learning process. Thus, they capture user preferences from large amounts of data points without the need to tap explicit contextual attributes [22]. Powerful as these machine learning tools may be, they predict and recommend at the expense of contextual factors. It has been argued elsewhere that deep learning models yield results that are less applicable to opportunities and challenges existing in the real world [10, 11]. Ignoring context may be equalled to missing out on the larger sociotechnical system in which recommendations are provided.

2 OPINION

We recognize that the research community nowadays often implicitly accounts for contextual aspects – cf. session-based or intent-aware recommender systems. In [18], for instance, a diversity-based algorithm taps the user's weak (implicit) rather than strong (explicit) preferences. In this approach, the recommender system is aware of the user's intent in a given context, and makes adjustments

This work is licensed under a Creative Commons Attribution International 4.0 License.

UMAP Adjunct '24, July 01–04, 2024, Cagliari, Italy
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0466-6/24/07
https://doi.org/10.1145/3631700.3665190

accordingly. Our position, however, is that such implicit approaches do not suffice to keep CARS research alive. What is needed is a bold focus on explicit context representation.

Public sharing of datasets with context features for recommendation would be a step in the right direction. Data sharing is an important step to overcome the reproducibility crisis in science [3], and this point is also recognized in the recommender systems community [5, 7]. Researchers should not only ask permission from ethics boards and participants to run the study itself (informed consent), but also explicitly obtain permission to share those data with the research community on later occasions.

People may have explicit or implicit attitudes towards, for instance, a product or brand, and still be willing to change their point of view depending on the context [14, 15]. Thus, the challenge to accurately extract context information from users comes first. Therefore, we call for a reappraisal of conversational recommender systems [9]. Traditionally, these systems explicitly asked users to share contextual aspects relevant for a specific domain. Human preferences and interests (i.e., attitudes) are not necessarily fixed and consistent. In the era of generative LLMs (Large Language Models), a new generation of users is – again – growing accustomed to explicitly communicating their context with the system, as illustrated in this hypothetical request: *I am planning to go for dinner with my friends [social context] after attending the Mets [location/situation context] game on Sunday [temporal context]. Any recommendations?*"

Already in the early days, conversational and knowledge-based recommender systems explicitly asked users about their needs and requirements, thus clarified on the context and intent of users as in [4, 6, 20]. Such an explicit context representation rendered transparently to users why one option was recommended over another. The lack of dynamism and high ramp-up costs of these early systems can nowadays be overcome with generative LLM technology. Therefore, envisioning interactive AI assistants capable of natural language interaction in the tradition of these conversational and knowledge-driven recommender systems is in our view the research opportunity for a next generation of CARS.

REFERENCES

[1] Gediminas Adomavicius and Alexander Tuzhilin. 2005. Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *IEEE Transactions on Knowledge and Data Engineering* 17, 6 (2005), 734–749.

[2] Gediminas Adomavicius and Alexander Tuzhilin. 2010. Context-aware recommender systems. In *Recommender Systems Handbook*. Springer, 217–253.

[3] Monya Baker. 2016. 1,500 scientists lift the lid on reproducibility. *Nature* 533, 7604 (2016).

[4] Robin D. Burke. 1999. The Wasabi Personal Shopper: A Case-Based Recommender System. In *Proceedings of the Sixteenth National Conference on Artificial Intelligence and Eleventh Conference on Innovative Applications of Artificial Intelligence*. 844–849.

[5] Emanuele Cavenaghi, Gabriele Sottocornola, Fabio Stella, and Markus Zanker. 2023. A Systematic Study on Reproducibility of Reinforcement Learning in Recommendation Systems. *ACM Transactions on Recommender Systems* 1, 3 (2023), 1–23.

[6] Alexander Felfernig, Gerhard Friedrich, Dietmar Jannach, and Markus Zanker. 2006. An Integrated Environment for the Development of Knowledge-Based Recommender Applications. *Int. J. Electron. Commer.* 11, 2 (2006), 11–34.

[7] Maurizio Ferrari Dacrema, Paolo Cremonesi, and Dietmar Jannach. 2019. Are we really making much progress? A worrying analysis of recent neural recommendation approaches. In *Proceedings of the 13th ACM Conference on Recommender Systems*. 101–109.

[8] Sergio Ilarri, Raquel Trillo-Lado, and Ramón Hermoso. 2018. Datasets for context-aware recommender systems: Current context and possible directions. In *2018 IEEE 34th International Conference on Data Engineering Workshops (ICDEW)*. 25–28.

[9] Dietmar Jannach, Ahtsham Manzoor, Wanling Cai, and Li Chen. 2022. A Survey on Conversational Recommender Systems. *ACM Comput. Surv.* 54, 5 (2022), 105:1–105:36.

[10] Dietmar Jannach, Pearl Pu, Francesco Ricci, and Markus Zanker. 2022. Recommender systems: Trends and frontiers. *AI Magazine* 43, 2 (2022), 145–150.

[11] Dietmar Jannach, Paul Resnick, Alexander Tuzhilin, and Markus Zanker. 2016. Recommender systems—beyond matrix completion. *Commun. ACM* 59, 11 (2016), 94–102.

[12] Dietmar Jannach, Markus Zanker, Mouzhi Ge, and Marian Gröning. 2012. Recommender systems in computer science and information systems–a landscape of research. In *13th International Conference on E-Commerce and Web Technologies, EC-Web 2012*. 76–87.

[13] Saurabh Kulkarni and Sunil F Rodd. 2020. Context Aware Recommendation Systems: A review of the state of the art techniques. *Computer Science Review* 37 (2020), 100255.

[14] Dhananjay Nayakankuppam and Joseph R Priester. 2014. Putting context effects in context: the construction and retrieval as moderated by attitude strength (CARMAS) model of evaluative judgment. In *Handbook of Brand Relationships*. Routledge, 327–348.

[15] Clint LP Pennings, Jan van Dalen, and Laurens Rook. 2019. Coordinating judgmental forecasting: Coping with intentional biases. *Omega* 87 (2019), 46–56.

[16] Shaina Raza and Chen Ding. 2019. Progress in context-aware recommender systems—An overview. *Computer Science Review* 31 (2019), 84–97.

[17] Laurens Rook, Adem Sabic, and Markus Zanker. 2020. Engagement in proactive recommendations: The role of recommendation accuracy, information privacy concerns and personality traits. *Journal of Intelligent Information Systems* 54, 1 (2020), 79–100.

[18] Saul Vargas, Pablo Castells, and David Vallet. 2011. Intent-oriented diversity in recommender systems. In *Proceedings of the 34th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1211–1212.

[19] Paul Voigt and Axel Von dem Bussche. 2017. The EU general data protection regulation (GDPR). *A Practical Guide, 1st Ed., Cham: Springer International Publishing* 10, 3152676 (2017), 10–5555.

[20] Markus Zanker and Markus Jessenitschnig. 2009. Case-studies on exploiting explicit customer requirements in recommender systems. *User Model. User Adapt. Interact.* 19, 1-2 (2009), 133–166.

[21] Markus Zanker, Laurens Rook, and Dietmar Jannach. 2019. Measuring the impact of online personalisation: Past, present and future. *International Journal of Human-Computer Studies* 131 (2019), 160–168.

[22] Shuai Zhang, Lina Yao, Aixin Sun, and Yi Tay. 2019. Deep learning based recommender system: A survey and new perspectives. *Comput. Surveys* 52, 1 (2019), 1–38.