

Hybrid AI for Human-Centric Personalization (HyPer)

Elisabeth Lex
elisabeth.lex@tugraz.at
Graz University of Technology
Graz, Austria

Kevin Innerebner
innerebner@student.tugraz.at
Graz University of Technology
Graz, Austria

Marko Tkalcic
marko.tkalcic@famnit.upr.si
University of Primorska
Koper, Slovenia

Dominik Kowald
dkowald@know-center.at
Know Center Research GmbH & Graz
University of Technology
Graz, Austria

Markus Schedl
markus.schedl@jku.at
Johannes Kepler University Linz &
LIT AI Lab
Linz, Austria

Abstract

Hybrid AI, which integrates symbolic and sub-symbolic methods, has emerged as a promising paradigm for advancing human-centric personalization. By combining machine learning with structured knowledge representations, hybrid AI enables interpretable and adaptive user models that account for complex human factors such as biases, mental models, and affective states. The HyPer workshop focuses on how hybrid AI approaches—combining neural architectures, symbolic representations, and cognitive/behavioral frameworks—can foster more explainable and personalized user experiences. Specifically, we aim to explore innovative applications of hybrid AI in personalization, bridging the gap between explainability, cognitive modeling, and automated adaptation to user preferences. The HyPer workshop will provide a venue for researchers and practitioners to discuss the latest advancements, challenges, and future directions in this interdisciplinary field.

Keywords

Hybrid User Models, Hybrid AI, Symbolic AI, Sub-symbolic AI, Cognitive Models

ACM Reference Format:

Elisabeth Lex, Kevin Innerebner, Marko Tkalcic, Dominik Kowald, and Markus Schedl. 2025. Hybrid AI for Human-Centric Personalization (HyPer). In . ACM, New York, NY, USA, 3 pages. <https://doi.org/10.1145/nnnnnnn>. nnnnnnn

1 Introduction and Motivation

As AI-driven personalization becomes increasingly pervasive, there is a growing demand for interpretable, trustworthy, and human-centered user modeling techniques. Traditional AI-based personalization systems rely heavily on data-driven, black-box machine

learning models, which often lack transparency and do not adequately capture complex human behaviors and cognitive processes [26]. Hybrid AI models integrate symbolic reasoning with sub-symbolic, data-driven learning [3], and have emerged as a promising solution for advancing human-centric personalization by combining structured knowledge representations with adaptive, data-driven models [31]. Corresponding hybrid AI systems can combine the interpretability of symbolic AI with the learning capabilities of sub-symbolic models, and thus, enable more transparent and explainable user modeling while maintaining the flexibility and adaptability of machine learning-based approaches.

Despite these advantages, hybrid AI-based personalization faces several challenges. Integrating symbolic and sub-symbolic AI requires designing architectures that efficiently combine discrete, logic-based representations with continuous, high-dimensional feature spaces [19]. Ensuring fairness, mitigating biases, and aligning hybrid models with human cognitive processes poses another challenge [17]. Additionally, dynamic user preferences and real-time decision-making introduce the need for adaptive hybrid models that can continuously update their symbolic reasoning frameworks while maintaining computational efficiency. Addressing these challenges is essential to making hybrid AI viable for real-world personalization systems across various domains, including healthcare, education, and entertainment.

The HyPer workshop aims to bridge the gap between sub-symbolic learning (e.g., neural networks) and symbolic knowledge representations (e.g., knowledge graphs, ontologies, logic-based models) to develop hybrid user models that better reflect human cognitive processes, social behaviors, and decision-making patterns.

HyPer is particularly timely, as it addresses challenges in designing AI-driven personalization systems that align with human cognitive processes while maintaining the algorithmic performance of machine learning-based approaches. It complements the main UMAP conference by offering a dedicated space to explore hybrid AI models for personalization, fostering discussions on novel hybrid AI methods, knowledge graphs, and cognitive frameworks, thereby advancing human-centric personalization.

