Context-Aware Recommender Systems: Challenges in Personalization and Fairness

Anna Dalla Vecchia
Supervised by Elisa Quintarelli and Sara Migliorini
University of Verona
Italy
anna.dallavecchia@univr.it

ABSTRACT

Recommender systems play a crucial role in various domains, from education to tourism, by helping users extract hidden knowledge and valuable insights from data. The relevance and effectiveness of these recommendations are closely tied to contextual information, such as user location, weather, and temporal factors, which can enhance both user experience and system performance. However, such contextual data is not always readily available for all users. To address this limitation, this work explores novel approaches to inferring context changes in user behavior. In addition, the concept of fairness is introduced to account for the diverse preferences and constraints of users, aiming to provide equally preferable recommendations across individuals. Specifically, I am investigating the concepts of local and global fairness, highlighting that a universally fair solution may not exist due to differing user needs. As part of my PhD project, I am applying the proposed solutions to realistic case studies from diverse domains, demonstrating the importance of context in recommender systems and the potential for enhancing user satisfaction through a fairness-focused approach.

VLDB Workshop Reference Format:
Anna Dalla Vecchia. Context-Aware Recommender Systems: Challenges in Personalization and Fairness. VLDB 2025 Workshop: VLDB Ph.D. Workshop.

1 INTRODUCTION

Recommender systems (RSs) have become essential tools across various domains, from tourism to healthcare, as they help users extract hidden knowledge and useful insights from datasets. With the growing availability of data from sources such as sensors, wearable devices, and other platforms, the need for intelligent techniques to guide users through data exploration has never been more crucial. As a result, RSs have been widely applied, becoming an invaluable assistant not only for end users, who may be disoriented in the presence of a huge number of different alternatives but also for service providers or sellers, who would like to be able to guide the choice of customers toward particular items with specific characteristics.

In many real-world domains, the temporal order of events, combined with contextual factors such as the user's location, weather, and holidays, can significantly enhance the relevance and effectiveness of recommendations. Therefore, Context-Aware Recommendation Systems (CARS) hold a distinct advantage over more traditional recommenders [1, 9, 21].

However, retrieving contextual information for a specific dataset is often challenging or even infeasible. The context model and its contextual features are typically defined at design time and applied uniformly for all users. This highlights the need for a methodology able to infer such information in its absence and in a personalized way. Recent approaches [8] have demonstrated that it is possible to detect contextual factors by identifying anomalies, enabling systems to better adapt to shifts in user behavior.

Despite RSs' focus on providing personalized and tailored suggestions to individual users or groups, it is essential to question whether they are fair to the broader category for which they are designed. Although the fairness concept depends on the context and domain [16, 20], it generally refers to mitigating biases and disparities while promoting inclusivity and diversity within systems.

This paper presents the first year and a half of my PhD research, illustrating the frameworks developed and the ongoing problem under investigation.

2 RELATED WORK

This section summarizes the state of the art related to the topic investigated in the Thesis, focusing on temporal association rules, context-aware recommender systems, and fairness.

Temporal association rules. In [24], a two-level taxonomy is proposed for integrating temporal aspects in association rule mining. In particular, the first level distinguishes whether time provides order to the data collected, locating some temporal constraints, or is considered as an attribute within the learning process. Our proposal falls in the first category since time is used to obtain totally ordered sequential rules, but at the same time, it escapes from the following low-level categorization in three ways: (i) antecedent and consequent are not only mutually ordered, but the items in the antecedent are also enriched with labels representing their relative order, (ii) thanks to the use of this relative order, we can properly handle holes in the mined rules, (iii) the absolute timestamps can be used to implement an aging mechanism during the computation of the rule support. Although in the literature many proposals extract only partially ordered sequential rules [14], for reducing the number of rules and thus simplifying the mining process, we deem that many real-world domains are based on data and events that occur in a specific order, which must be considered during the knowledge discovery processes. Similarly, [3] deals with the problem of Mining Precise-positioning Episode Rules, introducing

This work is licensed under the Creative Commons BY-NC-ND 4.0 International License. Visit https://creativecommons.org/licenses/by-nc-nd/4.0/ to view a copy of this license. For any use beyond those covered by this license, obtain permission by emailing info@vldb.org. Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.
Proceedings of the VLDB Endowment. ISSN 2150-8097.

fixed-gap episodes for specifying the exact time of the consequent events. To the best of our knowledge, no existing work combines both relative order and absolute timestamp in rule construction as our framework does.

