

PALS 2025

**Tailoring AI: Exploring Active and Passive LLM
Personalization**

Proceedings of the Workshop

November 9, 2025

The PALS organizers gratefully acknowledge the support from the following sponsors.

Apple

©2025 Empirical Methods in Natural Language Processing (EMNLP)

Order copies of this and other ACL proceedings from:

Association for Computational Linguistics (ACL)
317 Sidney Baker St. S
Suite 400 - 134
Kerrville, TX 78028
USA
Tel: +1-855-225-1962
acl@aclweb.org

ISBN N/A

Introduction

Large language models (LLMs) have demonstrated remarkable capabilities in various NLP tasks. However, the extent to which these models can and should adapt to individual users' needs remains an open question. Therefore, this workshop will focus on the personalization of LLMs to meet individual user's needs and preferences. From the user's perspective the personalization process can be either passive, where the LLM learns from observing user behavior, or active, where the user directly guides the personalization. The questions of how to personalize effectively, when to personalize, which personalization paradigm to apply (active vs. passive) remain open questions that are important to address. Further, active and passive personalization each have several challenges and open questions in their own right. Successfully answering these kinds of questions inherently requires an interdisciplinary approach, combining expertise from a wide variety of fields, such as: NLP, human-computer interaction, linguistics, cognitive science, behavioral science, psychology, ethics, etc. Thus, the workshop will promote interdisciplinary research necessary for effective, user-first personalization, driving towards solving the challenges of passive and active personalization.

This workshop invites submissions related, but not limited, to the following topics in the space of active and passive LLM personalization:

- Algorithmic approaches (user persona/profile creation, in-context learning, RAG, parameter tuning, user representation learning, user-specific reasoning, etc.)
- Evaluation methods for offline and online use cases
- Potential risks of personalized LLMs (creating an echo chamber, limiting user exploration)
- User studies or surveys to understand when, where, and how active vs. passive personalization is effective
- Dataset creation and curation (collecting and curating a dataset from real users, creating and validating a synthetic dataset, identifying salient samples for personalization, etc.)

Organizing Committee

Organizers

Andrew Silva, Apple
Lucie Charlotte Magister, University of Cambridge
Katherine Metcalf, Apple
Clemencia Siro, CWI
Maartje Ter Hoeve, Apple

Program Committee

Program Chairs

Katherine Metcalf, Apple
Lucie Charlotte Magister, University of Cambridge
Andrew Silva, Apple
Clemencia Siro, CWI
Maartje Ter Hoeve, Apple

Reviewers

Tolulope Olalekan Abiola, Imran Ahsan, Talha Rüzgar Akkuş, Arnav Arora, Reshma Ashok

Wanchen Cao

Jonathan DeCastro, Maureen de Seyssel

Masha Fedzechkina, Yahui Fu, Michael Färber

Deepak Edakkattil Gopinath, Zijin Gu

Tiancheng Hu

Xia Jiang

Mucheol Kim, Boshko Koloski

Daniel Lee

Ercong Nie

Vishakh Padmakumar, Zi Haur Pang, DongKyu Park, Maximilian Puelma Touzel

Jerome Ramos

Skyler Seto, Li Siyan

Pradyumna Tambwekar, Barry-John Theobald

Yaixin Wu

Yizhou Xu

Hiroaki Yamada, Zachary Yang, Hyunwook Yu, Shuzhou Yuan

Zihe Zhang, Kai Zhao

Keynote Talk

Rethinking Personalization and User Control in LLMs

Diyi Yang
Stanford University

Abstract: The growing capabilities of LLMs offer significant opportunities for personalized assistance, but a critical view is needed on how to achieve personalization in a generalized way without being constrained by specific contexts, and how much control end users can and should have. This talk examines how AI systems can be shaped around you, rather than enforcing a one-size-fits-all approach. The first part presents Generative User Modeling (GUM), which demonstrates how a general user model can infer preferences and beliefs from unstructured user interactions, such as recorded videos of people’s screen use. The second part introduces Knoll, a knowledge-module ecosystem that allows end users to curate, configure, and control what the model “knows,” and demonstrates its effectiveness in personalized recommendations, advice-seeking, and writing assistance. We conclude by reflecting on how future LLM agents should balance personalization with transparency, and how they can empower users to steer, correct, and control their LLM systems.

