HOANG ANH NGUYEN

Published as Anh Nguyen Hoang and Nguyen Hoang Anh

+84 398 272 620 ◇ Hanoi, Vietnam

anhnh.4work@gmail.com ◇ Homepage ◇ GitHub

RESEARCH INTERESTS

My research focuses on AI for Software Engineering, particularly foundation models for code, autonomous coding agents, and long-horizon software evolution. I aim to study principled approaches for training and adapting code language models and for designing reliable agentic systems for repository-scale tasks. I am also open to related areas.

EDUCATION

Hanoi University of Science and Technology 2020 – 2024

Bachelor of Science in Computer Science, GPA: 3.75/4.0 Ranked 15th of 296 students (Top 5%)

PUBLICATIONS

[1] **CodeWiki: Evaluating AI's Ability to Generate Holistic Documentation for Large-Scale Codebases**
Anh Nguyen Hoang, *Minh Le-Anh, Bach Le, Nghi D. Q. Bui*
Findings of ACL 2026 [Paper] [Code]

[2] **Beyond Repository Boundaries: Cross-Repository Graph Retrieval for Code Generation**
Minh Le-Anh, Nam Le Hai, Quyen Tran, ***Anh Nguyen Hoang****, Linh Ngo Van, Bach Le, Nghi D. Q. Bui*
Findings of EMNLP 2026 [Paper]

[3] **GeoDPO: Geometric Token Importance for Stable LLM Alignment**
Dat Phi Van, ***Anh Nguyen Hoang****, Linh Ngo Van, Thien Huu Nguyen, Thanh-Toan Do, Thanh Hong Nguyen, Trung Le*
EACL 2027

[4] **Mutual-Pairing Data Augmentation for Few-Shot Continual Relation Extraction**
Nguyen Hoang Anh*, Quyen Tran*, Thanh Xuan Nguyen*, Nguyen Thi Ngoc Diep, Linh Ngo Van, Thien Huu Nguyen, Trung Le*
NAACL 2025 [Paper] [Code]

[5] **Preserving Generalization of Language Models in Few-Shot Continual Relation Extraction**
Quyen Tran, Nguyen Xuan Thanh*,* ***Nguyen Hoang Anh****, Nam Le Hai, Trung Le, Linh Van Ngo, Thien Huu Nguyen*
EMNLP 2024 [Paper] [Code]

[6] **Enhancing Few-Shot Continual Relation Extraction via Pairwise Augmentation and Guided Sharpness-Aware Optimization**
Hoang Thanh Dat,* ***Nguyen Hoang Anh****, Nam Le Hai*, Linh Ngo Van, Sang Dinh*
Artificial Intelligence (AIJ), accepted [Paper]

** denotes equal contribution.*

PREPRINTS AND UNDER REVIEW

[1] **Secure Agentic Coding through Counterexample-Grounded Feedback (FALCON)**
Tue Le, ***Anh Nguyen Hoang****, Gabriel Sasseville, Doseok Jang, Pierre-Luc Bacon, Nghi D. Q. Bui*
Under review [Paper] [Code]

[2] **Can Agentic Code Review Catch Vulnerabilities Before They Are Merged?**
Yihan Yang, Chenyuan Yang, Zijie Zhao, ***Anh Nguyen Hoang****, Ziqi Zhang, Weidong Wang, Zhiheng Lyu, Lingming Zhang*
Under review [Paper] [Code]

[3] **RustPrint: Documentation-Guided Agentic Codebase Migration from C to Rust**
Minh Le-Anh, ***Anh Nguyen Hoang****, Bach Le, Nghi D. Q. Bui*
Preprint, 2026 [Preprint]

HONORS AND AWARDS

Champion, AI for Real Action 2025 (Team Converged) January 2026

First prize among nearly 400 teams in Vietnam's first national artificial intelligence competition

