

Chenyu Wang 王晨语

jamesnolan17.github.io

LLMs for Software Engineering • Coding Agents & Inference Efficiency • AI Security

Singapore PR • chenyuwang@smu.edu.sg • +65 9396 5944 • github.com/JamesNolan17 • Google Scholar • LinkedIn

PhD candidate and Research Engineer at SMU's SOAR Lab (advisor Prof. David Lo). First-author on three lines: **(1) inference-time control of coding agents** – a 0.6B monitor that either early-terminates doomed runs for **15-20%** token savings or smart-restarts them to lift SWE-bench resolution **66.6%→71.8%** (+5.2 pts; under review, **CORE A***); **(2) code-model security** – data-poisoning backdoors (ASE 2025, **CORE A***); **(3) industry AI-quality practices** (EASE 2026; featured in *IEEE Software*). Co-authored at **ICSE, TSE, MSR**. Open to full-time roles from **Fall 2027**.

Experience

Singapore Management University – SOAR Lab

2023 – Present

Research Engineer • PI: Prof. David Lo

Singapore

- Lead the lab's **inference-time agent-control** and **code-model-security** research line as first author, taking each project **end-to-end** – empirical study, dataset construction, model training & fine-tuning, evaluation, and benchmarking tooling.
- Design, with **GovTech Singapore** (Jul 2026–), a **repository-level, agentic vulnerability-detection** pipeline – **code property graphs** for interprocedural analysis plus an **LLM-as-a-judge** stage that prunes candidate vulnerable paths to focus an expensive detection loop, designed around real CI/CD cost constraints.

Education

Singapore Management University – Ph.D., Computer Science (part-time)

2023 – 2027 (expected)

Software Engineering & Trustworthy AI • Advisor: Prof. David Lo

Singapore

Singapore University of Technology & Design – B.Eng., Computer Science & Design

May 2019 – Aug 2022

GPA **4.83/5.0**, Highest Distinction • Minor in Artificial Intelligence • NUS exchange (2021)

Singapore

Selected Publications (first-author)

Fail-Fast, Restart-Smart: Early Failure Prediction and Restart for SWE Agentic Tasks

Under review, 2027 **CORE A***

A two-stage controller for SWE agents with two complementary modes. **Fail-Fast (cost)**: a lightweight **0.6B monitor** early-terminates doomed runs from the trajectory prefix – **14.6-20.4% token savings** at a 5% false-positive rate, **transferring across 4 policies**, including a closed-API model. **Restart-Smart (quality)**: on an alarm, relaunch a fresh rollout offering the prior repo diff as an optional overlay – SWE-bench Verified resolution **66.6% → 71.8%**. arXiv:2608.03222

Backdoors in Code Summarizers: How Bad Is It?

ASE 2025 **CORE A***

First **systematic (9-factor)** study of data-poisoning backdoors in code LLMs: just **20 of 454,451 samples (0.004%)** – **25×** below rates assumed safe – implant a backdoor the spectral-signature defense **fails to remove**, and **smaller training batches make the attack stronger**. A concrete, under-estimated AI supply-chain risk. arXiv:2506.01825

Industry Practitioners' Perspectives on AI Model Quality

EASE 2026 • featured in *IEEE Software* **CORE A**

Mixed-methods study (**15 interviews + 50-practitioner survey**) of how industry actually manages AI quality – e.g., **most practitioners never monitor deployed models for data drift**, a critical blind spot for production ML. arXiv:2402.16391

Additional Publications (co-authored)

- Unveiling Memorization in Code Models ICSE 2024 **CORE A*** arXiv
- Gotcha! This Model Uses My Code! Evaluating Membership Leakage Risks in Code Models IEEE TSE 2024 **Top journal** arXiv
- What Do Users Ask in Open-Source AI Repositories? An Empirical Study of GitHub Issues MSR 2023 **CORE A** arXiv

Awards & Recognition

- SCIS PhD Research Excellence Award (Tier 1)**, SMU – for a first-author A* paper; AY2026 (50% tuition-fee waiver).
- Best Student Award**, SUTD – ranked **1st of 400+ students** (Term 3) • Honours with Highest Distinction • Honours List (multiple terms).
- Senior Middle 2 Scholarship**, Singapore Ministry of Education • Selected for **NUS exchange** via SUSEP (7 students from pillar).
- Reviewer**: ACM TOSEM (×2), MSR 2026 • EASE 2026 paper **featured in IEEE Software** (Practitioner's Digest, Nov/Dec 2026).

Technical Skills

AI-Agent Orchestration & AI-Native Software Engineering: design and orchestrate distributed multi-agent workflows for research and software engineering – autonomous task decomposition, coding-agent collaboration, literature analysis, experiment automation, multi-server GPU execution, and cross-agent synchronization via Git and cloud storage; the pipeline behind my first-author experiments (*above*), with sharp judgment for what to offload to agents vs. own.

Languages: Python, Java, JavaScript (*proficient*); C, TypeScript, Shell (*working*)

ML / LLM: PyTorch, Hugging Face, LoRA / PEFT fine-tuning, LLM agents & SWE-bench benchmarking, code LLMs (Qwen, CodeT5, Gemma)

Program Analysis & Security: code property graphs, program slicing, taint analysis, backdoor & adversarial robustness, vulnerability detection

Engineering: Docker, Git, Linux, React / full-stack web, reproducible data pipelines