Context-Aware Recommender Systems. RSs offer suggestions on items, services, or news that may interest users and affect their decisions based on their profile, history, and preferences [22]. For instance, in [7], the authors develop an RS that can suggest activities targeted to specific users to improve their health conditions starting from data collected by a Fitbit wearable device. The physical activity information collected by Fitbit is also used in [23] to correlate daily physical activity levels with predictions of sleep quality. Neither of the mentioned works considers contextual information.

CARSs have received a lot of attention in recent years since their aim is to decrease recommendation errors and improve the system's quality by enriching available item and user information with contextual data [1, 9, 21]. In the state of the art, contextual information is often integrated into RS to improve the precision of recommendations. In general, user preferences may vary depending on the environment and the situation in which they are acting [2]. Therefore, CARS use contextual information, such as time, location, and social situation, to add knowledge during the recommendation process, thus improving the personalization and the relevance of the suggestion [1]. A systematic literature review is proposed in [28], where the authors describe the integration of the context in RS, the main categories of contextual features, and the validation mechanism with respect to datasets, properties, metrics, and evaluation protocols. In the urban tourist scenario, multiple factors may be considered. However, in existing works the contextual information is quite limited and restricted to only a weather feature in [27] and to an hourly weather summary (e.g. cloudy), temperature (e.g. cold), and temporal information such as the time interval related to the visit (e.g. evening) in [18]. Although these proposals deem as important the notion of context, they do not consider other external information like the occupancy rate of each PoI, the day of the week, the presence of holidays, or other important events in the considered city.

Fairness. The authors in [16] examine the definition of fairness as the absence of discrimination for individuals with the same "merit" and fairness in algorithms as the absence of discrimination. However, they point out three weaknesses of this definition: disparities justified by "merit", the limitation of the algorithm, and the ignoring of the disparities within groups. It confirms that, although the problem has been explored for years, it is not easy to delineate in a uniform way the concept of fairness [15, 20] Fairness in task assignment is addressed in [31] through a coalition-based approach, where workers cooperate on spatial tasks to maximize their overall rewards. However, unlike our setting, these methods focus on searching for stable worker coalitions rather than assigning a predefined set of mandatory tasks for each stakeholder. Similarly, the fair allocation strategies proposed in [5, 30] do not apply to our scenario, as they do not accommodate sequential task allocation.

Furthermore, my work delves into a deeper discussion of the meaning of fairness constraints and preferences. The importance of evaluating fairness from multiple stakeholders' perspectives has been investigated in [4, 29] in the context of RSs, with a special focus on tourism.

3 PROBLEM STATEMENT

Nowadays, RSs are increasingly integrated into everyday users' lives across various domains. For this reason, the personalization process plays a crucial role in providing useful and customized recommendations. External contextual factors, such as time, location, weather conditions, or social setting, are essential to dynamically adapt suggestions to each user and situation.

The main objective of this work is the development of a CARS that can automatically recognize and extract context from data, even when such information is not explicitly provided. The main challenge concerns datasets coming from real-world scenarios, since contextual data is not always readily available. Therefore, in the absence of explicitly provided context, the system must autonomously detect and adapt to changes in user context and routine by analyzing data. Furthermore, users may have diverse needs, constraints, and preferences. A fair recommender system must avoid favoring a specific group of users. Considering fairness in the recommendation process introduces an additional layer of complexity that has to be addressed.

This research follows a stepwise path:

- First, the role of context is investigated and validated through two real-world scenarios: tourism and wearable device data.
- Then, we investigate anomaly detection techniques to infer context changes, in order to integrate this component into our recommendation framework.
- Finally, we define a fairness-aware evaluation metric to assess whether the recommendations satisfy users' individual requirements.

4 PRELIMINARY RESULTS

This section presents the preliminary results of my research. First, the role of context is investigated in two domains: tourism and wearable devices. Then, focusing on the latter use case, we apply anomaly detection techniques to automatically infer context changes. Finally, we discuss the definition of fairness metrics.

