

Yuan Tian

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

Education

Purdue University

Ph.D. in Computer Science (M.S. Awarded)

August 2021 – Present

GPA: 3.71 / 4.00

University at Albany, SUNY

B.S. in Computer Science

August 2019 – May 2020

GPA: 3.97 / 4.00 (Dean's List)

Chongqing University of Posts and Telecommunications

B.E. in Computer Science & Software Engineering

August 2016 – July 2019

GPA: 3.84 / 4.00 (Ranked 1st)

Work Experience

Research Assistant

Purdue University | West Lafayette, IN, USA

August 2021 – Present
(with internship/TA breaks)

- Independent research (Papers accepted to top-tier conferences)
- Build ontology-driven data extraction pipeline and knowledge graph for AI security
[NSF Proto-OKN] | [National AI Research Resource Pilot]

Applied Scientist [Intern]

Adobe, GenAI @ AEP | San Jose, CA, USA

May 2025 – August 2025

- R&D at a novel semantic enrichment system for under-specified fields
- Robust richness scoring & clarification generation & interactive evaluation by populated queries
- To integrate with the Adobe XDM system

Applied Scientist [Intern]

Adobe, GenAI @ AEP | San Jose, CA, USA

May 2024 – August 2024

- R&D at a novel text-to-SQL domain adaptation system
- Data schema management & high-quality, customized text-to-SQL data generation
- Patented & Open-sourced

Teaching Assistant

Purdue University | West Lafayette, IN, USA

January 2024 – May 2024

- Deliver 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 with a probe-and-verify manner.
- Propose a rule-based SQL verifier.
- Comprehensively analyze 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
- Integrate with Hugging Face API and support all Hugging Face LLMs
- Proposed an attention-based code generation pipeline, achieving new SOTA performance
- Received much 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
- 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

- Benchmark to test text-to-SQL robustness under schema evolution
- 10 Types of Schema Perturbations Based on a Hybrid Method (LLM + Heuristics)
- Fine-tuning LLMs on perturbed schemas benefits performance improvement

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
- 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

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

- Extended our previous text-to-SQL error analysis to include LLM analysis and attention studies.
- Demonstrate 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”, 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
- within-subjects study with 26 participants to evaluate three interactive systems and provide insights

Under Reviewed Papers

Attention-Aligned Reasoning for Large Language Models

[Under Review] | Attention Steering, Alignment, Reasoning

- Developed a novel key-value reasoning pattern based on step-by-step attention steering.
- Outperform existing prompting-based approaches and outperforms RL-based trained SLM.

Alloy: Generating Reusable Agent Workflows from User Demonstration

[Under Review] | Computer Use, Workflow Generalization, Programming by Demonstration, MCP

- An agentic system that converts computer use demonstrations into editable and reusable LLM workflows.
- Evaluate workflow generation and generalization by both a user study and a quantitative evaluation.

Large Language Models in Computational Mechanics: Toward an Autonomous Simulation Paradigm

[Under Review] | Finite Element Method, Code Generation, Multi-agent System, Finetuning

- 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

[Under Review] | 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

[Under Review] | 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 (e.g., Cursor) and user needs.

AI Chatroom for Mental Health: Generating Personalized Well-being Reports with LLMs and Taxonomy-Driven Checklists

[Under Review] | Multi-agent System, LLM-scoring, PERMA, Chatroom

- Built an agentic well-being analysis tool based on 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

[Under Review] | 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

[Under Review] | 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 Data Field Description

[Under Review] | Interactive Semantic Enrichment, Scoring System, Clarification Taxonomy

- Developed a robust richness scoring system, including 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

[Under Review] | VR, Robotics, Unity, Oculus, Haptics, Tactile Feedback, Calibration

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

Early Projects

Conversational Agent for SQL Generation

Conversational AI, Rule-based Clarification, Database Interfaces, UI

- Developed a conversational natural language interface for database query
- Developed a clarification method for disambiguation during generation and query repair after generation

Dancing Humanoid Robot

Robotics, Humanoid Robot, Calibration

- Designed the communication protocol between the upper computer and the ESP8266 chip inside the robot
- Enabled the robot to perform actions (e.g., walk) based on predefined parameters and calibrations.

Ultrasonic-powered Parking Management System

Automated Data Management, Ultrasonic Sensor

- Developed an automated parking management system based on ultrasonic sensors
- Link ultrasonic sensors, the Java virtual machine, and Microsoft SQL server through serial port communication.

Skills

- **Full-Stack Engineer** Independently developing frontend UIs, backend servers, and database systems.
- **Independent Researcher** Independently developing novel ML/NLP/HCI algorithms.

Service

Reviewer

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

Company Talks

- Datadog • [Interaction and Adaptation between Agents and Data-intensive Systems](#) April 30, 2026
- Meta • [Attention Steering with KV Cache Optimization](#) August 21, 2025
- Adobe • [Selective Prompt Anchoring for Code Generation](#) July 22, 2025
- Bloomberg • [Grounding Natural Language Database Queries with Interactive Explanations](#) January 31, 2025