Essam Mansour

Enterprise agentic AI: governed memory, action contracts, validation, and agent security

Concordia Data Systems Lab (CoDS)
Concordia University
2155 Guy Street (Room 1045)
Montreal, Quebec H3H 2L9, Canada
General Chair, IEEE ICDE 2026

emansour.com
LinkedIn, Google Scholar, GitHub: CoDS-GCS

essam.mansour@concordia.ca
Citizenship: Canadian

Research Vision: Enterprise-Aware Agents

My research builds agents that operate safely on organizational data. The barrier to enterprise agentic deployment is not model capability. It is that agents have no governed way to accumulate what an organization knows, and no way to establish that what they accumulated is trustworthy. Our SIGMOD 2027 study established this empirically across nine production agent frameworks at fixed LLM and task. Orchestration alone changes latency by over 60×. Schema-constrained planning costs up to 32 accuracy points through parsing failures rather than reasoning failures. Communication topology swings coordination success from above 90% to below 30%. Not one framework supports controlled knowledge revision. **Architecture, not only model capability, governs agent performance.** That is a data and systems problem. I attack it at the data layer, across four coupled fronts.

Memory is a data-management workload, not storage. Every framework we measured treats memory as append-only, so an agent cannot revise what it believes. Our vision paper answers that gap with Governed Evolving Memory, where correctness is a property of the state trajectory rather than of individual records. Record-level systems cannot satisfy these conditions regardless of substrate.

Governance must be structural. Agents reach enterprise data only through mediated contracts, so policy is enforced by construction rather than by configuration. Every read, write, and promotion of a claim passes a declared gate.

Validation must be grounded in data. Every system I have shipped replaces LLM self-judgment with assertion-based checks against the actual data. Classical functional dependencies break under non-deterministic producers, so derivation-aware constraints replace them.

Security is a first-class workload. Agent memory is an attack surface. The promotion of agent-generated claims to observed status without a governed gate is where enterprise trust fails silently.

This program is grounded in seven open-source systems and two diagnostic benchmarks, published at SIGMOD, VLDB, ICDE, EMNLP, EACL, and CCS across 2025 to 2027.

Deployed Impact and Adoption

- **Production adoption at Google.** KGLiDS (ICDE 2024) was adopted by Google and Kaggle teams for production data science workflows. The follow-on system RAGvis is released inside Google's own GitHub organization (google/ragvis) and published at EMNLP 2025, main conference.
- **Systems shipped into enterprise constraints.** Six years as a research scientist at QCRI, technical lead on three joint MIT CSAIL projects that delivered running systems. Since 2019, shipping collaborations with Google, IBM Research, National Research Council Canada, National Bank of Canada, RBC Borealis AI, and Qohash, securing over $3M in competitive funding.
- **Measured gains against real workloads.** ReSequel (VLDB 2026) delivers up to 16× over native PostgreSQL, MySQL, and DuckDB and 22× over LLM baselines across eight benchmarks. Chatty-KG (SIGMOD 2026) doubles accuracy on conversational KG question answering, from 39 to 80 F1, using the same LLM under a better architecture.
- **Benchmarks used by the community.** MAFBench (SIGMOD 2027) for multi-agent framework diagnosis across five architectural dimensions, yielding six evidence-based design principles, and Chatty-Gen (SIGMOD 2025), which reduced dialogue benchmark generation cost from thousands of dollars to under one dollar.
- **Talent placed into industry.** Graduates and team members now at IBM Research, Verily (Alphabet), Mila–McGill, University of Waterloo, National Bank of Canada, Desjardins, and Snowflake.
- **Agentic tooling in daily research practice.** I use agentic coding tools across the lab for benchmark construction, experiment orchestration, and artifact reproduction, and I set the team's practices for their safe use.