- **Techcombank Overseas PhD Study Sponsorship:** awarded for winning the competition; up to USD 333,000 for overseas PhD study (April 2026) [Offer Letter]
- **Vietnamese education LLM:** fully fine-tuned Qwen3-8B on 2K web and textbook preference pairs with a weighted SFT + DPO loss (my proposal) to learn new knowledge and answer in natural, respectful Vietnamese; scored about 12 percentage points above base Qwen3-8B on the competition's education benchmark
- **STEM animation platform:** built the Manim code generation that turns a STEM concept into an explainer video in about 2 minutes: the fine-tuned model writes a scenario, then Manim code fixed in a render loop

EXPERIENCE

FPT Software AI Center May 2025 – Present
AI Research Resident *Hanoi, Vietnam*
Advisors: Associate Prof. Bach Le, University of Melbourne; Assistant Prof. Nghi D. Q. Bui, VinUniversity

- **CodeWiki** (first author, Findings of ACL 2026): built an agentic framework for repository-level documentation across seven languages by decomposing the dependency graph into a module tree, delegating leaf modules to documentation agents, and synthesizing parent-level text, architecture and data-flow diagrams bottom-up
- Built **CodeWikiBench**, a benchmark with hierarchical rubrics and agentic LLM-judge evaluation; CodeWiki reaches a 68.79% quality score, +4.73 points over closed-source DeepWiki and +21.66 points over OpenDeepWiki, with gains of up to 18.54% on individual repositories
- **RustPrint**: co-developed agentic C-to-Rust migration that uses CodeWiki documentation as a blueprint and documentation mismatches as repair signals; on eight C repositories (11K–84K LoC) it compiles every target where prior LLM translators fail and, with Kimi-K2, beats Claude Code on feature preservation (93.26% vs. 52.52%) and test pass rate (95.17% vs. 79.85%)
- **FALCON** (under review): co-developed a security verifier that separates attack exploration (proof-of-concept probes) from judging their execution traces; as repair feedback it raises CWEval Func-Sec@1 by 25.2 points where LLM-as-a-judge brings almost no gain, and as an RL reward it beats the REAL and SecCodePRM rewards
- **SecPRBench** (under review, with Prof. Lingming Zhang, UIUC): curated the FFmpeg, FreeType, and SQLite tasks of a benchmark of 265 vulnerability-introducing changes from nine C, C++, and Rust projects, each paired with its repaired version; agents on eight model backends detect only 32.5%–78.1% of the target vulnerabilities

Research Assistant March 2024 – April 2025
Continual Learning and Natural Language Processing Research *Hanoi, Vietnam*
Advisors: Assistant Prof. Linh Ngo Van; Associate Prof. Thien Huu Nguyen

- **Mutual-pairing data augmentation** (co-first author, NAACL 2025): combined old and new relation samples into composite inputs with objectives that induce adversarial-training effects, cutting the accuracy drop to 26.40% on FewRel and 24.36% on TACRED (2.49% better than CPL on TACRED); showed that Sharpness-Aware Minimization costs over 6% accuracy with LLM2Vec
- **Generalization preservation** (co-first author, EMNLP 2024): reused discarded pre-trained LM heads through mutual-information maximization with the classifier, improving final accuracy by about 2% for CPL and ConPL, 1% for SCKD, and up to 3.75% with LLaMA-2-7B

Oraichain Labs May 2024 – April 2025
AI Engineer *Hanoi, Vietnam*

- Led the development of Cupiee, an emotional AI companion platform with 300,000+ users and about 100,000 messages per day
- Built Cupiee's voice-agent companion; 30-day user retention rose from 13% to 47% after its release

SELECTED OPEN-SOURCE SOFTWARE

CodeWiki 1,700+ GitHub stars
Repository-level documentation framework for seven languages, with CodeWikiBench

XTTSv2-Finetuning-for-New-Languages 200+ GitHub stars
Adapts XTTSv2 text-to-speech to any language, with a released Vietnamese model; sliding-window streaming cuts time to first audio from about 1 s to 200 ms on the same hardware

LANGUAGES

Vietnamese (native), English (IELTS 7.0, TOEFL 5.0)