4.1 Role of Context in Tourism RSs

As highlighted in the previous sections, context-aware recommender systems (CARSs) are becoming essential tools in various domains. In this first case study, we focus on the tourism domain by studying the role of contextual information in determining both Points of Interest (PoI) occupations and user preferences. We investigate how machine learning and deep learning techniques can improve the recommendation quality by enriching historical information with its contextual counterpart.

As a result, we propose a context-Aware Recommender sysTEM wIth crowding forecaSting called ARTEMIS, a system designed to learn and forecast user preferences and occupation levels based on historical contextual features [19]. Specifically, we apply this system to a real-world scenario, analyzing tourist visits in Verona, Italy, between 2014 and 2019. The general architecture of ARTEMIS is shown in Fig. 1, and includes the following components: (i) $C1$ is the crowding forecaster, a Deep Neural Network model trained

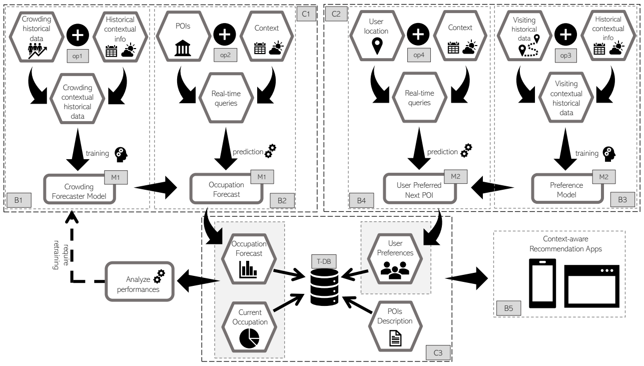

Figure 1: ARTEMIS Architecture.

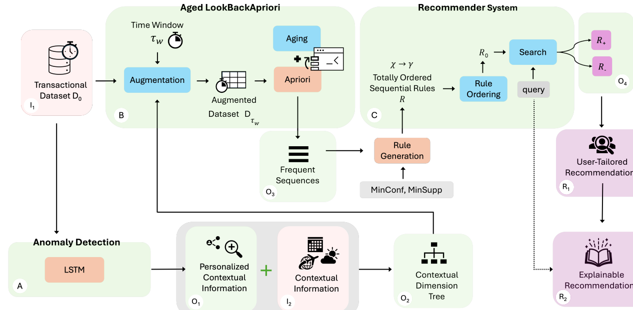

Figure 2: ACTER Architecture.

on the historical and contextual information in order to predict real-time PoI occupation, which is updated with the actual user affluence; (ii) $C2$ is the contextual preferences estimator, which learns the user habits from visiting historical data and context; lastly, (iii) $C3$ combines the information produced by the contextual crowding forecaster and the contextual preference estimator to produce the final set of recommendations, which will be returned to the user through a user interface or application.

Using the same dataset, we also developed a recommender system to promote sustainable tourism [13]. The concept of sustainable tourism covers many aspects, from economic, social, and environmental issues to attention to improving tourists' experience and the needs of host communities. In this regard, one of the most important aspects is the prevention of over-tourism, i.e., overcrowding at specific attractions or locations. Moreover, instead of suggesting the next PoI to visit in a given situation, this system proposes a complete sequence of PoIs (tourist itinerary) that covers an entire day or vacation period. In order to address this task, we develop a Deep Reinforcement Learning approach, where the tourist's reward function depends on the specific spatial and temporal context in which the itinerary has to be performed.

4.2 Role of Context in Wearable Device RSs

The second domain of interest is healthcare, where we analyze a real-world dataset collected from Fitbit wearable devices. The final goal is to improve the user's sleep score after performing fitness activities in different contextual conditions.

Firstly, we introduce ALBA (AgedLookBackApriori)[10, 17], a general extension of sequential rule mining. Unlike other literature methods, ALBA not only ensures that the antecedent precedes the consequent but also enriches itemsets with an explicit representation of their relative temporal order. This enrichment allows the generation of more precise, timely recommendations. Using ALBA, we validate the relevance of the contextual information in the Fitbit dataset, highlighting the importance of an aging mechanism [11]. Our methodology improves the accuracy of the recommendation, especially when enriched with contextual information.