Professional Experience

- **Concordia University**, Montreal, Canada. **Head, Concordia Data Systems Lab (CoDS)**, 2019–present.
  Tenured Associate Professor (June 2024–present); Assistant Professor (July 2019–May 2024).
  Build and lead a group of five PhD students, one MSc student, and one postdoctoral researcher. Direct the GEMS research program and the CoDS-GCS open-source organization. Set the research agenda, secure the funding, and run the delivery of shipped systems with industrial partners.
- **Qatar Computing Research Institute**, Doha, Qatar. **Research Scientist**, June 2013–July 2019.
  Distributed Systems and Data Analytics Group. Technical lead on three joint MIT CSAIL–QCRI projects with ***Michael Stonebraker***, ***Samuel Madden***, and ***Tim Berners-Lee***. Delivered E-Store, Data Civilizer, and Solid as running systems, leading mixed teams of researchers and engineers.
- **King Abdullah University of Science and Technology (KAUST)**, Thuwal, Saudi Arabia. **Postdoctoral Fellow**, October 2009–June 2013.
  Led large-scale string analytics research; co-supervised two graduate students.
- **International University in Germany**, Bruchsal, Germany. **Postdoctoral Fellow**, August 2008–October 2009.
  Mobile data management and energy-efficient database systems.

Industry Research Collaborations

- **IBM Research, USA (2020–present).** With Dr. Ibrahim Abdelaziz (technical lead on IBM Granite and agentic AI) and Dr. Kavitha Srinivas. Multi-agent AI systems for conversational question answering and enterprise data science assistance. Prior work yielded KGpip (PVLDB 2022).
- **Google, USA (Jan 2024–Aug 2025).** With Dr. Jichuan Chang (Director and Principal Engineer). RAG-based systems harnessing KGs and Gemini to automate EDA pipelines. Outcome: RAGvis (EMNLP 2025), released in Google's GitHub organization.
- **National Research Council Canada (Sept 2025–Aug 2028).** With Dr. Shahrear Iqbal. $450K project (my share: $150K) on secure knowledge sharing and agentic threat detection for the digitized construction industry.
- **National Bank of Canada (2023–2024).** LLM-based risk assessment system over project data logs.
- **RBC Borealis AI, Canada (2022–2023).** Linked data science platform for feature discovery and automated data preparation.
- **MIT CSAIL, USA (2013–2018).** Technical lead on three joint projects with QCRI: E-Store (VLDB 2015), Data Civilizer (ICDE 2018), and Solid (WWW 2016).

Research Program 2026–2030: Enterprise-Aware Agents (GEMS)

GEMS (Governed Enterprise Metadata for Agentic Systems) is my active research platform. It moves agents from individual intelligence to organizational intelligence. The program is funded through 2030 by NSERC Discovery and NSERC CREATE, released open-source, and staffed by five PhD students and one MSc student recruited specifically against these thrusts. Three PhD students and one MSc student joined in September 2026.

- **Governed memory and tacit knowledge.** Long-term agent memory as a governed state machine, with typed dependencies, semantic revision, declarative forgetting, and multi-tenant privacy. *L. Guelmani (PhD, 2026–2030); Dr. A. Orogat (postdoctoral fellow).*
- **Action contracts and governed mediation.** Structural enforcement of what an agent may read, write, and promote. This addresses belief laundering, the promotion of agent-generated claims to observed status without a governed gate. *F. Aminnezhad (MSc, 2026–2028).*
- **Validation of agent-generated code.** Assertion-based and derivation-aware validation that replaces probabilistic self-judgment with checks against source data. *Q. Wang (PhD, 2026–2030).*
- **Agent security and threat detection.** Multi-agent attack surfaces, memory poisoning, and prompt injection under enterprise deployment. Builds on OCR-APT (CCS 2025) and MEGR-APT (TIFS 2024). *Y. Gharib (PhD, 2026–2030); A. Aly (PhD, to 2027).*

- **Graph foundation models for agentic reasoning.** Motif-centric inference over large enterprise knowledge graphs, giving agents structured recall at scale. *W. Afandi (PhD, to 2027).*

Open-Source Systems and Benchmarks

I direct the CoDS-GCS GitHub organization, which open-sources the portfolio across more than 30 repositories. Each repository ships with reproducible benchmarks, documentation, and an associated peer-reviewed publication.

