Yuan Tian

tian211@purdue.edu | yuan-tian.com | (+1) 765-426-6107

Research Statement

My research focuses on agentic database systems, where I build coding and database agents for data-intensive tasks. Such tasks often involve overwhelming contexts that are challenging for agents to consume and process. I design and scale the interaction between AI agents and data environments through different approaches, such as post-training, scaffolding, attention steering, reasoning with verification, etc. Complementing these, I also work on data-centric approaches for domain adaptation, including semantic enrichment and programmatic data synthesis.

Education

Purdue University *August 2021 – Present*
Ph.D. in Computer Science *M.S. awarded, Ph.D. expected 2026*

University at Albany, State University of New York *August 2019 – May 2020*
Bachelor of Science, Computer Science *GPA: 3.97/4.00*

Chongqing University of Posts and Telecommunications *August 2016 – July 2019*
Bachelor of Engineering, Computer Science & Software Engineering *Ranked* 1^{st}

Work Experience

Research Assistant *August 2021 – Present*
Purdue University, HCSS Lab | West Lafayette, IN, USA

- Independent research in code and data language models
- Built an ontology-driven data extraction pipeline and knowledge graph for AI security (NSF Proto-OKN), selected for the NSF National AI Research Resource (NAIRR) Pilot

Applied Scientist [Intern] *May 2025 – August 2025*
Adobe, GenAI @ AEP | San Jose, CA, USA

- Built an end-to-end semantic enrichment pipeline for Adobe XDM system, which scores semantic quality, detects under-specified fields, and enriches field descriptions w/ and w/o human feedback
- Improved downstream task accuracy by 1.5x for entity linking and 1.2x for text-to-SQL

Applied Scientist [Intern] *May 2024 – August 2024*
Adobe, GenAI @ AEP | San Jose, CA, USA

- Built an end-to-end text-to-SQL domain adaptation system that populates sandbox databases with verified queries for arbitrary schemas
- Improved data accuracy by 1.5× and data diversity by 1.9×, while enabling up to 44× faster annotation w/ human validation or fully automated w/o humans

Teaching Assistant *January 2024 – May 2024*
Purdue University, Computer Science | West Lafayette, IN, USA

- Delivered lectures on discrete mathematics, algorithms, and data structures (CS 182)

Peer-Reviewed Publications

PV-SQL: Synergizing Database Probing and Rule-based Verification for Text-to-SQL Agents
Yuan Tian, Tianyi Zhang
[ACL 2026] Findings | Text-to-SQL, Agent, Verification-driven, Database Probing

- Developed a novel text-to-SQL agent using a probe-and-verify approach
- Proposed a rule-based SQL verifier that provides reliable feedback signals for database queries
- Comprehensively analyzed text-to-SQL errors made by LLMs and existing SOTA text-to-SQL methods

ALL-FEM: Agentic Large Language Models Fine-tuned for Finite Element Methods
Rushikesh Deotale, Adithya Srinivasan, Yuan Tian, Tianyi Zhang, Pavlos Vlachos, Hector Gomez

[Computer Methods in Applied Mechanics and Engineering] – Journal | Fine-tuning, Code Generation, Finite Element Methods

- Developed an LLM-based multi-agent system to generate and iteratively refine finite element code
- Fine-tuned agents within a multi-agent framework for domain adaptation

Selective Prompt Anchoring for Code Generation

Yuan Tian, Tianyi Zhang

[ICML 2025] | Attention Steering, Taylor Expansion, Logit Arithmetic, Attention Analysis, Code Generation

- Identified the attention dilution phenomenon as a root cause of code generation errors in LLMs
- Proposed and mathematically proved a general attention-steering method for LLMs
- Integrated with the Hugging Face API and supports all Hugging Face LLMs
- Proposed an attention-based code generation pipeline, achieving new SOTA performance
- Received significant attention in the open-source community

Text-to-SQL Domain Adaptation via Human–LLM Collaborative Data Annotation

Yuan Tian, Daniel Lee, Fei Wu, Tung Mai, Kun Qian, Siddhartha Sahai, Tianyi Zhang, Yunyao Li