Relevant topics include, but are not limited to:

- Methods for integrating symbolic knowledge and sub-symbolic learning in user modeling
- Applications of cognitive theories and behavioral insights in hybrid AI models for personalization

Permission to make digital or hard copies of all or part 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 components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
Conference'17, Washington, DC, USA

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-XXXX-X/2018/06
<https://doi.org/10.1145/nnnnnnn>

- Techniques for interpretability, explainability, and trust in hybrid AI systems
- Methods for detecting and mitigating biases and unfairness in hybrid AI, e.g., by using symbolic approaches such as counterfactual fairness
- Behavioral data analysis and user studies of cognitively-inspired modeling approaches
- Domain-specific implementations of hybrid AI models (e.g., e-learning, healthcare, finance, music)

2 Workshop Format, Intended Audience, and Participant Selection

The HyPer workshop is a half-day workshop in the format of a mini-conference, including a keynote to kick-start the discussion and a panel discussion. The workshop targets researchers and practitioners from academia and industry who are interested in using symbolic and sub-symbolic AI methods for user modeling and engaging in an interdisciplinary exchange between experts in various fields. We anticipate approximately 20 to 30 participants, with participant selection based on their submission of research papers, extended abstracts, or position papers relevant to the workshop topics. Furthermore, we will also encourage the submission of applications or statements of interest for interested individuals who may not have a formal paper submission but aim to contribute to discussions.

3 Previous Editions

This will be the first edition of the International Workshop on Hybrid AI for Human-Centric Personalization (HyPer). While previous workshops at UMAP have explored related topics, such as user-centered AI (UCAI), fairness (FairUMAP), and explainable user models (ExUM), our workshop is unique in its focus on the intersection of hybrid AI, cognitive modeling techniques, and personalization.

4 Biographies of Organizers

Elisabeth Lex (in-person attendance: sure) is an associate professor at Graz University of Technology, Austria. Besides, she is the head of the AI for Society Lab at the Institute of Human-Centred Computing. Her research interests include user modeling, recommender systems, information retrieval, and data science, with a particular focus on psychology-informed recommender systems [8, 9, 11, 16, 18, 24, 29], responsible recommender systems [12, 14, 22, 23] or music consumption [10]. She has (co-)authored more than 160 peer-reviewed publications on the aforementioned topics. She frequently takes on the role of a track chair at distinguished conferences such as *ACM Web Conference* or *ACM UMAP*. Besides, she regularly organizes workshops, such as the series *Workshop on Interfaces and Human Decision Making for Recommender Systems (IntRS)* [1]. She is a key researcher in the Cluster of Excellence Bilateral AI¹, funded by the Austrian Science Fund (FWF), where she leads a work package on combining symbolic and sub-symbolic AI to create novel user modeling and recommender systems.

Kevin Innerebner (in-person attendance: sure) is a PhD student at

Graz University of Technology and a member of the AI for Society Lab. His research focuses on combining symbolic and sub-symbolic AI techniques to design explainable recommender systems.

Marko Tkalcic (in-person attendance: sure) is a full professor at the University of Primorska, Slovenia, specializing in affective computing, user modeling, psychology-informed user modeling, and hybrid AI methods. He has published extensively on the intersection of emotion recognition, personality, and personalization [4, 5, 32, 34–36], as well as on the intersection of psychological models and recommender systems [24, 33]. He also regularly organizes workshops on theory-guided personalization and user modeling. Additionally, he has served as program chair at UMAP 2021 as well as other chair roles in conference (UMAP, IUI, RecSys). **Dominik Kowald** (in-person attendance: likely) is a research area manager at Know Center Research GmbH for the FAIR-AI research area. Additionally, he acts as a senior researcher and lecturer at Graz University of Technology, Austria. His research focuses on establishing trustworthy and reproducible AI [13, 30], privacy, bias, and fairness in recommender systems [10, 12, 22, 23], and cognitive user models [8, 9, 11, 16, 18, 29]. His research on fairness in AI and bias in recommender systems was awarded the Mind-the-gap gender and diversity award of Graz University of Technology in 2022 and was presented in several news outlets. He has served as an organizer for several interdisciplinary workshops in the field of personalization and AI. Additionally, he regularly acts as a program committee member of related conferences and workshops, as well as a topic editor for journals in the field.