As previously discussed, contextual information is limited, particularly when personal information is involved, as in this domain. To address this issue, we propose a methodology based on anomaly detection [12] to automatically identify context changes directly from existing temporal data. Specifically, we adopt a Long Short-Term Memory approach, which detects context change in temporal datasets and tests it against a dataset of Fitbit data collected from the participants in this case study. The results show that the technique easily highlights periods of time when the users' context was indeed different from their normal routine, thus successfully detecting a change in their context.

All these insights are integrated into a unified framework called ACTER (Activity Customization through Timely and Explainable Recommendations). The framework can, on one side, learn directly from sensor data, both user habits and the main contextual perspectives that have an impact on each user's behavior, and, on the other side, leverage such knowledge to provide contextual, personalized, timely, and explainable recommendations. In Fig. 2, the overall architecture is depicted: (i) the component A leverages historical data to detect time intervals where the user behavior has changed from the most frequent trends, thus refining the context notion defined at design time; (ii) the second component B generates frequent and ordered contextual sequences by applying ALBA; (iii) lastly, C represents the recommender system. It mines totally ordered sequential rules, which are ordered and split into two sets of positive and negative rules (R+ and R-). These rule sets are helpful in providing suggestions that are coherent with the user's current situation, with their explanation.

4.3 Fairness in RSs with Multiple Stakeholders

Recommendation tasks often involve multiple stakeholders with diverse objectives, preferences, and constraints. Ensuring fairness in these scenarios is crucial to maintain user satisfaction, improve efficiency, and mitigate biases. Although the definition of fairness is context-dependent, it generally aims to ensure an equal distribution of tasks among participants while respecting user-defined constraints and minimizing disparities.

In our work, we explore fairness in sequential task assignments, introducing the notions of local and global fairness. Firs of all we distinguish between two kinds of constraints: *Hard constraints* that are a set of mandatory constraints that must be satisfied by any valid solution (e.g., a user cannot be assigned to two tasks at the same time); and *Soft constraints* that are a set of constraints, including preferences, that are not strictly mandatory, but stakeholders can assign higher or lower importance to them (e.g., a user prefers a task in a specific time slot). Based on these definitions, the **Local**

Fairness is the degree to which a valid solution satisfies the set of soft constraints for a specific stakeholder, while **Global Fairness** evaluates how a valid solution satisfies the sets of soft constraints for each stakeholder involved.

To address this problem, we propose FaST-MOSA, a multi-objective algorithm based on simulated annealing [25], which starts from solutions that satisfy the local fairness, and generates a unique assignment satisfying the global fairness. Our evaluation shows that this approach improves fairness with fewer iterations and better acceptance among users, as confirmed through a questionnaire-based user study.

5 CONCLUSION AND FUTURE WORKS

This work presents two frameworks aimed at tackling the challenges of integrating context into real-world datasets to provide personalized recommendations to users. These methodologies are validated in two domains: tourism and wearable devices, highlighting the advantages of incorporating contextual informations into recommendations. Furthermore, concepts of global and local fairness are introduced to improve overall user satisfaction.

I am currently investigating the relationship between fairness and context to gain deeper understanding of user behavior and contextual dynamics. In particular, I am exploring alternative techniques for detecting changes in context, such as the ConQuest approach proposed in [8], in order to automatically discover and incorporate contexts. Furthermore, I am studying the application of machine unlearning techniques, as introduced in [6]. Inspired by the approach taken in [26], where unlearning was used to detect unfairness in binary classifiers, I aim to adapt this strategy to identify unbiased patterns in the context to improve recommendations.

REFERENCES

[1] Gediminas Adomavicius, Bamshad Mobasher, Francesco Ricci, and Alexander Tuzhilin. 2011. Context-Aware Recommender Systems. *AI Magazine* 32, 3 (Oct. 2011), 67–80. https://doi.org/10.1609/aimag.v32i3.2364

[2] Gediminas Adomavicius, Ramesh Sankaranarayanan, Shahana Sen, and Alexander Tuzhilin. 2005. Incorporating contextual information in recommender systems using a multidimensional approach. *ACM Trans. Inf. Syst.* 23, 1 (2005), 103–145.