Agentic AI systems and benchmarks (2025–2027)

- **MAFBench** (SIGMOD 2027) — Unified evaluation suite for multi-agent LLM frameworks, with an architectural taxonomy over five dimensions and a controlled study across nine frameworks including production coding agents. Yields six evidence-based design principles. Foundational reference for the architectural thesis above.
- **MemState** (vision, 2026) — Prototype realizing Governed Evolving Memory on a property-graph backend, with topic-based storage, typed dependencies, and declarative policies. Exposes what a native long-term memory engine must provide.
- **Chatty-KG** (SIGMOD 2026) — Multi-agent AI system for on-demand conversational question answering over knowledge graphs, with multi-turn dialogue across large-scale KGs.
- **ReSequel** (VLDB 2026) — LLM-assisted SQL rewriting over PostgreSQL, MySQL, and DuckDB. Templatizes workloads, infers rules from catalog and statistics, generates variants, verifies on downsampled data, picks Top-1.
- **OCR-APT** (CCS 2025) — Reconstruction of advanced persistent threat stories from audit logs using subgraph anomaly detection and LLMs, with hard-gating assertion validation.
- **RAGvis** (EMNLP 2025, main conference) — Knowledge-grounded RAG for exploratory data analysis and visualization. Reliable and cost-effective alternative to large-model EDA.
- **CatDB** (VLDB 2025) — Data-catalog-guided, LLM-based generator of data-centric ML pipelines with iterative self-correction validated on the actual data.
- **Chatty-Gen** (SIGMOD 2025) — Cost-effective dialogue benchmark generation from knowledge graphs using retrieval-augmented LLMs.

Foundational systems

- **KGLiDS** (ICDE 2024) and **KGpip** (VLDB 2022) — Semantic abstraction, linking, and automation of data science. *Adopted by Google and Kaggle teams.* Underlies CatDB and RAGvis.
- **KGNet** (ICDE 2024, ICDE 2026) — GNN-enabled knowledge graph engine with task-oriented training and LLM-guided query-aware inference.
- **KGQAn** (SIGMOD 2023) — Universal KG question answering, and GNN-based scalable AutoML with IBM Research.
- **External.** Contributor to aurum-datadiscovery (MIT DBG), and to E-Store (VLDB 2015) and Solid (WWW 2016, with Tim Berners-Lee).

Selected Publications

More than 50 papers in top-tier venues. Full list at DBLP and Google Scholar. Students supervised by or working closely with me are marked with an asterisk (*).

Agentic AI, memory, and safety

[A1] Abdelghny Orogat*, Ana Rostam*, and Essam Mansour. "Architectural Design, Not Only Model Intelligence, Governs Multi-Agent LLM Performance," *Proceedings of the ACM on Management of Data*, SIGMOD 2027. `https://arxiv.org/abs/2602.03128`

[A2] Abdelghny Orogat* and Essam Mansour. "Is Agent Memory a Database? Rethinking Data Foundations for Long-Term AI Agent Memory," 2026. https://arxiv.org/abs/2605.26252

[A3] Reham Omar*, Abdelghny Orogat*, Ibrahim Abdelaziz, Omij Mangukiya*, Panos Kalnis, and Essam Mansour. "Chatty-KG: A Multi-Agent AI System for On-Demand Conversational Question Answering over Knowledge Graphs," *Proceedings of the ACM on Management of Data*, SIGMOD 2026. https://arxiv.org/abs/2511.20940

[A4] Ahmed Aly*, Essam Mansour, and Amr Youssef. "OCR-APT: Reconstructing APT Stories from Audit Logs using Subgraph Anomaly Detection and LLMs," *ACM Conference on Computer and Communications Security (CCS)*, 2025. https://www.arxiv.org/abs/2510.15188

[A5] Ahmed Aly*, Shahrear Iqbal, Amr Youssef, and Essam Mansour. "MEGR-APT: A Memory-Efficient APT Hunting System Based on Attack Representation Learning," *IEEE Transactions on Information Forensics and Security (TIFS)*, April 2024. https://ieeexplore.ieee.org/document/10517754

LLM and knowledge graph systems

[L6] Saeed Fathollahzadeh*, Essam Mansour, and Matthias Boehm. "ReSequel: Robust LLM-assisted Query Rewriting and Optimization using Templatization and Sampling," *PVLDB*, VLDB 2026. https://arxiv.org/abs/2606.20853