[IUI 2025] | Human-AI Collaboration, Domain Adaptation, Interactive Systems, Text-to-SQL, Data Augmentation

- Proposed a comprehensive framework (UI + backend) for schema editing, interactive text-to-SQL annotation, automated text-to-SQL data augmentation, and text-to-SQL dataset analysis
- Conducted a rigorous user study with 12 participants to evaluate usability and annotation efficiency

EvoSchema: Towards Text-to-SQL Robustness Against Schema Evolution

Tianshu Zhang, Kun Qian, Siddhartha Sahai, Yuan Tian, Shaddy Garg, Huan Sun, Yunyao Li

[VLDB 2026] | Schema Evolution, Text-to-SQL, Benchmark, Perturbation Taxonomy

- Built a benchmark to test text-to-SQL robustness under schema evolution
- Defined 10 types of schema perturbations based on a hybrid method (LLM + heuristics)
- Showed that fine-tuning LLMs on perturbed schemas improves performance

AgentPbD: Interactive Agentic Workflow Generation from User Demonstration on Web Browsers

Jiawen Li, Zheng Ning, Yuan Tian, Toby Jia-Jun Li

[VL/HCC 2026 Poster] | LLM Workflow, Programming by Demonstration

- Developed a multi-agent system that converts user browser demonstrations into LLM workflows
- Developed an interactive interface for visualizing, refining, and reusing workflows through visual programming

Supporting Construction Worker Well-Being with Multi-Agent Conversational AI

Fan Yang, Yuan Tian, Jiansong Zhang

[CRC 2025] | Multi-Agent Systems, Conversational AI, Well-being

- Developed a conversational multi-agent system for construction workers' mental health support
- Agent customization with internal prompt generation and external RAG-based document upload
- Conducted a user study with 12 participants to demonstrate improved user engagement and support effectiveness through group chat with AI agents

SQLucid: Grounding Natural Language Database Queries with Interactive Explanations

Yuan Tian, Jonathan K. Kummerfeld, Toby Jia-Jun Li, Tianyi Zhang

[UIST 2024] | Interactive Systems, Text-to-SQL, Database Interfaces, Grounding Theory

- Built a novel interactive SQL generation tool based on editable step-by-step explanations, visual grounding, and intermediate query executions
- Conducted two comprehensive user studies (30 participants) validating system effectiveness

Insights into NL Database Query Errors: From Attention Misalignment to User Strategies

Zheng Ning*, Yuan Tian*, Zheng Zhang, Tianyi Zhang, Toby Jia-Jun Li

* Equal Contribution

[TiiS 2024] – Journal | Error Analysis, Attention Mechanisms, User Behavior

- Extended our previous text-to-SQL error analysis to include LLM analysis and attention studies
- Demonstrated that models make errors when their attention does not align with human attention

Interactive SQL Generation via Editable Step-by-Step Explanation

Yuan Tian, Zheng Zhang, Zheng Ning, Toby Jia-Jun Li, Jonathan K. Kummerfeld, Tianyi Zhang

[EMNLP 2023] | Grammar/Rule-based method, SQL Parsing, Explanations, Text-to-SQL, Clause Generation

- Proposed "editable step-by-step explanation", a novel mechanism for SQL generation and repair

- Built a robust grammar-based SQL parser, a rule-based NL explanation generator, and a neural-symbolic clause-level SQL editing model for error correction

An Empirical Study of Model Errors and User Repair Strategies in NL-to-SQL
Zheng Ning*, Zheng Zhang*, Tianyi Sun, Yuan Tian, Tianyi Zhang, and Toby Jia-Jun Li
[IUI 2023] | Error Taxonomy, User Studies, Interactive Repair

- Developed a taxonomy of SQL errors produced by SOTA text-to-SQL models
- Conducted a within-subjects study with 26 participants to compare different interactive systems

Preprints

Attention-Aligned Reasoning for Large Language Models
[Preprint] | Attention Steering, Alignment, Reasoning

