Hybrid AI for Human-Centric Personalization (HyPer)

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

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

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

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

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

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 human factors such as biases, mental models, and affective states. The HyPer workshop provides a venue to discuss how hybrid AI approaches, combining neural architectures, symbolic representations, and cognitive/behavioral frameworks, can bridge the gap between explainability, cognitive modeling, and automated adaptation to user preferences.

CCS Concepts

• **Computing methodologies** → **Artificial intelligence**; • **Information systems** → **Decision support systems**.

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 *Adjunct Proceedings of the 33rd ACM Conference on User Modeling, Adaptation and Personalization (UMAP Adjunct '25), June 16–19, 2025, New York City, NY, USA.* ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/3708319.3727563

1 Introduction and Workshop Goals

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 [30]. Hybrid AI models integrate symbolic reasoning with sub-symbolic, data-driven learning [4], and have emerged as a promising solution for advancing human-centric personalization by combining structured knowledge representations with adaptive, data-driven models [34]. Corresponding hybrid AI systems can integrate 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 [20]. Ensuring fairness, mitigating biases, and aligning hybrid models with human cognitive processes poses another challenge [19]. 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.

The HyPer workshop aims to bridge the gap between sub-symbolic learning (e.g., neural networks) and symbolic knowledge representations (e.g., knowledge graphs) to develop hybrid user models that better reflect human cognitive processes, social behaviors, and decision-making patterns. Additionally, HyPer 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. Summed up, this workshop includes eight papers that contribute to the following research topics:

- 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
- Techniques for interpretability, explainability, and trust in hybrid AI systems

UMAP Adjunct '25, New York City, NY, USA

ACM ISBN 979-8-4007-1399-6/25/06
https://doi.org/10.1145/3708319.3727563

- 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

2 Workshop Summary and Accepted Contributions

We accepted eight papers and grouped them into three themes.

2.1 Cognitive Architectures and Neuro-Symbolic Approaches for Transparent Hybrid AI

Four papers in this workshop emphasize the need for hybrid AI systems that reflect human cognitive processes and offer greater transparency. One contribution critiques the procedural limitations of current large language models and proposes augmenting them with semantic memory and associative learning to better handle complex, dynamic environments [42]. Another paper presents fuzzy neural networks as a neuro-symbolic approach for recommender systems, enabling human-readable logic rules and explainable decision-making [1]. A vision paper builds on cognitive science by integrating declarative and procedural memory from the ACT-R architecture to simulate human-like decision-making and personalization [7]. Finally, a position paper advocates for cognitively aligned human-AI decision-making, proposing interactive AI assistance that supports domain experts in co-constructing decisions by combining symbolic reasoning with sub-symbolic learning [27].

2.2 Enhancing Personalization and Recommendation through Automation and Cognition in Hybrid AI

Two contributions explore the intersection of automation and cognitive alignment in recommender systems and user modeling within hybrid AI. One study benchmarks a wide range of AutoML and AutoRecSys libraries on multiple datasets, revealing that AutoML approaches often outperform traditional recommender systems,especially for non-expert users [40]. Another study investigates exploratory search behavior in e-commerce, identifying distinct user orientations and linking them to specific search stages [9].

2.3 Trust, Interaction, and Interpretability in Human-Hybrid AI Collaboration

Two papers address how users interact with hybrid AI systems and how transparency can influence trust and effectiveness. One study introduces a confidence rating interface for chatbot interactions, suggesting that visualizing model certainty and offering prompt improvement suggestions can improve trust and performance in verification and reasoning tasks [41]. Another paper proposes a symbolic motion representation system based on Labanotation and Laban Movement Analysis, aimed at improving human-robot interaction and preserving expressive, interpretable movement data [26].

3 Biographies of Organizers

Elisabeth Lex is a full 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 focus on psychology-informed recommender systems [3, 10, 11, 13, 18, 19, 28, 31], responsible recommender systems [14, 16, 23, 24] or music [12, 25]. She regularly organizes workshops, research tracks, and gives tutorials at the core venues in her field.