Bio: Diyi Yang is an assistant professor in the Computer Science Department at Stanford University, also affiliated with the Stanford NLP Group, Stanford HCI Group and Stanford Human Centered AI Institute. Her research focuses on human-centered natural language processing and human-AI interaction. She is a recipient of Microsoft Research Faculty Fellowship (2021), NSF CAREER Award (2022), and a Sloan Research Fellowship (2024). Her work has received multiple paper awards or nominations at top NLP and HCI conferences.

Keynote Talk

Personalising Language Models for the Most of the World: Lessons from Low-Resource Contexts

Vukosi Marivate
University of Pretoria

Abstract: Most personalisation research assumes abundant data and well-resourced languages. This reason challenges that paradigm, showing how personalisation in low-resource contexts requires rethinking what user adaptation means when users, data, and languages are underrepresented. Drawing from African NLP and the work of African initiatives, the talk discusses how personalisation can become a vehicle for inclusion rather than reinforcement of global data imbalances.

Bio: Prof Vukosi Marivate is a Professor of Computer Science and holds the ABSA UP Chair of Data Science at the University of Pretoria. He specialises in developing Machine Learning (ML) and Artificial Intelligence (AI) methods to extract insights from data, with a particular focus on the intersection of ML/AI and Natural Language Processing (NLP). His research is dedicated to improving the methods, tools and availability of data for local or low-resource languages. As the leader of the Data Science for Social Impact research group in the Computer Science department, Vukosi is interested in using data science to solve social challenges. He has worked on projects related to science, energy, public safety, and utilities, among others. Prof Marivate is a co-founder of Lelapa AI, an African startup focused on AI for Africans by Africans. Vukosi is a chief investigator on the Masakhane Research Foundation, which aims to develop NLP technologies for African languages. Vukosi is also a co-founder of the Deep Learning Indaba, the leading grassroots Machine Learning and Artificial Intelligence conference on the African continent that aims to empower and support African researchers and practitioners in the field.

Keynote Talk
**Tailoring LLMs for lesser-used language users: challenges
and opportunities**

Federico Gobbo
University of Amsterdam

Abstract: The multiplicity in the languages of the world is only partially represented in its digital counterparts, creating new forms of inequalities and digital divides. In particular, LLMs rely on big data and only a few languages in the world provide sufficient data to train them adequately. This talk presents a case study of using Esperanto, a lesser-used language with a relatively small community but widespread global presence, within the context of LLMs from the perspective of its users. Drawing from the case study, the talk will present challenges and opportunities for lesser-used languages in general.

Bio: Federico Gobbo is full professor of Interlinguistics and Esperanto at the University of Amsterdam. His research interests are at the crossroads of Linguistics and Computer Science, in particular in human-machine communication using lesser-used languages.

Keynote Talk

Hannah Rose Kirk
University of Oxford

Bio: Hannah Rose Kirk currently researches AI alignment in her PhD at the University of Oxford and as a Research Scientist at the UK AI Safety Institute. Her current research centres on human-and-model-in-the-loop feedback and data-centric alignment of AI. She is passionate about the societal impact of AI systems as we scale across model capabilities, domains and human populations.