Markus Schedl (in-person attendance: sure) is a full professor at the Johannes Kepler University Linz (JKU), where he leads the Multimedia Mining and Search group. In addition, he heads the Human-centered AI group at the Linz Institute of Technology (LIT) AI Lab. His research interests include recommender systems, user modeling, information retrieval, machine learning, multimedia, and trustworthy AI, with a particular focus on detecting and mitigating bias in retrieval and recommendation algorithms [15, 20, 21, 25] and on psychological models for recommendation [16, 18, 24]. He (co-)authored more than 250 refereed conference papers, journal articles, and book chapters. He served as a program chair for the International Society for Music Information Retrieval (ISMIR) Conference 2020 and will serve as a general chair of ACM Recommender Systems (RecSys) 2025. He has extensive experience in organizing workshops, having served as co-organizer of the workshop series *Theory-Informed User Modeling for Tailoring and Personalizing Interfaces (HUMANIZE)* [6], *Intelligent Music Interfaces for Listening and Creation (MILC)* [7], *Advances in Music Information Research (Ad-MIRe)* [27], *Search and Mining User-generated Contents (SMUC)* [2], and *Social Media Retrieval and Analysis (SoMeRA)* [28]. He has given numerous tutorials in top venues including ACM RecSys (2018 on “New Paths in Music Recommender Systems Research”), ACM SIGIR (2013 on “Music Similarity and Retrieval” and 2015 on “Music Retrieval and Recommendation”), ACM Multimedia (2013 on “Multimedia Information Retrieval: Music and Audio”), and the World Wide Web conference (2018 on “Complex Recommendations”).

¹<https://www.bilateral-ai.net/home>

Acknowledgments

This research was funded in whole or in part by the Austrian Science Fund (FWF): 10.55776/COE12, Cluster of Excellence [Bilateral Artificial Intelligence](#), the Austrian FFG COMET program, and the FFG HybridAir project (project #FO999902654).

