WSDM 2024 Workshop on Large Language Models for Individuals, Groups, and Society

Michael Bendersky
bemike@google.com
Google Research

Cheng Li
chgli@google.com
Google Research

Qiaozhu Mei
qmei@umich.edu
University of Michigan

Vanessa Murdock
vmurdock@amazon.com
Amazon Research

Jie Tang
jietang@tsinghua.edu.cn
Tsinghua University

Hongning Wang
hw5x@virginia.edu
University of Virginia

Hamed Zamani
zamani@cs.umass.edu
University of Massachusetts Amherst

Mingyang Zhang
mingyang@google.com
Google Research

ABSTRACT

This workshop discusses the cutting-edge developments in research and applications of personalizing large language models (LLMs) and adapting them to the demands of diverse user populations and societal needs. The full-day workshop includes several keynotes and invited talks, a poster session and a panel discussion.

ACM Reference Format:
Michael Bendersky, Cheng Li, Qiaozhu Mei, Vanessa Murdock, Jie Tang, Hongning Wang, Hamed Zamani, and Mingyang Zhang. 2024. WSDM 2024 Workshop on Large Language Models for Individuals, Groups, and Society. In *Proceedings of the 17th ACM International Conference on Web Search and Data Mining (WSDM '24), March 4–8, 2024, Merida, Mexico.* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3616855.3635726

1 INTRODUCTION

The recent advancements in large language models (LLMs), such as GPT, PaLM, and Llama, along with the generative AI capabilities they possess, have garnered significant attention within both the research community and the public sphere. Although these models are easily accessible to users and researchers through conventional prompting interfaces, API calls, or static snapshots, there is an increasing demand for these models to provide personalized and context-aware responses. This requirement arises from diverse application scenarios where assistive creation and tailored generation are essential for individual and groups/sub-populations of users with even more diverse backgrounds and preferences. Merely relying on generic responses is insufficient in addressing the specific needs and constraints of users in personal, group, or even societal contexts. Instead, such scenarios demand the models' ability to consider and align their responses to the preferences and objectives of the users in these aforementioned contexts.

For example, it has been widely recognized that the success of many machine learning applications on the Web, such as search engines, recommender systems, and online advertising heavily relies on effective personalization, and tailoring their outputs to the unique user preferences and intents. As another example, in scenarios where LLMs are being used as productivity assistants, it is highly desirable for them to be attuned to the requirements of a specific enterprise, business role, and user communication style. Finally, recently LLMs have also been adapted for creating virtual personas that can provide companionship and entertainment to users that come from different demographics and backgrounds.

Making LLMs and their downstream applications personalized, or aligning them with individual and societal needs [1], presents significant new challenges for the research community. Addressing these challenges is vital for enhancing the user experience in applications and achieving optimal outcomes in diverse domains that are related to the community, and eventually delivering a positive long-term societal impact through these practices. As core members within the Web search and data mining community, we are eager to rally our fellow researchers and unite our efforts toward this pressing and crucial direction.

This workshop, as the first in a series, aims to create a collaborative and interdisciplinary platform that brings together creators, researchers, and practitioners of large language models. By fostering an open and forward-looking environment, the workshop seeks to facilitate discussions on the current landscape of personalizing LLMs, adapting LLMs to individual and group contexts, and aligning LLMs with the value and objectives of the society at large. It provides an opportunity for participants to share insights, exchange ideas, and explore innovative approaches in the field. The ultimate goal is to drive progress and shape the future of large language models for individuals, groups, and the society.

Topics covered in the call for papers include:

- Novel models and algorithms for adapting large language models to personal contexts.
- New developments in aligning large language models with the preferences and objectives of individuals, sub-populations, or the society at large.
- Theoretical and empirical results of applying reinforcement learning from the feedback of individuals and groups of human users to LLMs.
- Evaluation of personalization and societal alignment of LLMs, including datasets, metrics, and benchmarks.
- Personalizing and aligning LLMs under resource constraints. For example, deploying personalized LLMs on mobile devices or aligning the output of frozen LLMs through APIs.

WSDM '24, March 4–8, 2024, Merida, Mexico

ACM ISBN 979-8-4007-0371-3/24/03.
https://doi.org/10.1145/3616855.3635726

- Applications of personalization and societal-alignment of LLMs, such as search engines, recommender systems, email and writing assistants, social networking, entertainment, education, healthcare, scientific discovery, and future of work.
- Ethics of personalizing LLMs, such as privacy, fairness, bias, transparency, diversity, and other potential impacts of LLMs to individuals, groups, and the society.
- Equitable applications of LLM to diverse user groups.

2 ORGANIZERS

Michael Bendersky. Michael is a Principal Software Engineer / Engineering Director at Google Research. He is currently managing a team whose mission is improving algorithms, models, and metrics for information discovery and quality across Google products. His recent research interests include neural ranking and retrieval, query understanding, dynamic content understanding, personalization in search and recommendation systems, and more. Michael is a Distinguished Member of the ACM. He holds a Ph.D. from the University of Massachusetts Amherst, and a B.Sc. and M.Sc. from the Technion, Israel Institute of Technology. Michael co-authored over 80 publications. He served on program and organizing committees for multiple academic conferences, and co-organized tutorials and workshops at multiple venues including SIGIR, WSDM, and ICTIR. He co-authored two books in the "Foundations and Trends in Information Retrieval" series: "Information Retrieval with Verbose Queries", and "Search and Discovery in Personal Email Collections".

