Zhongyin Wang

+1 (217) 607-7987 | zw92@illinois.edu | GitHub: zoeywang02 | LinkedIn

Education

University of Illinois Urbana-Champaign — Urbana, IL, USA
Master of Science in Information Management, School of Information Sciences (GPA: 4.0/4.0) — *Aug. 2025 – May 2027*

- **Coursework**: Database Systems, Information Modeling, Software Engineering I, Statistical Models and Information.

Zhejiang University — Hangzhou, China
Bachelor in Civil Engineering (Smart Transportation & Systems) (GPA: 3.6/4.0) — *Sept. 2020 – July 2024*

- **Awards**: Zhejiang University Excellent Graduate, Structural Design Competition Second Prize.
- **Coursework**: Data Structures and Algorithm Analysis, C Programming.

Publications (Preprints & Under Review)

[1] **Truncated Reasoning Under a Budget: Evaluating System-1 Models as Decision-makers.**
Zhongyin Wang*, Zhaoqian Xue, Junting Mu, Haozheng Liu, Ye Yuan, Haolun Wu, Lingyao Li*. arXiv:2610.25768 (link); under review, ACL Rolling Review (Oct. 2026 cycle; target: NAACL 2027). **First author.**

[2] **Whose Responsibility? How People Judge Their Own and Others' Agents in A2A Conflicts.**
Hangyue Zhang†, Hanzhang (Wendy) Ma†, **Zhongyin Wang**†, Yun Huang. Under review, ACM CHI 2027. **Co-first author.**

[3] **InsightToChat: Supporting Post-Videoconferencing Reflection through a Conversational AI Agent.**
Shaoze Zhou, **Zhongyin Wang**, Junting Mu, Bojun Zhang, Rawan Alghofaili, Janet G. Johnson, Lingyao Li, Renkai Ma, Christine Lisetti, Chen Chen. Under review, ACM CHI 2027.

[4] **Decision Models as Judges for Hallucination Detection.**
Haozheng Liu*, Zhaoqian Xue, **Zhongyin Wang**, Junting Mu, Zihui Ma, Shuya Feng, Lingyao Li*. Under review, ACL Rolling Review (Oct. 2026 cycle; target: NAACL 2027).

[5] **ReviewPilot: A Multi-Agent System for Screening and Data Extraction in Systematic Reviews.**
Zhaoqian Xue, Shaoze Zhou, Xuexin Li, Junting Mu, Haozheng Liu, **Zhongyin Wang**, Yichun He, Renkai Ma, Hao Zhou, Sarah Bratt, Siyuan Ma, Chen Chen, Lingyao Li. arXiv:2610.21042 (link); under review, ACL Rolling Review (Oct. 2026 cycle; target: NAACL 2027).

†Equal contribution. *Corresponding author.

Research Experience & Selected Projects

Promises Across Layers: Safeguard Alignment & User Agency in AI Chatbots — Oct. 2026 – Present
Project Lead | Advisor: Prof. Ge (Tiffany) Wang (University of Illinois Urbana-Champaign)

- Leading this project independently: evaluating how chatbot providers' publicly stated safeguards against manipulative conversational and interface designs align with their implementation in chatbot responses and interface controls.
- Studying whether users can discover, understand, and exercise these safeguards, and what shapes the gap between intended protections and actual user control; targeting CSCW/ToCHI.

Efficient & Reliable LLM Decision-Making: Reasoning, Judging, Reviews — Aug. 2026 – Present
Research Assistant | Advisor: Prof. Lingyao Li (University of Arizona)

- First-authored a study of **truncated reasoning under token budgets**, using a System-1 decision model (Jev) to answer from the partial traces of seven LLM reasoners across five multiple-choice benchmarks. At a 1,500-token budget, accuracy improved by 4.1 pp over forced answering, and offline replay estimated 34.5% of full-run cost with a 1.13 pp accuracy loss.
- Analyzed **cost–accuracy trade-offs** of static budgets and adaptive stopping policies; a policy selected on training folds saved 71.3% of estimated full-run cost with a 0.58 pp accuracy loss on held-out folds, with significance tested via Holm-corrected bootstrap tests.
- Contributed to the experiment design and data analysis of an **LLM-as-a-Judge evaluation** of decision models for hallucination detection across five hallucination types (9,852 items, six datasets), assessing accuracy, probability calibration, and confidence-based cascades against a generative LLM judge at 1–3% of its cost.
- Co-authored **ReviewPilot**, a conversational multi-agent system that automates the PRISMA workflow from screening to information extraction; responsible for the **data extraction** component, which attaches a source to each extracted value (extraction F1 up to 0.788, outperforming a single-prompt LLM and AgentSLR on epidemiological expert reviews).

