Supporting Our AI Overlords (SAO): AI Agents and Data Systems

Elaine Ang
Columbia University
USA

Shu Liu
UC Berkeley
USA

Aditya Parameswaran
UC Berkeley
USA

John Dickerson
Mozilla AI
USA

Jonathan Frankle
Databricks
USA

Jacopo Tagliabue*
Bauplan Labs, NYU
USA

Abstract

AI agents are rapidly becoming first-class users of data platforms, but today's OLAP systems were designed for a small number of careful human operators. This workshop brings together researchers and practitioners to study what changes when OLAP systems meet AI agents: stronger safety requirements, self-serve analytics, and new opportunities for agents to optimize the systems they run on. We solicit contributions spanning infrastructure, user workflows, and agent-assisted system design, with an emphasis on evaluation and deployment lessons. Our goal is to map the agent-first design space for data management and catalyze collaborations across the stack.

CCS Concepts: • **Information systems → Data management systems**; • **Computing methodologies → Intelligent agents**.

Keywords: data agents, OLAP, lakehouse, safety, text-to-SQL

ACM Reference Format:
Elaine Ang, Shu Liu, Aditya Parameswaran, John Dickerson, Jonathan Frankle, and Jacopo Tagliabue. 2018. Supporting Our AI Overlords (SAO): AI Agents and Data Systems. In *Proceedings of Make sure to enter the correct conference title from your rights confirmation email (Conference acronym 'XX)*. ACM, New York, NY, USA, 3 pages. https://doi.org/XXXXXXX.XXXXXXX

1 Topic and Scope

In late 2025, [5] argued that existing OLAP systems are not ready for a world in which a large fraction of workflows are driven by AI agents — from self-serve analytics to data engineering and data science automation. The *Supporting AI Overlords (SAO)* vision highlighted three system-level pressure points for agentic analytics: correctness under concurrency, efficient support for semantic probing, and predictable performance under new access patterns.

Building on these foundations, this workshop organizes research directions by looking both "up" (toward users and use cases) and "down" the data stack (toward systems and optimization):

1. **Productionize Agentic Workloads.** While agents are increasingly capable of tackling tasks that historically required data analysts, data engineers, or data scientists [1], a deployment gap remains between benchmarks and production systems across analytics, engineering, and ML pipelines. We welcome work that captures the nuances of agentic reasoning and engineering techniques that bridge the production gap across the entire data lifecycle [6].
2. **Optimizing Agent Semantics.** Agents are increasingly defining new semantics for processing natively unstructured or ambiguous data [7]. We welcome work that explores the transition from deterministic SQL execution to agent-driven pipelines that may involve heavy data-driven reasoning.
3. **Data systems for agents.** Building an agent-first infrastructure requires rethinking core system guarantees from the ground up to support non-human workloads [10] and low-to-zero trust actors [8]. We seek designs that make agent-driven work safe, reproducible, and cost-efficient.
4. **Agents for data system design.** LLM-based design loops are starting to outperform hand-tuned heuristics [2], including agents optimizing OLAP internals [3, 4]. We want to explore the "self-driving" potential of the stack, where agents autonomously design, tune, and maintain the very infrastructure they inhabit.

2 Significance and Relevance

As industry begins to deploy AI-first data platforms at scale [9], this workshop will provide a focused forum on how agentic workloads change the design and operation of data systems. Because data management is a major driver of cloud compute and storage consumption, even incremental improvements for agentic workloads can translate into outsized efficiency gains at scale. The workshop will bring together designers, system builders, and practitioners deploying agents in production.

*Corresponding author: jacopo.tagliabue@bauplanlabs.com

3 Organizers and Audience

3.1 Organizers

Organizers are a mix of academic and industry practitioners, with a wide seniority distribution, ranging from young researchers to established industry leaders. The organizing team has extensive experience in organizing successful multidisciplinary workshops with hundreds of participants in conjunction with major conferences (SIGIR, KDD, CIKM).

Elaine Ang (Columbia University). Ph.D. student at Columbia and a member of the Data, Agents, and Processes Lab. Her research focuses on architecting system paradigms for agentic workloads and leveraging autonomous agents to optimize data-centric systems. She is co-advised by Eugene Wu and Kostis Kaffes.

Shu Liu (UC Berkeley). Ph.D. student in Computer Science at UC Berkeley, where she is a member of the Sky Computing Lab. Her research lies at the intersection of large language models, data systems, and cloud computing, with particular emphasis on agentic workloads and data analytics infrastructure.