[L7] Hussein Abdallah*, Ibrahim Abdelaziz, Panos Kalnis, and Essam Mansour. "Leveraging LLM-GNN Integration for Open-World Question Answering over Knowledge Graphs," *19th Conference of the European Chapter of the ACL (EACL)* 2026: 575–592. https://arxiv.org/abs/2604.13979

[L8] Mossad Helali*, Yutai Luo, Tae Jun Ham, Jim Plotts, Ashwin Chaugule, Jichuan Chang, Parthasarathy Ranganathan, and Essam Mansour. "Beyond Large Models: Reliable and Cost-Effective Exploratory Data Analysis via Knowledge-Grounded RAG," *EMNLP 2025, main conference*. https://aclanthology.org/2025.emnlp-main.836.pdf

[L9] Saeed Fathollahzadeh*, Essam Mansour, and Matthias Boehm. "CatDB: Data-catalog-guided, LLM-based Generation of Data-centric ML Pipelines," *PVLDB* 18(4), VLDB 2025. https://www.vldb.org/pvldb/vol18/p2639-fathollahzadeh.pdf

[L10] Reham Omar*, Omij Mangukiya*, and Essam Mansour. "Dialogue Benchmark Generation from Knowledge Graphs with Cost-Effective Retrieval-Augmented LLMs," *Proceedings of the ACM on Management of Data*, SIGMOD 2025. https://dl.acm.org/doi/10.1145/3709681

[L11] Waleed Afandi*, Hussein Abdallah*, Ashraf Aboulnaga, and Essam Mansour. "KG-WISE: An LLM-Guided Query-Aware Inference System for GNN Models on Large Knowledge Graphs," *IEEE ICDE 2026*. https://arxiv.org/abs/2603.04545

[L12] Reham Omar*, Ishika Dhall*, Panos Kalnis, and Essam Mansour. "A Universal Question-Answering Platform for Knowledge Graphs," *Proceedings of the ACM on Management of Data* 1(1): 1–25, SIGMOD 2023. https://doi.org/10.1145/3588911

Data systems foundations

[D13] Mossad Helali*, Niki Monjazeb*, Shubham Vashisth*, Philippe Carrier*, Ahmed Helal*, Antonio Cavalcante, Khaled Ammar, Katja Hose, and Essam Mansour. "KGLiDS: A Platform for Semantic Abstraction, Linking, and Automation of Data Science," *IEEE ICDE 2024*. https://ieeexplore.ieee.org/document/10597753

[D14] Hussein Abdallah*, Waleed Afandi*, Panos Kalnis, and Essam Mansour. "Task-Oriented GNNs Training on Large Knowledge Graphs for Accurate and Efficient Modeling," *IEEE ICDE 2024*. https://ieeexplore.ieee.org/document/10597998

[D15] Mossad Helali*, Essam Mansour, Ibrahim Abdelaziz, Julian Dolby, and Kavitha Srinivas. "A Scalable AutoML Approach Based on Graph Neural Networks," *PVLDB* 15(1): 2428–2436, VLDB 2022.

[D16] Essam Mansour, Kavitha Srinivas, and Katja Hose. "Federated Data Science to Break Down Silos," *ACM SIGMOD Record* 50(4): 16–22, January 2022.

[D17] Ibrahim Abdelaziz*, Essam Mansour, Mourad Ouzzani, Ashraf Aboulnaga, and Panos Kalnis. "Lusail: A System for Querying Linked Data at Scale," *PVLDB* 11(4): 485–498, VLDB 2018.

[D18] Rebecca Taft, Essam Mansour, Marco Serafini, Jennie Duggan, Aaron J. Elmore, Ashraf Aboulnaga, Andrew Pavlo, and Michael Stonebraker. "E-Store: Fine-Grained Elastic Partitioning for Distributed Transaction Processing," *PVLDB* 8(3): 245–256, VLDB 2015.

Funded Research

Approximately $3M in competitive research funding from federal and industry sources. Amounts in Canadian dollars; share noted where multi-PI.