[3] Xiang Ao, Ping Luo, Jin Wang, Fuzhen Zhuang, and Qing He. 2018. Mining Precise-Positioning Episode Rules from Event Sequences. *IEEE Transactions on Knowledge and Data Engineering* 30, 3 (2018), 530–543. https://doi.org/10.1109/TKDE.2017.2773493

[4] Ashmi Banerjee, Paromita Banik, and Wolfgang Wörndl. 2023. A review on individual and multistakeholder fairness in tourism recommender systems. *Frontiers in Big Data* 6 (2023). https://doi.org/10.3389/fdata.2023.1168692

[5] Fuat Basık, Buğra Gedik, Hakan Ferhatosmanoğlu, and Kun-Lung Wu. 2021. Fair Task Allocation in Crowdsourced Delivery. *IEEE Transactions on Services Computing* 14, 4 (2021), 1040–1053. https://doi.org/10.1109/TSC.2018.2854866

[6] Jonathan Brophy and Daniel Lowd. 2021. Machine Unlearning for Random Forests. In *Proceedings of the 38th International Conference on Machine Learning (Proceedings of Machine Learning Research)*, Marina Meila and Tong Zhang (Eds.), Vol. 139. PMLR, 1092–1104. https://proceedings.mlr.press/v139/brophy21a.html

[7] André Calero Valdez, Martina Ziefle, Katrien Verbert, Alexander Felfernig, and Andreas Holzinger. 2016. Recommender systems for health informatics: state-of-the-art and future perspectives. *Machine Learning for Health Informatics: State-of-the-Art and Future Challenges* (2016), 391–414.

[8] Ece Calikus, Slawomir Nowaczyk, and Onur Dikmen. 2025. Context discovery for anomaly detection. *International Journal of Data Science and Analytics* 19, 1 (2025), 99–113.

[9] Guanliang Chen and Li Chen. 2015. Augmenting service recommender systems by incorporating contextual opinions from user reviews. *User Model. User Adapt. Interact.* 25, 3 (2015), 295–329. https://doi.org/10.1007/s11257-015-9157-3

[10] Anna Dalla Vecchia, Niccolò Marastoni, Sara Migliorini, Barbara Oliboni, and Elisa Quintarelli. 2023. Mining Totally Ordered Sequential Rules to Provide Timely Recommendations. In *New Trends in Database and Information Systems*, Alberto Abelló, Panos Vassiliadis, Oscar Romero, Robert Wrembel, Francesca Bugiotti, Johann Gamper, Genoveva Vargas Solar, and Ester Zumpano (Eds.). Springer Nature Switzerland, Cham, 197–207.

[11] Anna Dalla Vecchia, Niccolò Marastoni, Barbara Oliboni, and Elisa Quintarelli. 2023. The Synergies of Context and Data Aging in Recommendations. In *Big Data Analytics and Knowledge Discovery*. Springer Nature Switzerland, Cham, 80–87.

[12] Anna Dalla Vecchia, Niccolò Marastoni, and Elisa Quintarelli. 2024. Anomaly detection to infer context changes in temporal data. In *2024 IEEE 18th International Conference on Application of Information and Communication Technologies (AICT)*. 1–6. https://doi.org/10.1109/AICT61888.2024.10740439

[13] Anna Dalla Vecchia, Sara Migliorini, Elisa Quintarelli, Mauro Gambini, and Alberto Belussi. 2024. Promoting sustainable tourism by recommending sequences of attractions with deep reinforcement learning. *Information Technology & Tourism* 26, 3 (2024), 449–484.

[14] Philippe Fournier-Viger, Peng Yang, Zhitian Li, Jerry Chun-Wei Lin, and Rage Uday Kiran. 2020. Discovering rare correlated periodic patterns in multiple sequences. *Data Knowl. Eng.* 126 (2020), 101733.