Aditya Parameswaran (UC Berkeley). Associate Professor of Computer Science and co-director of the EPIC Data Lab. His research spans data management and human-in-the-loop analytics, and he has led widely used open-source systems (e.g., Modin, Lux, IPyFlow, DocETL).

John Dickerson (Mozilla.ai). CEO of Mozilla.ai, with experience building and deploying enterprise AI systems; previously, he was co-founder and Chief Scientist at Arthur and a Professor of Computer Science at the University of Maryland. He has extensive experience co-organizing successful, multi-year workshops at NeurIPS, AAMAS, and ICLR.

Jonathan Frankle (Databricks). He is the Chief AI Scientist at Databricks, leading research on efficient and reliable deep learning for enterprise settings; he joined Databricks via the acquisition of MosaicML.

Jacopo Tagliabue (Bauplan, New York University). CTO and co-founder of *Bauplan*, the first agentic lakehouse, he is Adj. Professor of MLSys, speaks regularly at top-tier conferences (e.g. VLDB, NAACL, ICML, KDD) and organized popular workshops at SIGIR, CIKM, KDD. *Jacopo is the main point of contact*: jacopo.tagliabue@bauplanlabs.com.

3.2 Audience and history

This is the first edition of the *SAO* workshop. We believe it is a strong fit for the first ACM conference on Agentic Systems, as agentic systems increasingly both depend on *and* bring changes to data management infrastructure. The workshop targets both (i) researchers developing new abstractions and algorithms for agent-first data systems, and (ii) practitioners deploying agentic analytics who want to compare notes on where today's stacks impose bottlenecks in safety and reliability.

Diversity and inclusion. The organizing team spans international backgrounds and combines perspectives from academia and industry, with diversity across seniority, career stage, and gender. We will leverage this breadth to actively reach out to participants across communities and to encourage submissions from groups that are historically underrepresented. In addition to our sponsorship plan, we will publicize the workshop broadly through multiple professional networks and mentoring channels.

4 Call for Papers

Submitted papers will be reviewed in a single-blind process to avoid excluding real-world systems. Submissions will be evaluated for originality, technical quality, clarity, and relevance to the workshop. Accepted contributions will be presented at the workshop; the organizers will also maintain a workshop website where accepted materials can be linked and shared.

We solicit short research and position papers (up to 4 pages + references) and late-breaking results (extended abstracts, up to 2 pages + references). In addition to the directions outlined in Section 1, the following are natural topics, but we will be open to creative interpretations of the workshop themes:

- Agent-first OLAP architectures for safe, reproducible, and cost-efficient agentic analytics.
- Agentic analytics workflows, including human-in-the-loop patterns and production failure modes.
- Evaluation methodologies, benchmarks, and workload traces for data agents (scientists, engineers, analysts).
- LLM-assisted optimization and tuning for data systems.
- War stories and postmortems from production deployments (e.g., incidents, failure modes, debugging and remediation workflows, performance regressions, and operational lessons learned).
- Agentic workflows for data engineering and data science (ETL, feature pipelines, monitoring, notebook-to-production, model/data iteration loops).
- Operational reliability for agent-driven automation (observability, guardrails, incident response, governance, cost controls).

4.1 Prizes and student support

We will leverage relationships with industrial partners to sponsor best-paper awards and to support student and practitioner participation from underrepresented or resource-constrained backgrounds (e.g., via travel support and registration subsidies).

Time	Activity	Details
15 Minutes	**Welcome**	Introductory remarks
40 Minutes	**1st Keynote**	Academia keynote
40 Minutes	**Contributions**	Lightning talks, spotlights
40 Minutes	**2nd Keynote**	Industry keynote
40 Minutes	**Contributions**	Lightning talks, spotlights
60 Minutes	**Panel**	Moderated by organizers
5 Minutes	**Closing remarks**	Closing
After hours	**Pizza and awards**	Pizza night

Table 1. Tentative schedule (half-day workshop, 4 hours).

5 Tentative Schedule

5.1 Keynote speakers and program committee

To reflect a broader view on data agents, we are assembling a diverse program committee of practitioners from **industry** and **academia**. Invited speakers will be selected from a shortlist of world-renowned experts. We plan to invite one keynote from academia and one from industry. We also plan a panel designed to represent a wide range of perspectives in agentic data management, including model builders, coding-agent platforms, infrastructure providers, and enterprise adopters.

