

Panel Title: Large Language Models, Knowledge Graphs, and Vector Databases: Synergy and Opportunities for Data Management

Panel Members:

Haofen Wang (Tongji University)

Haofen Wang is a Distinguished Researcher and Ph.D. supervisor under the "100 People Plan" at Tongji University. He is one of the initiators of OpenKG, the world's largest alliance for Chinese open knowledge graphs. He has participated in and led several national-level AI-related projects, published over 100 high-level papers in the AI field with more than 3,900 citations and an H-index of 29. He developed the world's first interactive virtual idol—"Amber Xuyan." Additionally, the intelligent customer service robots he built have served over 1 billion users. Currently, he holds several social positions including Vice Chairman of the Terminology Committee of the Chinese Computer Federation (CCF), Secretary-General of the Natural Language Processing Society, Director of the Chinese Information Society of China, Executive Committee member of the Large Model Committee, Deputy Secretary-General of the Language and Knowledge Computing Committee, and Deputy Director of the Natural Language Processing Committee of the Shanghai Computer Society.

Wei Hu (Nanjing University)

Wei Hu is a full professor in the School of Computer Science at Nanjing University. His main research areas include knowledge graph, data integration, and intelligent software. He has conducted visiting research at VU University Amsterdam, Stanford University, and University of Toronto. He has published over 50 papers in top-tier conferences and journals, such as SIGMOD, VLDB, ICDE, WWW, SIGIR, ICML, NeurIPS, AAAI, IJCAI, ACL, EMNLP, NAACL, ICSE, TKDE, VLDBJ, TSE, and TNNLS. He has received the Best Paper Awards at JIST, CCKS, and CHIP, and the Best Paper Nomination at ISWC.

Jianguo Wang (Purdue University)

Jianguo Wang is an Assistant Professor of Computer Science at Purdue University. He obtained his Ph.D. from the University of California, San Diego. He has worked or interned at Zilliz, Amazon AWS, Microsoft Research, Oracle, and Samsung on various database systems. His current research interests include database systems for the cloud and large language models, especially disaggregated databases and vector databases. He regularly publishes and serves as a program committee member at premier database conferences such as SIGMOD, VLDB, and ICDE. He also served as a panel moderator for the VLDB'24 panel on vector databases. He is a recipient of the NSF CAREER Award. More information can be found at: <https://cs.purdue.edu/homes/csjpgwang/>.

Shreya Shankar (UC Berkeley)

Shreya Shankar is a PhD student in computer science at UC Berkeley, advised by Dr. Aditya Parameswaran. Her research addresses data challenges in production ML pipelines through a human-centered lens, focusing on data quality, observability, and more recently, leveraging large language models for data preprocessing. Shreya's work has appeared in top data management and HCI venues, including SIGMOD, VLDB, CIDR, CSCW, and UIST. She is a recipient of the NDSEG Fellowship and co-organizes the DEEM workshop at SIGMOD, which focuses on data management in end-to-end machine learning. Prior to her PhD, Shreya worked as an ML engineer and completed her undergraduate degree in computer science at Stanford University. In her free time, she enjoys roasting coffee and is actively trying to reduce her Twitter usage. More information can be found at: <https://www.sh-reya.com/> and https://x.com/sh_reya.