- $1,650,000 *(share: $210K)*. NSERC CREATE, "Sustainable Data Systems for Data Science." Theme lead and co-investigator, 2024–2030.
- $210,000. NSERC Discovery Grant, "Systems and Algorithms for Integrating LLMs and GNNs into Data Science and Knowledge Graph Platforms." PI, 2025–2030.
- $450,000 *(share: $150K)*. NRC Canada, "Cybersecurity Framework for the Digitized Construction Industry." Co-investigator, 2025–2028.
- $240,000. Concordia University Faculty Research Support Funds. PI, 2020–2027.
- $132,500. NSERC Discovery Grant, "Systems and Algorithms for Easy-to-use Federated Data Science." PI, 2020–2025.
- $150,000 total. MITACS partnerships with Qohash, National Bank of Canada, and RBC Borealis AI. PI, 2022–2024.

Scientific Leadership and Service

- **General Chair, IEEE ICDE 2026** (Montreal, May 2026), the flagship data engineering conference, with Fei Chiang.
- **Editor, ACM SIGMOD Blog.**
- **Invited participant and technical lead, Dagstuhl Seminar 26161** "Managing Vector Data for RAG" (April 2026), an invitation-only venue for senior computer scientists. Led the *Agentic Systems* cross-topic session and delivered a Gong Show talk.
- **Invited position panelist, ICDE 2026 panel "What can agents do for databases?"** with Fatma Özcan (moderator), Wenjie Zhang, Aditya Parameswaran, and Bolin Ding.
- **General Co-Chair**, International Workshop on Data Systems Meet Data Science (DSDS), 2022–2024.
- **Program committees.** ACM SIGMOD (2016, 2021, 2023 Associate Editor, Demo Committee 2025); VLDB / PVLDB (2017–2026); IEEE ICDE (2016, 2023); TheWebConf (2023); KDD (2021).
- **Journal reviewing.** ACM TODS, VLDB Journal, IEEE TKDE, IEEE TPDS. Proposal reviewer for NSERC and MITACS.
- **Recognition.** Distinguished Reviewer Awards: SIGMOD 2025, PVLDB 2018.

Selected Invited Talks

- "Architecture Over Model: Agentic Systems for Data-Intensive Domains." Agentic Systems cross-topic session lead, *Dagstuhl Seminar 26161*, Germany, April 2026.
- "What Agents Can Do for Databases." Position panelist, *ICDE 2026 panel*, Montreal, May 2026.
- "Advancing Domain-Specific Intelligence with LLMs and KGs: A Five-Year Journey." *DBML Workshop at ICDE 2025*, Hong Kong, May 2025.
- "ChatGPT versus Traditional Question Answering for KGs." *Apple Knowledge Platform team*, March 2023. Related seminars at Simon Fraser University (2023) and the University of Waterloo Data Systems Group (2022).

Talent Development

Postdoctoral researcher. Dr. Abdelghny Orogat (2025–2026, Principal Supervisor). Project: The Role of Agentic AI in Enterprise Systems.
Current PhD and MSc students. Ahmed Aly (to 2027, agent security), Waleed Afandi (to 2027, graph foundation models for agents), Yomna Gharib (2026–2030, agentic security), Loubna Guelmani (2026–2030, agent harness and long-term memory), Qizhi Wang (2026–2030, validation of agent-generated code), Farbod Aminnezhad (MSc, 2026–2028, action contracts).
Graduated PhD students and placements. Hussein Abdallah (2025, thesis graded Outstanding) — Postdoctoral Fellow, Mila–McGill. Mossad Helali (2025, thesis graded Excellent) — AI Applied Scientist, Verily (Alphabet). Saeed Fathollahzadeh (2026, co-supervised with Matthias Boehm) — Postdoctoral Scholar, Data Intelligence Lab, University of Waterloo. Reham Omar (2026, thesis graded Outstanding).
Master's graduates and placements. Omijkumar Mangukiya — IBM Canada. Niki Monjazeb — Desjardins. Shubham Vashisth — PhD student, McGill. Ishika Dhall — National Bank of Canada. Ahmed Helal — Snowflake.

Education and Languages

- **PhD in Computer Science**, Technological University Dublin, Ireland, 2004–2008.
- **MS in Information Systems**, Cairo University, Egypt, 2000–2003.
- **BS in Information Systems**, Cairo University, Egypt, 1996–2000.
- **Languages.** English and Arabic (fluent). French (working proficiency, actively improving in a bilingual Montreal environment).

Last updated: September 2026