User Agency, Delegation & Responsibility in AI Agents — May 2026 – Present
Research Assistant | Advisor: Prof. Yun Huang (University of Illinois Urbana-Champaign)

- Investigating dark patterns in AI-agent interactions, focusing on user awareness, delegation, and authorization in algorithmic decision-making.
- Co-first-authored a CHI 2027 submission on responsibility in agent-to-agent (A2A) conflicts: a **design fiction study** (eight speculative scenarios) with semi-structured interviews of 16 students, examining how people interpret responsibility, delegation, and manipulation, with implications for traceability and enforceable delegation boundaries.
- Designing and evaluating an AI-assisted platform that supports research pathway discovery and collaborative team formation.
- Developing LLM-based workflows and AI-agent collaboration strategies to generate personalized research and team recommendations.

Nexus: LLM-based Academic Planning System Mar. 2026 – Present

- Developed an intelligent scheduling system that automates the parsing of heterogeneous syllabus data (PDF/JSON) from Canvas and Coursera using **LLMs** for semantic extraction.
- Engineered a conflict-detection algorithm to generate optimized study plans, utilizing Redis caching for course metadata to ensure low-latency API responses.
- Designed an LLM-based agent workflow using Spring AI to coordinate task decomposition and personalized schedule generation.

LLM-based Multi-table Learning for Credit Risk Modeling Sept. 2024 – Dec. 2024

Research Assistant | Advisor: Prof. Yang Yang (Zhejiang University)

- Replicated the downstream regression task from the paper *Towards Cross-Table Masked Pretraining for Web Data Mining* to validate the effectiveness of cross-table pretraining.
- Extended the original row-level formulation to support table-level regression and classification on heterogeneous multi-table datasets.
- Implemented data preprocessing and sampling pipelines for cross-table representation learning across structured financial datasets.
- Evaluated model performance under different sampling strategies to support credit risk prediction.

Transformer-based Urban Road Network Design (Generative AI) Dec. 2023 – June 2024

Bachelor's Thesis | Advisor: Prof. Hongsheng Qi (Zhejiang University)

- Developed a generative AI pipeline using **Transformer** and **Graph Neural Networks (GNN)** to automate the design of urban road networks from zero-to-one.
- Processed large-scale geospatial data via OSM APIs, converting raw geographic nodes into optimized **graph-based topologies** to evaluate connectivity and network resilience.
- Validated model performance by comparing generated layouts with existing traffic networks to support autonomous driving simulations.

Smart Tunnel Dynamic Emergency Response System May 2022 – May 2023

Project Leader | Advisor: Prof. Ke Wu (Zhejiang University)

- Integrated tunnel fire dynamics models with multi-sensor data fusion to predict fire temperature fields and smoke spread in real-time.
- Leveraged **YOLO** algorithms and computer vision to automate fire source detection and scale estimation, reconstructing fire scenes based on physical dynamics equations.
- Developed a dynamic response platform using **Java** that allows real-time emergency coordination and response for critical infrastructure safety.

Industry Experience

IDX Exchange | AI Agentic Engineer Intern Boise, ID (Remote) | Jun. 2026 – Present

- Built a multi-agent AI assistant with OpenClaw over 58K+ MLS records, implementing NL query parsing, SQL search, and multi-turn conversational memory.
- Developed a hybrid recommendation engine (structured + OpenAI embedding similarity) and a RAG knowledge assistant grounded in real MLS documentation.

Baidu, Inc. | NLP Algorithm Engineer Intern Beijing, China | Apr. 2024 – June 2024

- Evaluated representative large language models through API-based experiments, analyzing inference latency, token efficiency, and long-context performance across multiple evaluation metrics.
- Investigated representative LLM agent frameworks through literature review, open-source implementation analysis, and architectural comparison.

Technical Skills

Programming Languages: Python, Java, C/C++, SQL, JavaScript, Cypher
AI: LLMs, AI Agents, LLM Evaluation, Prompt Engineering
Databases: Neo4j, MySQL, PostgreSQL, Redis
Software & Tools: Spring Boot, RESTful APIs, Docker, Git, Linux