References

- [1] Peter Brusilovsky, Marco de Gemmis, Alexander Felfernig, Elisabeth Lex, Pasquale Lops, Giovanni Semeraro, and Martijn C. Willemsen. 2021. Joint Workshop on Interfaces and Human Decision Making for Recommender Systems (IntRS'21). In *Fifteenth ACM Conference on Recommender Systems* (Amsterdam, Netherlands) (RecSys '21), 783–786. doi:10.1145/3460231.3470927
- [2] Iván Cantador, Francisco M. Carrero, José Carlos Cortizo, Paolo Rosso, Markus Schedl, and José A. Troyano (Eds.). 2011. *Proceedings of the 3rd International CIKM Workshop on Search and Mining User-Generated Contents, SMUC 2011, Glasgow, United Kingdom, October 28, 2011*. ACM. doi:10.1145/2065023
- [3] Brandon C Colelough and William Regli. 2025. Neuro-Symbolic AI in 2024: A Systematic Review. *arXiv preprint arXiv:2501.05435* (2025).
- [4] Bruce Ferwerda and Marko Tkalcić. 2018. Predicting users' personality from instagram pictures: Using visual and/or content features?. In *Proceedings of the 26th Conference on User Modeling, Adaptation and Personalization*. 157–161.
- [5] Bruce Ferwerda, Marko Tkalcić, and Markus Schedl. 2017. Personality Traits and Music Genres: What Do People Prefer to Listen To?. In *Proc. UMAP*.
- [6] Mark P. Graus, Bruce Ferwerda, Markus Schedl, Marko Tkalcić, Martijn C. Willemsen, and Panagiotis Germanakos (Eds.). 2017. *Proceedings of the 2017 ACM Workshop on Theory-Informed User Modeling for Tailoring and Personalizing Interfaces, HUMANIZE@UI 2017, Limassol, Cyprus, March 13, 2017*. ACM. doi:10.1145/3039677
- [7] Peter Knees, Markus Schedl, and Rebecca Fiebrink. 2019. Intelligent music interfaces for listening and creation. In *Proceedings of the 24th International Conference on Intelligent User Interfaces: Companion, Marina del Ray, CA, USA, March 16–20, 2019*. ACM, 135–136. doi:10.1145/3308557.3313110
- [8] Simone Kopeinik, Dominik Kowald, Ilire Hasani-Mavriqi, and Elisabeth Lex. 2017. Improving Collaborative Filtering Using a Cognitive Model of Human Category Learning. *Journal of Web Science* 4, 2, 45–61.
- [9] Dominik Kowald and Elisabeth Lex. 2016. The influence of frequency, recency and semantic context on the reuse of tags in social tagging systems. In *Proc. of ACM HT*. 237–242.
- [10] Dominik Kowald, Peter Muellner, Eva Zangerle, Christine Bauer, Markus Schedl, and Elisabeth Lex. 2021. Support the underground: characteristics of beyond-mainstream music listeners. *EPJ Data Science* 10, 1, 1–26.
- [11] Dominik Kowald, Subhash Chandra Pujari, and Elisabeth Lex. 2017. Temporal effects on hashtag reuse in twitter: A cognitive-inspired hashtag recommendation approach. In *Proceedings of the 26th International Conference on World Wide Web*. 1401–1410.
- [12] Dominik Kowald, Markus Schedl, and Elisabeth Lex. 2020. The Unfairness of Popularity Bias in Music Recommendation: A Reproducibility Study. In *European Conference on Information Retrieval*. 35–42.
- [13] Dominik Kowald, Sebastian Scher, Viktoria Pammer-Schindler, Peter Müllner, Kerstin Waxnegger, Lea Demelius, Angela Fessl, Maximilian Toller, Inti Gabriel Mendoza Estrada, Ilija Šimić, et al. 2024. Establishing and evaluating trustworthy AI: overview and research challenges. *Frontiers in Big Data* 7 (2024), 1467222.
- [14] Emanuel Lacić, Dominik Kowald, Matthias Traub, Granit Luzhnica, Joerg Simon, and Elisabeth Lex. [n. d.]. Tackling Cold-Start Users in Recommender Systems with Indoor Positioning Systems.
- [15] Oleg Lesota, Alessandro Melchiorre, Navid Rekabsaz, Stefan Brandl, Dominik Kowald, Elisabeth Lex, and Markus Schedl. 2021. Analyzing item popularity bias of music recommender systems: are different genders equally affected?. In *Proceedings of RecSys*. 601–606.
- [16] Elisabeth Lex, Dominik Kowald, and Markus Schedl. 2020. Modeling Popularity and Temporal Drift of Music Genre Preferences. *Transactions of ISMIR* 3, 1, 17–30.
- [17] Elisabeth Lex, Dominik Kowald, Paul Seitlinger, Thi Ngoc Trang Tran, Alexander Felfernig, and Markus Schedl. 2021. Psychology-informed Recommender Systems. *Foundations and Trends® in Information Retrieval* 15, 2 (2021), 134–242. <https://elizabethlex.info/docs/2021fntir-psychology.pdf>
- [18] Elisabeth Lex, Dominik Kowald, Paul Seitlinger, Thi Ngoc Trang Tran, Alexander Felfernig, and Markus Schedl. 2021. Psychology-informed Recommender Systems. *Found. Trends Inf. Retr.* 15, 2, 134–242.
