

Ningzhi Tang

✉ ntang@nd.edu  nztang.com  [TTangNingzhi](https://github.com/TTangNingzhi)  [in/ningzhi-tang](https://www.linkedin.com/in/ningzhi-tang)  [Google Scholar](#)

Ph.D. Candidate in CSE · Coding Agents, Human-AI Interaction · 10+ publications (EMNLP, UIST, ASE)

INDUSTRY EXPERIENCE

Google, User Experience Engineer (UXE) Intern

May–Aug 2026

Host: Tao Dong, Co-hosts: Alan Ross, Sam Cheng

San Jose, CA

- Mapped a 16-dimensional interaction design space for trajectory-grounded agent customization, distilling agent misbehaviors and developer preferences from past trajectories into rules for AGENTS.md and skills.
- Built a runnable probe: an agentic pipeline that analyzes a developer’s own coding-agent sessions and generates an interactive, data-grounded website to explore the design space.
- Delivered per-team design insights to builders of internal customization pipelines and engineering teams.
- Received 2 peer-recognition awards (gThanks) from a partner team for building a trajectory-annotation dashboard and piloting a diary study to collect ground-truth agent-behavior labels.

EDUCATION

University of Notre Dame

Aug 2023–Present

Ph.D. Candidate in Computer Science and Engineering

South Bend, IN

Advisor: Toby Jia-Jun Li, GPA: 4.0/4.0  Edison Innovation Fellowship (\$16,500)

Southern University of Science and Technology

Aug 2019–Jun 2023

B.E. in Computer Science and Engineering

Shenzhen, China

Advisor: Yuhui Shi, GPA: 3.93/4.0, Rank: 1/207

Undergraduate research on recommender systems published at [SIGIR 2022](#) and [WWW 2023](#).

SELECTED RESEARCH & PROJECTS

Large-Scale Trajectory Analysis of Developer Interaction with Coding Agents

2026

- Created the first [data-scraping pipeline](#) for developers’ in-the-wild chats with coding agents (e.g., Cursor, Copilot) from public GitHub repositories, collecting 14,789 real-world trajectories across 1,441 repositories.
- ([ASE 2026](#)) Characterized developer behavioral intents via a two-level taxonomy of 75K prompts with human validation, and modeled session structure using sequence clustering and Markov transition analysis.
- ([EMNLP 2026](#)) Analyzed failure patterns of developer-agent misalignment (surfaced as user pushback) using an LLM-based pipeline to extract and annotate structured records covering forms, causes, costs, and resolutions ([interactive data viewer](#)), informing future agent evaluation and interaction design.

NaturalEdit: Natural Language Summaries as Interaction Surface for Code Modification

2025–2026

- ([VL/HCC 2025](#)) Conducted a 15-developer study on editing code through natural language summaries, as a precursor to “spec-driven development,” identifying key developer experience and design requirements.
- ([UIST 2026](#)) Built NaturalEdit, a [VS Code extension](#) that treats adaptive summaries as an editable interface to code, with AST-based NL-to-code mapping, format switching, and incremental summary diffs. Validated the technical pipeline on two LLM code-editing benchmarks and demonstrated its utility in a 20-developer study, with reported improvements in comprehension, intent articulation, and sense of control.

CodeGRITS: Developer Behavior and Eye-Tracking Infrastructure for JetBrains IDEs

2023–2024

- ([ICSE 2024 Demo](#)) Developed [CodeGRITS](#), an open-source [JetBrains plugin](#) (30+ GitHub stars, 300+ Marketplace downloads) that captures developer IDE interactions and maps eye-gaze coordinates to source-code elements (line, col, AST token), as a research instrument adopted by 5+ published studies.
- ([VL/HCC 2024](#)) Conducted a 28-participant controlled study of how developers validate and repair LLM-generated code, using CodeGRITS to examine their behavioral and eye-tracking patterns.