ZEWEI DENG

Ph.D. Student in Industrial and Systems Engineering | University of Minnesota Twin Cities
deng0379@umn.edu | Website | Google Scholar | GitHub
U.S. Work Authorization · No Employer Sponsorship or CPT Required
Summer 2027 Availability: May–September 2027

PROFILE

Develop statistical, coding-theoretic, and optimization methods for trustworthy AI, with research on LLM watermarking and provenance, distributionally robust inference, uncertainty-aware sequential modeling, and temporal point processes.

EDUCATION

University of Minnesota Twin Cities **2025–Present**
Ph.D. in Industrial and Systems Engineering; Advisor: Prof. Liyan Xie *Expected graduation: Dec 2028 · Minneapolis, MN*

University of Chicago **2023–2024**
M.S. in Applied Data Science *Chicago, IL*

University of Illinois Urbana-Champaign **2019–2023**
B.S. in Statistics; Minors in Computer Science and Economics; GPA: 3.9/4.0 *Urbana-Champaign, IL*

RESEARCH & PROFESSIONAL EXPERIENCE

Graduate Research Assistant, University of Minnesota Twin Cities **Summer 2026–Present**
Advisor: Prof. Liyan Xie *Minneapolis, MN*

- Develop statistical and coding-theoretic methods for LLM watermark provenance, edit detection, local integrity verification, and watermarkability.
- Led an ECC watermark project accepted to EMNLP 2026 Findings, spanning method design, implementation, experiments, and manuscript preparation.
- Develop uncertainty-aware sequential models and distributionally robust inference for delayed point processes under reporting delays and model misspecification.

Research Assistant, Chicago Booth School of Business **2024**
Advisor: Prof. Dacheng Xiu *Chicago, IL*

- Built financial-text forecasting pipelines using BERT, LLaMA, and T5 representations with regression models for return and debt-related prediction.

Machine Learning Capstone, Alliant Credit Union / UChicago DSI **2024**
Client-facing machine-learning project *Chicago, IL*

- Contributed to model development and evaluation for a credit-union analytics application.

PUBLICATIONS & WORKING PAPERS

Detecting Post-generation Edits to Watermarked LLM Outputs via Error-Correcting Codes.
Zewei Deng, Muhammad Siddeek, Liyan Xie, Mohamed Seif, Mengdi Wang, H. Vincent Poor, and Andrea Goldsmith.
Findings of EMNLP 2026, pp. 45788–45796. [Paper]

Anchor-ECC: Local Integrity Checking for Watermarked LLM Outputs via Error-Correcting Codes.
Zewei Deng, Muhammad Siddeek, Liyan Xie, Mohamed Seif, Mengdi Wang, H. Vincent Poor, and Andrea Goldsmith.
arXiv:2609.38722, 2026; submitted to Transactions on Machine Learning Research (TMLR). [Paper]

Distributionally Robust Inference for Delayed Point Processes.
Zewei Deng, Shuyi Chen, Shixiang Zhu, and Liyan Xie.
Working paper in preparation for submission to the Annals of Applied Statistics (AOAS); workshop version forthcoming at the INFORMS Data Mining and Decision Analytics Workshop, 2026.

Watermarkability: Measuring and Improving Language-Model Compatibility with Logit-Based Watermarks.
Zewei Deng, Mohamed Seif, and Liyan Xie.
arXiv:2610.22810, 2026. [Paper]

HawkesLLM: Semantic Uncertainty Propagation in Agentic Text Simulation.
Zewei Deng, Tinghan Ye, and Liyan Xie.
ICML Workshop on Statistical Frameworks for Uncertainty in Agentic Systems, 2026. [Paper]

Learning to Test: Physics-Informed Representation for Dynamical Instability Detection.
Minxing Zheng, **Zewei Deng**, Liyan Xie, and Shixiang Zhu.
arXiv:2604.10967, 2026. [Paper]

TEACHING & HONORS

Teaching: Teaching Assistant, University of Minnesota Twin Cities, Sep 2025–May 2026

Honors: Highest Distinction in Statistics; Dean's List (2020–2023); COMAP MCM Honorable Mention (2021)