M. Tamer Özsu, University of Waterloo

Title: Exploring the Duality Between Large Language Models and Database Systems

Abstract: This talk examines the reciprocal influence between large language models (LLMs) and database systems. In one direction, LLMs act as intelligent front-ends, translating natural language into executable queries and explanations across diverse data structures, thereby democratizing access to advanced data management. In the other direction, database architectures provide the scalability, low latency, and reliability required to support modern LLM workloads, particularly in retrieval-augmented generation scenarios that rely on efficient vector indexing. Framing these developments as two sides of the same coin, the talk proposes a cohesive research agenda that integrates language intelligence with foundational data management principles.

Bio: M. Tamer Özsu is a University Professor in the David R. Cheriton School of Computer Science at the University of Waterloo and Co-Director of the Faculty of Mathematics graduate program on Data Science and AI. He served as the Director of the Cheriton School from January 2007 to June 2010 and as the Associate Dean of Research for the Faculty of Mathematics from January 2014 to June 2016. His PhD degree is from the Ohio State University (1983). He is a Fellow of the Royal Society of Canada, American Association for the Advancement of Science (AAAS), Science Academy, Türkiye, Asia-Pacific Artificial Intelligence Association (AAIA), and the Balsillie School of International Affairs (BSIA); he is also Life Fellow of Association for Computing Machinery (ACM) and The Institute of Electrical and Electronics Engineers (IEEE). He is the recipient of the University of Waterloo Excellence in Graduate Supervision Award (2025), ACM Presidential Award (2024), IEEE Technical Committee on Data Engineering (TCDE) Education Award (2024), IEEE Innovation in Societal Infrastructure Award (2022), CS-Can/Info-Can Lifetime Achievement Award (2018), ACM SIGMOD Test-of-Time Award (2015), the ACM SIGMOD Contributions Award (2006), The Ohio State University College of Engineering Distinguished Alumnus Award (2008), and multiple Outstanding Performance Awards at the University of Waterloo His publications have received four best paper awards and one honorable mention.

Philip S. Yu, University of Illinois at Chicago

Title:
Towards Graph Foundation Models with Riemannian Geometry

Abstract:
Graph is a ubiquitous non-Euclidean structure, describing the intercorrelated objects in the complex system, ranging from social networks, transportation systems, financial transaction networks to biochemical and molecule structures. Nowadays, graph neural networks are becoming the de facto solution for learning on graphs, generating node or graph embeddings in representation space, such as the traditional Euclidean space. However, a natural and fundamental question that has been rarely explored is: which representation space is more suitable for complex graphs? In fact, the "flat" Euclidean space is suitable for grid structures but is not geometrically aligned with generic graphs with complex structures. Thus, it is interesting to explore deep graph learning in different geometric spaces. This talk will delve into the fascinating properties of mixing various geometric spaces (e.g., hyperbolic and hyperspherical spaces), grounded in the elegant framework of Riemannian geometry, and will discuss recent advancements in tasks such as classification, clustering, contrastive learning, graph structure learning, and continual graph learning. With graph foundation models drawing increasing attention, the talk will also cover preliminary work on building a foundation model for graph structures by exploring Riemannian geometric space. These endeavors pave the way for the next generation of deep graph learning.

Bio:
Dr. Philip S. Yu is a Distinguished Professor and the Wexler Chair in Information Technology at the Department of Computer Science, University of Illinois at Chicago. He is a Fellow of the ACM and IEEE. Dr. Yu is the recipient of ACM SIGKDD 2016 Innovation Award for his influential research and scientific contributions on mining, fusion and anonymization of big data, the IEEE Computer Society's 2013 Technical Achievement Award for "pioneering and fundamentally innovative contributions to the scalable indexing, querying, searching, mining and anonymization of big data" and the Research Contributions Award from ICDM in 2003 for his pioneering contributions to the field of data mining. Dr. Yu has published more than 1,600 referred conference and journal papers cited more than 195,000 times with an H-index of 197. He has applied for more than 300 patents. Dr. Yu was the Editor-in-Chiefs of ACM TKDD (2011-2017) and IEEE TKDE (2001-2004).

Angela Bonifati, Lyon 1 University

Title:
Reasoning over Property Graphs: Leveraging Large Language Models for Automated Data Consistency

Abstract:
Graph data structures are foundational for modeling complex relationships across a wide range of domains, including the life sciences, social media, healthcare, finance, security, and planning. Property graphs, in particular, have emerged as a dominant paradigm for encoding semantic relationships due to their expressiveness and flexibility. However, the increasing adoption of graph databases has intensified the need for robust mechanisms to ensure data quality and consistency.
Traditional consistency maintenance techniques—such as domain-specific rules and mined constraints like functional and entity dependencies—face significant limitations in scalability, adaptability to evolving data, and interpretability by non-experts. In this talk, I will discuss the emerging role of Large Language Models (LLMs) as a tool for automatically discovering and refining consistency constraints in property graphs through guided natural language prompts.
By bridging symbolic graph representations with the reasoning capabilities of LLMs, I will also pinpoint promising directions for automating integrity management in graph systems, along with a wide array of graph-based reasoning tasks.