Program

Sunday, November 9, 2025

09:00 - 09:45 *Invited Talk: Diyi Yang*

09:45 - 10:30 *Oral Presentations*

One-Topic-Doesn't-Fit-All: Transcreating Reading Comprehension Test for Personalized Learning

Jieun Han, Daniel Lee, Haneul Yoo, Jinsung Yoon, Junyeong Park, Suin Kim, So-Yeon Ahn and Alice Oh

Enhancing Rating Prediction with Off-the-Shelf LLMs Using In-Context User Reviews

Koki Ryu and Hitomi Yanaka

Modeling Layered Consciousness with Multi-Agent Large Language Models

Sang Hun Kim, Jongmin Lee, DongKyu Park, So Young Lee and Yosep Chong

10:30 - 11:00 *Coffee Break*

11:00 - 12:30 *Poster Session*

PediaMind-R1: A Temperament-Aware Language Model for Personalized Early Childhood Care Reasoning via Cognitive Modeling and Preference Alignment

Zihe Zhang, Can Zhang and XU Yanheng

PromptTailor: Multi-turn Intent-Aligned Prompt Synthesis for Lightweight LLMs

Yizhou Xu and Janet Davis

Enhancing Rating Prediction with Off-the-Shelf LLMs Using In-Context User Reviews

Koki Ryu and Hitomi Yanaka

Modeling Layered Consciousness with Multi-Agent Large Language Models

Sang Hun Kim, Jongmin Lee, DongKyu Park, So Young Lee and Yosep Chong

Hateful Person or Hateful Model? Investigating the Role of Personas in Hate Speech Detection by Large Language Models

Shuzhou Yuan, Ercong Nie, Mario Tawfelis, Helmut Schmid, Hinrich Schuetze and Michael Färber

Sunday, November 9, 2025 (continued)

When Should Agents Ask? Decision-Theoretic Adaptive Communication for LLM Agents

Yijiang River Dong, Tiancheng Hu, Zheng Hui, Caiqi Zhang, Ivan Vulić, Andreea Bobu and Nigel Collier

One-Topic-Doesn't-Fit-All: Transcreating Reading Comprehension Test for Personalized Learning

Jieun Han, Daniel Lee, Haneul Yoo, Jinsung Yoon, Junyeong Park, Suin Kim, So-Yeon Ahn and Alice Oh

BluePrint : A Social Media User Dataset for LLM Persona Evaluation and Training

Aurélien Bück-Kaeffer, Je Qin Chooi, Dan Zhao, Maximilian Puelma Touzel, Kellin Pelrine, Jean-François Godbout, Reihaneh Rabbany and Zachary Yang

Augmenting Dialog with Think-Aloud Utterances for Modeling Individual Personality Traits by LLM

Seiya Ishikura, Hiroaki Yamada, Tatsuya Hiraoka, Hiroaki Yamada and Takenobu Tokunaga

Is Passive Expertise-Based Personalization Enough? A Case Study in AI-Assisted Test-Taking

Li Siyan, Jason Zhang, Akash V Maharaj, Yuanming Shi, Daniel Lee and Yunyao Li

Is Active Persona Inference Necessary for Aligning Small Models to Personal Preferences?

Zilu Tang, Afra Feyza Akyürek, Ekin Akyürek and Derry Tanti Wijaya

Minority-Aware Satisfaction Estimation in Dialogue Systems via Preference-Adaptive Reinforcement Learning

Yahui Fu, Zi Haur Pang and Tatsuya Kawahara

Collaborative User Prompt for Personalized Generative Recommendation

Jerome Ramos, Bin Wu and Aldo Lipani

Analyzing Trade-offs Between Faithfulness and Correctness in LLM Personalization

Tiasa Singha Roy and Vishakh Padmakumar

Personality Editing for Language Models through Relevant Knowledge Editing

Seojin Hwang, Yumin Kim, Byeongjeong Kim, Donghoon Shin and Hwanhee Lee

12:30 - 14:00

Lunch Break

Sunday, November 9, 2025 (continued)

14:00 - 14:45 *Invited Talk: Vukosi Marivate*

14:45 - 15:30 *Invited Talk: Federico Gobbo*

15:30 - 16:00 *Coffee Break*

16:00 - 16:45 *Invited Talk: Hannah Rose Kirk*

Author Index