Kevin Innerebner is a PhD student at Graz University of Technology and a member of the AI for Society Lab. His research combines symbolic and sub-symbolic AI techniques to design information retrieval and recommender systems [2, 8].

Marko Tkalcic 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 [5, 6, 35, 37–39], as well as on the intersection of psychological models and recommender systems [28, 36]. 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 conferences (UMAP, IUI, RecSys).

Dominik Kowald is a research area manager at Know Center Research GmbH for the FAIR-AI research area. He is also a senior researcher and lecturer at Graz University of Technology, Austria. His research focuses on establishing trustworthy and reproducible AI [15, 32, 33], privacy, bias, and fairness in recommender systems [12, 14, 23, 24], and cognitive user models [10, 11, 13, 18, 19, 31]. He has organized several interdisciplinary workshops on personalization and AI. Additionally, he regularly serves on the program committees of related conferences and workshops and is a topic editor for journals in the field.

Markus Schedl is a full professor at Johannes Kepler University Linz (JKU), leading the Multimedia Mining and Search group. He also 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, natural language processing, multimedia, and trustworthy AI, with a focus on detecting and mitigating bias in retrieval and recommendation algorithms [17, 21, 22, 29] and on psychological models for recommendation [18, 19, 28]. Markus is general co-chair of the ACM Conference on Recommender Systems (RecSys) 2025.

4 Conclusion

The eight contributions to the HyPer workshop demonstrate how hybrid AI, combining symbolic reasoning, sub-symbolic learning, and cognitive modeling, can advance explainable and human-centered personalization. They highlight promising directions for building adaptive systems that align with human reasoning in a cognitive-plausible and transparent way.

Acknowledgments

This research is 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, the FFG HybridAir project (project #FO999902654), the University of Primorska grant CogniCom, and the Slovenian Research and Innovation Agency through grants BI-NO/25-27-007 and N2-0354 (C).

References

[1] Stephan Bartl, Kevin Innerebner, and Elisabeth Lex. 2025. Differentiable Fuzzy Neural Networks for Recommender Systems. In *UMAP Adjunct '25*. ACM, New York, NY, USA.

[2] Daniel Borst, Stefan Sobernig, and Kevin Innerebner. 2024. Protocol for a Systematic Literature Review on the Evaluation of Code Recommender Systems. (2024).

[3] 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

[4] Brandon C Colelough and William Regli. 2025. Neuro-Symbolic AI in 2024: A Systematic Review. *arXiv preprint arXiv:2501.05435* (2025).

[5] Bruce Ferwerda and Marko Tkalcic. 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.

[6] Bruce Ferwerda, Marko Tkalcic, and Markus Schedl. 2017. Personality Traits and Music Genres: What Do People Prefer to Listen To?. In *Proc. UMAP*.

[7] Kevin Innerebner, Dominik Kowald, Markus Schedl, and Elisabeth Lex. 2025. Hybrid Personalization Using Declarative and Procedural Memory Modules of the Cognitive Architecture ACT-R. In *UMAP Adjunct '25*. ACM, New York, NY, USA.

[8] Benedikt Kantz, Kevin Innerebner, Peter Waldert, Stefan Lengauer, Elisabeth Lex, and Tobias Schreck. 2025. OnSET: Ontology and Semantic Exploration Toolkit. In *SIGIR Demo Paper Track 2025*.

[9] Eunhye Kim, Kiroong Choe, Guangjing Yan, and Mingyu Kang. 2025. User Orientations and Stage-Specific Behaviors in E-commerce Exploratory Search - A Formative Study. In *UMAP Adjunct '25*. ACM, New York, NY, USA.

[10] 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.

[11] 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.

[12] 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.

[13] 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*.

[14] 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.