Bio:
Angela Bonifati is a Distinguished Professor of Computer Science at Lyon 1 University and at the CNRS Liris research lab, where she leads the Database Group. She is also an Adjunct Professor at the University of Waterloo in Canada from 2020 and a Senior member of the French University Institute (IUF) from 2023. Her current research interests are on several aspects of data management, including graph databases, knowledge graphs, data integration and their applications to data science and artificial intelligence. She has co-authored more than 200 publications in top venues of the data management field, including five Best Paper Awards, two books and an invited paper in ACM Sigmod Record 2018. She is a recipient of an ERC Advanced Grant 2024 dedicated to leading researchers in Europe. She is the youngest recipient of the prestigious IEEE TCDE Impact Award 2023 and a co-recipient of an ACM Research Highlights Award 2023. She is the General Chair of VLDB 2026 and was the Program Chair of IEEE ICDE 2025, ACM Sigmod 2022 and EDBT 2020. She is currently an Associate Editor for the Proceedings of VLDB Vol. 19 and for IEEE TKDE and ACM TODS. She is the Chair of the ACM Sigmod Executive Committee (2025-2029) and was the President of the EDBT Executive Board and Association (2020-2024). She is a member of the IEEE Technical Committee on Data Engineering (2024-2029) and a member of the PVLDB Board of Trustees (2024- 2029).

Cheng Chen, ByteDance

Talk Title:
Applications and Challenges of GraphRAG and Graph Foundation Models at ByteDance

Abstract:
With the widespread adoption of Large Language Models (LLMs), the integration of LLMs with graph-structured data—referred to as LLM+Graph—has emerged as a promising direction for enhancing machine reasoning, retrieval, and knowledge representation. At ByteDance, we have observed a growing number of internal applications that leverage this synergy. This talk presents our ongoing exploration of two representative paradigms: GraphRAG (Graph Retrieval-Augmented Generation) and Graph Foundation Models. We begin by discussing the potential of GraphRAG to improve retrieval quality and context-aware generation in complex question answering and enterprise knowledge management, along with the technical challenges it introduces—such as dynamic graph construction and multi-hop reasoning accuracy. In parallel, we examine efforts to build Graph Foundation Models that aim to embed multimodal information (e.g., text, images, behavioral signals) into graph node representations, and to empower LLMs with a deeper understanding of graph topology. These models are expected to facilitate downstream applications in recommendation systems, fraud detection, and other graph-centric tasks. Through this talk, we share insights from real-world scenarios, architectural considerations, and preliminary empirical results, with the goal of fostering a deeper understanding of LLM+Graph integration and its implications for future AI systems.

Bio:
Dr. Chen Cheng is a Tech Lead from the ByteGraph team at ByteDance. He received his Ph.D. from the National University of Singapore. His research interests span graph neural networks (GNN), AI infrastructure, graph computing, graph databases, and storage systems. He has published around 40 papers in top-tier conferences and journals, accumulating over 1,400 citations. His work has been recognized with several honors, including the Best Paper Runner-up Award at EuroSys 2024 and the Best Industry Paper Runner-up Award at VLDB 2023. He actively contributes to the academic community and has served as a program committee member for both research track (ICDE 2025) and industrial track (VLDB 2025, SIGMOD 2024, ICDE 2025/2024).

Heng Lin, Ant Group

Title:
Chat2Graph: a Graph Native Agentic System.

Abstract:
Chat2Graph is a Graph Native Agentic System. By leveraging the inherent advantages of graph data structures such as relationship modeling and interpretability, it enhances key capabilities of agents including reasoning, planning, memory, and tools. This enables diverse functionalities in intelligent R&D, O&M, Q&A, and generation within graph databases, helping users, developers, PD, SA, and SRE to efficiently utilize graph databases, lower the barriers to using graphs, accelerate content generation, and achieve chat with graphs. It truly realizes the deep integration of graph computing technology and artificial intelligence technology, known as "Graph + AI".

Bio:
Heng LIN is a Staff Engineer at Ant Group, where he leads the graph database product (branded as TuGraph) and open-source initiatives. He holds a PhD from Tsinghua University and previously co-founded FMA Technology. His work focuses on distributed graph databases and large-scale heterogeneous graph computing, including a graph system processing 70 trillion edges.

Brian (Borun) Shi, Neo4j

Title:
Retrieval and Reasoning with LLMs on Neo4j: Progress and Challenges

Abstract:
Transformer-based LLMs struggle with complex reasoning on graphs and lack the ability to interact with external systems. Retrieval-augmented methods and agentic systems have largely addressed these limitations. In this talk, we will explore a range of methods and tools Neo4j has developed in these areas in recent years, including retrieval-augmented generation on graphs (GraphRAG) and Text-to-Cypher. We will also introduce and demo the latest agentic system on Neo4j that brings us closer to general-purpose retrieval and reasoning on graphs. The talk will cover both research and engineering perspectives on Neo4j′s current progress, and highlight remaining bottlenecks and research challenges.

Bio:
Brian (Borun) Shi is a senior software engineer at Neo4j, working on both research and engineering in graph machine learning and graph algorithms. More recently, he leads multiple initiatives to improve GraphRAG using techniques from graph learning and algorithms, as well as developing a general-purpose graph reasoning agent.