Cheng Li. Cheng is a Staff Software Engineer at Google Research. Cheng's research interests include understanding of queries and documents for information retrieval, understanding of creators on social platforms, and generation using large language models. Her areas of expertise include developing machine learning models to improve search quality by understanding queries and documents, learning creator representation with application to various downstream tasks, and personalized generation.

Qiaozhu Mei. Qiaozhu is s a professor in the School of Information and the Department of EECS at the University of Michigan. His research focuses on large-scale data mining, machine learning, information retrieval, and natural language processing, with broad applications. Qiaozhu is an ACM distinguished member (2017) and a recipient of the NSF Career Award (2011). His work has received multiple best paper awards at WWW, ICML, KDD, WSDM, and other major conferences. He served as the founding director of the master degree of applied data science at the University of Michigan. He has rich experience organizing workshops and related events, including being the General Co-Chair of SIGIR 2018.

Vanessa Murdock. Vanessa manages a science team in Alexa Shopping at Amazon, partnering with Amazon's Choice, Alexa Shopping List and others. She is Chair of the ACM Special Interest Group on Information Retrieval (SIGIR) and is serving as an Editor-In-Chief of the Journal of Information Retrieval. Her research spans a wide range of information retrieval and recommender systems topics, focusing on connecting people's online and offline lives. Vanessa has her PhD in computer science from the University of Massachusetts. She began her professional life as a classical pianist.

Jie Tang. Jie is a Professor of the Department of Computer Science and Technology of Tsinghua University. He is a Fellow of the ACM, a Fellow of AAAI, and a Fellow of the IEEE. His research interests include artificial general intelligence (AGI), data mining, social networks, machine learning and knowledge graph, with an emphasis on designing new algorithms for information and social network mining and designing new paradigms for artificial general intelligence. Similar to Open AI's GPT serials, Jie, leading a big research team, have designed GLM-130B, ChatGLM, CogView&CogVideo, CodeGeex, toward teaching machines to think like humans. Jie also invented AMiner.org, which has attracted millions of users from 220 countries/regions in the world. He has been honored with the SIGKDD Test-of-Time Award for Applied Science (Ten-year Best Paper Award), the 2nd National Award for Science&Technology, NSFC for Distinguished Young Scholar, UK Royal Society-Newton Advanced Fellowship Award, and SIGKDD Service Award. He served as PC Co-Chair of CIKM'16, WSDM'15, Associate General Chair of KDD'18, and the General Co-Chair of WWW'23.

Hongning Wang. Hongning is an associate professor at the University of Virginia. His research lies in the intersection among machine learning, data mining and information retrieval, with a special focus on sequential decision optimization. His work has generated over 100 research papers in top venues in data mining and information retrieval areas. He is a recipient of 2016 National Science Foundation CAREER Award, 2020 Google Faculty Research Award, and SIGIR'2019 Best Paper Award. He has rich experience organizing workshops and related events. He is the General Co-Chair of SIGIR 2024, and workshop co-chair of KDD 2022 and 2023.

Hamed Zamani. Hamed Zamani is an Assistant Professor at the University of Massachusetts Amherst, where he also serves as the Associate Director of the Center for Intelligent Information Retrieval (CIIR), one of the top academic research labs in Information Retrieval worldwide. Prior to UMass, he was a Researcher at Microsoft working on search and recommendation problems. His research focuses on designing and evaluating (interactive) information access systems, including search engines, recommender systems, and question answering. His work has led to over 85 refereed publications in the field and has developed the LaMP benchmark (https://lamp-benchmark.github.io/) for language model personalization. He is a recipient of the NSF CAREER Award, ACM SIGIR Early Career Excellence in Research and Community Engagement awards, and Amazon Research Award. He is an Associate Editor of the ACM Transactions on Information Systems (TOIS), has organized multiple workshops at SIGIR, RecSys, WSDM, and WWW conferences, and served as a PC Chair at SIGIR 2022 (Short Papers).

Mingyang Zhang. Mingyang is a Senior Staff Software Engineer, Tech Lead Manager at Google Research. His research interests include understanding search queries, documents and content creators, neural information retrieval and discovery, personalized generation approaches and metrics, etc. He leads a Google Research team which invents new technologies and applies them to major Google products. He holds a PhD in Computer Science from The George Washington University. He also won the best student paper of CIKM 2013 and the best short paper of SIGIR 2023.

REFERENCES

[1] Ruibo Liu, Ge Zhang, Xinyu Feng, and Soroush Vosoughi. 2022. Aligning generative language models with human values. In *Findings of the Association for Computational Linguistics: NAACL 2022*. 241–252.