[15] Ben Hutchinson and Margaret Mitchell. 2019. 50 Years of Test (Un)fairness: Lessons for Machine Learning. In *Proceedings of the Conference on Fairness, Accountability, and Transparency* (Atlanta, GA, USA) *(FAT* '19)*. Association for Computing Machinery, New York, NY, USA, 49–58. https://doi.org/10.1145/3287560.3287600

[16] Maximilian Kasy and Rediet Abebe. 2021. Fairness, Equality, and Power in Algorithmic Decision-Making. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency* (Virtual Event, Canada) *(FAccT '21)*. Association for Computing Machinery, New York, NY, USA, 576–586. https://doi.org/10.1145/3442188.3445919

[17] Niccoló Marastoni, Barbara Oliboni, and Elisa Quintarelli. 2022. Explainable Recommendations for Wearable Sensor Data. In *Big Data Analytics and Knowledge Discovery*, Robert Wrembel, Johann Gamper, Gabriele Kotsis, A. Min Tjoa, and Ismail Khalil (Eds.). Springer International Publishing, Cham, 241–246.

[18] David Massimo and Francesco Ricci. 2021. Next-POI Recommendations Matching User's Visit Behaviour. In *Inform. and Communication Technologies in Tourism 2021*. Springer International Publishing, 45–57.

[19] Sara Migliorini, Anna Dalla Vecchia, Alberto Belussi, and Elisa Quintarelli. 2024. ARTEMIS: a Context-Aware Recommendation System with Crowding Forecaster for the Touristic Domain. *Information Systems Frontiers* (2024), 1–27.

[20] Arvind Narayanan. 2018. Translation tutorial: 21 fairness definitions and their politics. In *Proc. conf. fairness accountability transp., new york, usa*, Vol. 1170. 3.

[21] Hossein A. Rahmani, Yashar Deldjoo, and Tommaso Di Noia. 2022. The role of context fusion on accuracy, beyond-accuracy, and fairness of point-of-interest recommendation systems. *Expert Syst. Appl.* 205 (2022), 117700. https://doi.org/10.1016/j.eswa.2022.117700

[22] Francesco Ricci, Lior Rokach, and Bracha Shapira. 2011. Introduction to recommender systems handbook. In *Recommender systems handbook*. Springer, 1–35.

[23] Aarti Sathyanarayana, Shafiq Joty, Luis Fernandez-Luque, Ferda Ofli, Jaideep Srivastava, Ahmed Elmagarmid, Teresa Arora, and Shahrad Taheri. 2016. Sleep quality prediction from wearable data using deep learning. *JMIR mHealth and uHealth* 4, 4 (2016), e125.

[24] Alberto Segura-Delgado, María José Gacto, Rafael Alcalá, and Jesús Alcalá-Fdez. 2020. Temporal association rule mining: An overview considering the time variable as an integral or implied component. *WIREs Data Mining and Knowledge Discovery* 10, 4 (2020), e1367. https://doi.org/10.1002/widm.1367

[25] Balram Suman. 2004. Study of simulated annealing based algorithms for multiobjective optimization of a constrained problem. *Computers & Chemical Engineering* 28, 9 (2004), 1849–1871.

[26] Tanmay Surve and Romila Pradhan. 2025. Explaining Fairness Violations using Machine Unlearning. (2025).

[27] Christoph Trattner, Alexander Oberegger, Lukas Eberhard, Denis Parra, and Leandro Balby Marinho. 2016. Understanding the Impact of Weather for POI Recommendations. In *Proceedings of the Workshop on Recommenders in Tourism (CEUR Workshop Proc.)*, Vol. 1685. 16–23.

[28] Norha M. Villegas, Cristian Sánchez, Javier Díaz-Cely, and Gabriel Tamura. 2018. Characterizing context-aware recommender systems: A systematic literature review. *Knowledge-Based Systems* 140 (2018), 173–200.

[29] Yifan Wang, Weizhi Ma, Min Zhang, Yiqun Liu, and Shaoping Ma. 2023. A Survey on the Fairness of Recommender Systems. *ACM Trans. Inf. Syst.* 41, 3 (2023), 52:1–52:43.

[30] Qing Chuan Ye, Yingqian Zhang, and Rommert Dekker. 2017. Fair task allocation in transportation. *Omega* 68 (2017), 1–16. https://doi.org/10.1016/j.omega.2016.05.005

[31] Yan Zhao, Kai Zheng, Ziwei Wang, Liwei Deng, Bin Yang, Torben Bach Pedersen, Christian S Jensen, and Xiaofang Zhou. 2024. Coalition-based task assignment with priority-aware fairness in spatial crowdsourcing. *The VLDB Journal* 33, 1 (2024), 163–184.