- Developed a novel key-value reasoning pattern based on step-by-step attention steering
- Outperformed existing prompting-based approaches and RL-trained SLMs

Alloy: Generating Reusable Agent Workflows from User Demonstration
[Preprint] | Computer Use, Workflow Generalization, Programming by Demonstration, MCP

- Built an agentic system that converts computer use demonstrations into editable and reusable LLM workflows
- Evaluated workflow generation and generalization through a user study and quantitative evaluation

Large Language Models in Computational Mechanics: Toward an Autonomous Simulation Paradigm
[Preprint] | Finite Element Method, Code Generation, Multi-agent System, Fine-tuning

- Fine-tuned multiple base LLMs on FEniCS code (finite element method).
- Helped develop a multi-agent system for code generation for finite element simulations

Conversational Agents for Construction Safety Education: Supporting Affective and Attitudinal Learning through Real-World Accidents
[Preprint] | Conversational AI, Multi-agent System, AI Education

- Built a conversational multi-agent system for safety education using real accident data
- Two user studies showed that conversation can reshape users' empathy and awareness but cannot change their fundamental mindset

MUSE: A User-Side Meta-Agent for Controlling and Evaluating Multi-Agent Data-Science Systems
[Preprint] | Data Science, Multi-agent System, Interactive Design

- Built a novel user-side agent providing understandable explanations hierarchically
- Conducted a formative study to analyze the limitations of existing agentic systems and user needs

AI Chatroom for Mental Health: Generating Personalized Well-being Reports with LLMs and Taxonomy-Driven Checklists
[Preprint] | Multi-agent System, LLM-scoring, PERMA, Chatroom

- Built an agentic well-being analysis tool based on the PERMA mental model
- An LLM-based simulated psychologist initiates and guides conversations with users using a clarification checklist based on the PERMA model
- Confidence-aware mental health report generation based on conversation history, featuring a pentagon PERMA analysis chart, detailed explanations, and actionable suggestions

TACO: Task-Aware Column Description Generation Using LLMs
[Preprint] | Semantic Enrichment, Tabular Data, Downstream-Aware Generation

- Helped develop a system that expands abbreviations, generates, and revises tabular column descriptions
- Proposed a novel downstream-aware column description generation method

Ask More, Score Higher: Towards Robust Semantic Enrichment through Interactive Query Population and Richness Scoring
[Preprint] | UI Design, Human-AI Collaboration, Semantic Enrichment, Query Population

- Designed the UI and multi-agent-based backend for semantic enrichment
- Developed a pipeline to enhance semantics by verifying populated queries and real-time richness evaluation

ISEE: Interactive Semantic Enrichment for Database Fields
[Preprint] | Interactive Semantic Enrichment, Scoring System, Clarification Taxonomy

- Developed a robust richness scoring system with 5 distinct metrics, each based on a unique algorithm

- Proposed a clarification taxonomy for semantic enrichment
- Conducted a user study with industrial engineers to evaluate the usability and effectiveness

Encountered-Type Haptic Display via Tracking-Calibrated Robot
[Preprint] | VR, Robotics, Unity, Oculus, Haptics, Tactile Feedback, Calibration

- Developed a novel method to simulate haptic feedback in VR using the UR16e robotic arm
- Developed a robust tracking calibration algorithm for virtual-physical synchronization

Industry Talks

Datadog | AI Research *Apr 30, 2026*
Interaction and Adaptation between Agents and Data-intensive Systems

Meta | Reality Lab *Aug 21, 2025*
Attention Steering with KV Cache Optimization

Adobe | Experience Platform *Jul 22, 2025*
Selective Prompt Anchoring for Code Generation

Bloomberg | UX Team *Jan 31, 2025*
Grounding Natural Language Database Queries with Interactive Explanations

Service

Reviewer

- ACL, EMNLP, NeurIPS, ICLR, CHI, UIST *2026*
- NAACL, IUI, UIST, TNNLS *2025*
- ACL, NAACL *2024*