- [19] Gary Marcus. 2020. The next decade in AI: four steps towards robust artificial intelligence. *arXiv preprint arXiv:2002.06177* (2020).
- [20] Alessandro B. Melchiorre, Navid Rekabsaz, Emilia Parada-Cabaleiro, Stefan Brandl, Oleg Lesota, and Markus Schedl. 2021. Investigating gender fairness of recommendation algorithms in the music domain. *Inf. Process. Manag.* 58, 5.
- [21] Alessandro B. Melchiorre, Eva Zangerle, and Markus Schedl. 2020. Personality Bias of Music Recommendation Algorithms. In *Proc. of ACM RecSys*. 533–538.
- [22] Peter Muellner, Dominik Kowald, and Elisabeth Lex. 2021. Robustness of Meta Matrix Factorization Against Strict Privacy Constraints. In *Proc. of ECIR*. 107–119.
- [23] Peter Müllner, Elisabeth Lex, Markus Schedl, and Dominik Kowald. 2023. ReuseKNN: Neighborhood Reuse for Differentially Private KNN-Based Recommendations. *ACM Transactions on Intelligent Systems and Technology* 14, 5 (2023), 1–29.
- [24] Markus Reiter-Haas, Emilia Parada-Cabaleiro, Markus Schedl, Elham Motamedi, Marko Tkalcić, and Elisabeth Lex. 2021. Predicting Music Relisting Behavior Using the ACT-R Framework. In *Proc. of ACM RecSys*. 702–707.
- [25] Navid Rekabsaz, Simone Kopeinik, and Markus Schedl. 2021. Societal Biases in Retrieved Contents: Measurement Framework and Adversarial Mitigation of BERT Rankers. In *Proc. of ACM SIGIR*. 306–316.
- [26] Markus Schedl, Vito Walter Anelli, and Elisabeth Lex. 2024. Trustworthy User Modeling and Recommendation From Technical and Regulatory Perspectives. In *Adjunct Proceedings of the 32nd ACM Conference on User Modeling, Adaptation and Personalization*. 17–19.
- [27] Markus Schedl, Peter Knees, and Öscar Celma. 2012. Session Details: AdMIRE'12 Workshop 6. In *Proceedings of the 21st International Conference on World Wide Web* (Lyon, France) (WWW '12 Companion). ACM, New York, NY, USA. doi:10.1145/3259172
- [28] Markus Schedl, Peter Knees, and Jialie Shen. 2014. SoMeRA 2014: social media retrieval and analysis workshop. In *The 37th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '14, Gold Coast, QLD, Australia - July 06 - 11, 2014*, Shlomo Geva, Andrew Trotman, Peter Bruza, Charles L. A. Clarke, and Kalervo Järvelin (Eds.). ACM, 1297. doi:10.1145/2600428.2600739
- [29] Paul Seitlinger, Dominik Kowald, Simone Kopeinik, Ilire Hasani-Mavriqi, Tobias Ley, and Elisabeth Lex. 2015. Attention please! a hybrid resource recommender mimicking attention-interpretation dynamics. In *Proc. of WWW*. 339–345.
- [30] Harald Semmelrock, Tony Ross-Hellauer, Simone Kopeinik, Dieter Theiler, Armin Haberl, Stefan Thalmann, and Dominik Kowald. 2024. Reproducibility in Machine Learning-based Research: Overview, Barriers and Drivers. *arXiv preprint arXiv:2406.14325* (2024).
- [31] Giuseppe Spillo, Cataldo Musto, Marco de Gemmis, Pasquale Lops, and Giovanni Semeraro. 2024. Recommender systems based on neuro-symbolic knowledge graph embeddings encoding first-order logic rules. *User Modeling and User-Adapted Interaction* (2024), 1–45.
- [32] Marko Tkalcić and Li Chen. 2022. Personality and Recommender Systems. In *Recommender Systems Handbook*, Francesco Ricci, Lior Rokach, and Bracha Shapira (Eds.). Springer US, 757–787. doi:10.1007/978-1-0716-2197-4_20
- [33] Marko Tkalcić and Bruce Ferwerda. 2018. Theory-driven Recommendations: Modeling Hedonic and Eudaimonic Movie Preferences. In *Proceedings of the 9th Italian Information Retrieval Workshop, Rome, Italy, May, 28–30, 2018*. <http://ceur-ws.org/Vol-2140/paper7.pdf>
- [34] Marko Tkalcić, Andrej Kosir, and Jurij Tasic. 2011. Affective recommender systems: the role of emotions in recommender systems. In *Proceedings of the RecSys 2011 Workshop on Human Decision Making in Recommender Systems and User-Centric Evaluation of Recommender Systems and Their Interfaces*. 9–13.
- [35] Marko Tkalcić, Jurij Tasic, and Andrej Košir. 2009. Emotive and personality parameters in multimedia recommender systems. *ACII 2009 Affective Computing and Intelligent Interaction* 33 (2009).
- [36] Marko Tkalcić. 2018. Emotions and Personality in Recommender Systems: Tutorial. In *Proceedings of the 12th ACM Conference on Recommender Systems (RecSys '18)*. ACM, 535–536. doi:10.1145/3240323.3241619