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 interests center on post-training large language models for long-horizon agentic tasks. I aim to study how models learn to maintain relevant information, use tools, and revise decisions from demonstrations and trustworthy feedback, while generalizing to new tasks and retaining earlier capabilities. My background spans continual learning, model adaptation, and software agents.

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%)

SELECTED POST-TRAINING PROJECT

Vietnamese Education LLM AI for Real Action 2025
Proposed a weighted objective combining supervised fine-tuning (SFT) and direct preference optimization (DPO); fully fine-tuned Qwen3-8B on 2K web and textbook preference pairs for educational knowledge and natural, respectful Vietnamese responses. Improved the competition's education benchmark by about 12 percentage points over the base model.

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] **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]

[3] **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]

[4] **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]

[5] **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

[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.*

RESEARCH EXPERIENCE: CONTINUAL LEARNING

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

- Studied **few-shot continual relation extraction**: learning new relation types from limited labeled examples while retaining previously learned relations and pretrained language knowledge
- **Mutual-pairing data augmentation** (co-first author, NAACL 2025): designed augmentation and training objectives combining old and new relation examples into richer inputs without additional labels; reduced the accuracy drop to 26.40% on FewRel and 24.36% on TACRED, with a 2.49% improvement over CPL on TACRED
- **Preserving pretrained generalization** (co-first author, EMNLP 2024): reused pretrained LM heads through mutual-information maximization with the task classifier; improved final accuracy by about 2% for CPL and ConPL, 1% for SCKD, and up to 3.75% with LLaMA-2-7B
- Investigated the interaction between representation learning and continual adaptation; found that Sharpness-Aware Minimization reduced accuracy by over 6% when used with LLM2Vec

RESEARCH EXPERIENCE: LANGUAGE-MODEL AGENTS

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 LLM-agent framework that generates repository-level documentation across seven programming languages through hierarchical context decomposition and bottom-up synthesis of module explanations
- **CodeWikiBench**: built hierarchical evaluation rubrics and agentic LLM-judge evaluation; CodeWiki achieves a 68.79% quality score, +4.73 points over DeepWiki and +21.66 points over OpenDeepWiki
- **FALCON** (under review): refined execution-grounded checks that reject fabricated evidence in generated attack tests and link reported events to target execution; the framework improved CWEval Func-Sec@1 by 25.2 points as repair feedback
- **RustPrint** (preprint): co-developed documentation-guided C-to-Rust migration; with Kimi-K2, preserved 93.26% of evaluated features across eight C repositories, compared with 52.52% for a Claude Code baseline
- **SecPRBench** (under review, with Prof. Lingming Zhang, UIUC): curated FFmpeg, FreeType, and SQLite tasks for a benchmark of 265 vulnerability-introducing code changes paired with their repairs

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]

INDUSTRY EXPERIENCE

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

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

HONORS AND AWARDS

Champion, AI for Real Action 2025 (Team Converged) January 2026
First prize among nearly 400 teams in Vietnam's first national AI competition

- **Techcombank Overseas PhD Study Sponsorship**: up to USD 333,000 for overseas PhD study, awarded following the competition (April 2026) [Offer Letter]

SELECTED OPEN-SOURCE SOFTWARE

XTTSv2-Finetuning-for-New-Languages 200+ GitHub stars
Toolkit for adapting XTTSv2 to new languages, including a released Vietnamese model; sliding-window streaming reduces time to first audio from about 1 s to 200 ms

LANGUAGES

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