5.2 Workshop schedule

Table 1 summarizes our plan for the workshop. Our four-hour workshop (half-day) will feature invited talks, contributed papers, and a final panel discussion.

5.3 After-hours and awards: pizza night

Following the success of similar initiatives we ran at KDD and NeurIPS, we plan to invite participants to stay for an informal pizza night sponsored by the organizers. We will announce awards at the end of the evening and continue discussions in a relaxed setting.

References

[1] Asim Biswal, Liana Patel, Siddarth Jha, Amog Kamsetty, Shu Liu, Joseph E. Gonzalez, Carlos Guestrin, and Matei Zaharia. 2024. Text2SQL is Not Enough: Unifying AI and Databases with TAG. arXiv:2408.14717 [cs.DB] https://arxiv.org/abs/2408.14717

[2] Audrey Cheng, Shu Liu, Melissa Pan, Zhifei Li, Bowen Wang, Alex Krentsel, Tian Xia, Mert Cemri, Jongseok Park, Shuo Yang, Jeff Chen, Lakshya Agrawal, Aditya Desai, Jiarong Xing, Koushik Sen, Matei Zaharia, and Ion Stoica. 2025. Barbarians at the Gate: How AI is Upending Systems Research. arXiv:2510.06189 [cs.AI]

[3] Mehmet Hamza Erol, Xiangpeng Hao, Federico Bianchi, Ciro Greco, Jacopo Tagliabue, and James Zou. 2026. Making Databases Faster with LLM Evolutionary Sampling. arXiv:2602.10387 [cs.DB]

[4] Jiale Lao and Immanuel Trummer. 2026. GenDB: The Next Generation of Query Processing – Synthesized, Not Engineered. arXiv:2603.02081 [cs.DB] https://arxiv.org/abs/2603.02081

[5] Shu Liu, Soujanya Ponnapalli, Shreya Shankar, Sepanta Zeighami, Alan Zhu, Shubham Agarwal, Ruiqi Chen, Samion Suwito, Shuo Yuan, Ion Stoica, Matei Zaharia, Alvin Cheung, Natacha Crooks, Joseph E. Gonzalez, and Aditya G. Parameswaran. 2025. Supporting Our AI Overlords: Redesigning Data Systems to be Agent-First. arXiv:2509.00997 [cs.DB] https://arxiv.org/abs/2509.00997

[6] Cedric Renggli, Ihab F. Ilyas, and Theodoros Rekatsinas. 2025. Fundamental Challenges in Evaluating Text2SQL Solutions and Detecting Their Limitations. arXiv:2501.18197 [cs.LG]

[7] Shreya Shankar, Tristan Chambers, Tarak Shah, Aditya G Parameswaran, and Eugene Wu. 2024. Docetl: Agentic query rewriting and evaluation for complex document processing. arXiv preprint arXiv:2410.12189 (2024).

[8] Natalie Shapira, Chris Wendler, Avery Yen, Gabriele Sarti, Koyena Pal, Olivia Floody, Adam Belfki, Alex Loftus, Aditya Ratan Jannali, Nikhil Prakash, Jasmine Cui, Giordano Rogers, Jannik Brinkmann, Can Rager, Amir Zur, Michael Ripa, Aruna Sankaranarayanan, David Atkinson, Rohit Gandikota, Jaden Fiotto-Kaufman, EunJeong Hwang, Hadas Orgad, P Sam Sahil, Negev Taglicht, Tomer Shabtay, Atai Ambus, Nitay Alon, Shiri Oron, Ayelet Gordon-Tapiero, Yotam Kaplan, Vered Shwartz, Tamar Rott Shaham, Christoph Riedl, Reuth Mirsky, Maarten Sap, David Manheim, Tomer Ullman, and David Bau. 2026. Agents of Chaos. arXiv:2602.20021 [cs.AI] https://arxiv.org/abs/2602.20021

[9] Weiming Sheng, Jinlang Wang, Manuel Barros, Aldrin Montana, Jacopo Tagliabue, and Luca Bigon. 2026. Building a Correct-by-Design Lakehouse. Data Contracts, Versioning, and Transactional Pipelines for Humans and Agents. arXiv:2602.02335 [cs.DC]

[10] Jacopo Tagliabue, Federico Bianchi, and Ciro Greco. 2025. Trustworthy AI in the Agentic Lakehouse: from Concurrency to Governance. arXiv:2511.16402 [cs.AI] https://arxiv.org/abs/2511.16402

Received 9 March 2026