[15] Dominik et al. Kowald. 2024. Establishing and evaluating trustworthy AI: overview and research challenges. *Frontiers in Big Data* 7 (2024), 1467222.

[16] Emanuel Lacic, Dominik Kowald, Matthias Traub, Granit Luzhnica, Jörg Peter Simon, and Elisabeth Lex. 2015. Tackling Cold-Start Users in Recommender Systems with Indoor Positioning Systems. In *9th ACM RecSys 2015*.

[17] 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.

[18] Elisabeth Lex, Dominik Kowald, and Markus Schedl. 2020. Modeling Popularity and Temporal Drift of Music Genre Preferences. *Transactions of ISMIR* 3, 1, 17–30.

[19] 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.

[20] Gary Marcus. 2020. The next decade in AI: four steps towards robust artificial intelligence. *arXiv preprint arXiv:2002.06177* (2020).

[21] 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.

[22] Alessandro B. Melchiorre, Eva Zangerle, and Markus Schedl. 2020. Personality Bias of Music Recommendation Algorithms. In *Proc. of ACM RecSys*. 533–538.

[23] Peter Muellner, Dominik Kowald, and Elisabeth Lex. 2021. Robustness of Meta Matrix Factorization Against Strict Privacy Constraints. In *Proc. of ECIR*. 107–119.

[24] 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).

[25] Emilia Parada-Cabaleiro, Maximilian Mayerl, Stefan Brandl, Marcin Skowron, Markus Schedl, Elisabeth Lex, and Eva Zangerle. 2024. Song lyrics have become simpler and more repetitive over the last five decades. *Scientific Reports* 14, 1 (2024), 5531.

[26] Roberto Perez-Martinez, Alberto Casas-Ortiz, and Olga C. Santos. 2025. Towards a hybrid framework for Motion Representation and Analysis based on Labanotation and LMA. In *UMAP Adjunct '25*. ACM, New York, NY, USA.

[27] Muhammad Raees, Vassilis-Javed Khan, and Konstantinos Papangelis. 2025. Building Human-AI Reliance Through Cognitive Engagement and Exploratory AI Assistance. In *UMAP Adjunct '25*. ACM, New York, NY, USA.

[28] Markus Reiter-Haas, Emilia Parada-Cabaleiro, Markus Schedl, Elham Motamedi, Marko Tkalcic, and Elisabeth Lex. 2021. Predicting Music Relistening Behavior Using the ACT-R Framework. In *Proc. of ACM RecSys*. 702–707.

[29] 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.

[30] 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.

[31] 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.

[32] Harald Semmelrock, Simone Kopeinik, Dieter Theiler, Tony Ross-Hellauer, and Dominik Kowald. 2023. Reproducibility in machine learning-driven research. *arXiv preprint arXiv:2307.10320* (2023).

[33] 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).

[34] 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. *UMUAI* (2024), 1–45.

[35] Marko Tkalcic 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

[36] Marko Tkalcic 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.*

[37] Marko Tkalcic, 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.

[38] Marko Tkalcic, Jurij Tasic, and Andrej Košir. 2009. Emotive and personality parameters in multimedia recommender systems. *ACII 2009 Affective Computing and Intelligent Interaction* 33 (2009).

[39] Marko Tkalčič. 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

[40] Tobias Vente and Joeran Beel. 2025. The Potential of AutoML for Recommender Systems. In *UMAP Adjunct '25*. ACM, New York, NY, USA.

[41] Lifei Wang, Natalie Friedman, Chengchao Zhu, Zeshu Zhu, and S. Joy Mountford. 2025. The Impact of Confidence Ratings on User Trust in Large Language Models. In *UMAP Adjunct '25*. ACM, New York, NY, USA.

[42] Schaun Wheeler and Olivier Jeunen. 2025. Procedural Memory is not All You Need: Bridging Cognitive Gaps in LLM-Based Agents. In *UMAP Adjunct '25*. ACM, New York, NY, USA.