Yifan Zhang

U.S. Green Card holder | Homepage | LinkedIn | Twitter | Email | Phone: +1 919 638-0501

Research Interests

I make AI agents for coding and computer use reliable by grounding them in formal, checkable evidence. Many "verified" training signals are still one LLM judging another; I use proof checkers, compilers, static analyzers, execution, and conformal certification to train agents, vet their self-improvements, and control them at run time.

- **Formal Signals for Post-Training**: Sparse task success and LLM judges are weak teachers for agents. I work on SFT and RL in which every reward is backed by a checker, from certified verifiers and token-level supervision to dynamic RL environments whose tasks and rewards are proof- and type-checked.
- **Proofs for Recursive Self-Improvement**: Agents that improve themselves can also game or degrade themselves. I study self-improvement in which an agent may rewrite its own harness and plans, but every change must pass an independent check, and ultimately a proof checker.
- **Tooling for Reliable Agents**: A frontier model is only as reliable as the system around it. I build that layer: harnesses, simulated environments, retrieval, expert-level evaluation, and in-session routing that sends each step to the right model, tool, or checker.

Legally authorized to work in the U.S. without employer sponsorship | Last updated: 11/24/2026

Education

Vanderbilt University — *May 2022 – May 2026*
Ph.D. in Computer Science, Specializing in AI for Software Engineering — *Nashville, TN, USA*
- Advised by Prof. Kevin J. Leach & Prof. Yu Huang

Georgia Institute of Technology — *Aug. 2022 – Aug. 2026*
M.Sc. in Computer Science, Specializing in Interactive Intelligence/Artificial Intelligence — *Atlanta, GA, USA*
- Also completed electives in ML (DL/RL/NLP) & Computing Systems tracks

China University of Petroleum — *Sep. 2012 – Jun. 2019*
M.Eng. in Petroleum Engineering, Specializing in Industrial & System Engineering — *Beijing, CN*
- Double majored in English (TEM-8 holder) & minored in British Parliamentary debates

University of Calgary — *Dec. 2015 – Jun. 2016*
Undergraduate Exchange Program in Petroleum Engineering — *Calgary, AB, CA*
- Fully-funded by Chinese national fellowship for overseas studies

Publications (*Alphabetical; †Mentor)

Refereed Long Journal & Conference Papers — *18 papers*
NeurIPS, ACL, FSE, AAAI, WWW, NAACL, EACL, EuroS&P, TOSEM, TMLR, and others — *2017 – 2027*

- **EACL'27** | Manish Acharya, Zhenyu Liao, Yueke Zhang, Kevin Leach, Yu Huang, Yifan Zhang: *VERITAS: Verifier-Guided Proof Search for Zero-Shot Formal Theorem Proving.* In Findings of the 20th Conference of the European Chapter of the Association for Computational Linguistics, March 9-14, 2027.
  CORE A | Presented at AI4Math@ICML'26 | Links: [arXiv] [artifact] [poster] [slides]
- **NeurIPS'26** | Yifan Zhang, Yutong Dai, Viraj Prabhu, Zhiyuan Hu, Ran Xu, Zeyuan Chen: *CLIFT: Conformal Self-Verification for Web Agent Training and Test-Time Scaling.* In Proceedings of the 40th Annual Conference on Neural Information Processing Systems, December 6-12, 2026.
  CORE A* | **Featured in Sina Finance** | Links: [Paper] [arXiv] [poster] [slides] [video] [news]
- **TOSEM'26** | Yueke Zhang, Yifan Zhang, Wisniewski Pamela, Fengwei Zhang, Kevin Leach, Yu Huang: *Understanding and Detecting GitHub Impersonation through Automated Authorship Attribution.* The ACM Transactions on Software Engineering and Methodology, 2026.
  JCR Q1 | Links: [Paper]

- **ACL'26** | Yifan Zhang, Chen Huang, Yueke Zhang, Jiahao Zhang, Toby Jia-Jun Li, Collin McMillan, Kevin Leach, Yu Huang: *EyeMulator: Improving Code Language Models by Mimicking Human Visual Attention.* In Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics, July 2-7, 2026.
  CORE A* | **Oral Presentation** | Links: [Paper] [arXiv] [artifact] [poster] [slides] [video]
- **ACL'26** | Yifan Zhang, Jieyu Li, Kexin Pei, Kevin Leach, Yu Huang: *SynthFix: A Hybrid Neuro-Symbolic Framework for Code Vulnerability Repair.* In Findings of the 64th Annual Meeting of the Association for Computational Linguistics, July 2-7, 2026.
  CORE A* | Links: [Paper] [arXiv] [artifact] [poster] [slides] [video]
- **FSE'26** | Yueke Zhang*, Yifan Zhang*†, Zihan Fang, Greg Trafton, Daniel Levin, Kevin Leach, Yu Huang: *ScanCoder: Leveraging Human Attention Patterns to Enhance LLMs for Code.* In Proceedings of the ACM International Conference on the Foundations of Software Engineering, July 5-9, 2026.
  CORE A* | Links: [Paper] [slides]
- **ICPC'26** | Jiahao Zhang, Yifan Zhang†, Kevin Leach, Yu Huang: *EyeLayer: Integrating Human Attention Patterns into LLM-Based Code Summarization.* In Proceedings of the 34th IEEE/ACM International Conference on Program Comprehension. April 12-13, 2026.
  CORE A | Links: [arXiv] [video]
- **AAAI'26** | Yifan Zhang, Giridhar Ganapavarapu, Srideepika Jayaraman, Bhavna Agrawal, Dhaval Patel, Achille Fokoue: *SPIRAL: Symbolic LLM Planning via Grounded and Reflective Search.* In Proceedings of the Fortieth AAAI Conference on Artificial Intelligence, January 20-January 27, 2026.
  CORE A* | **IBM AssetOpsBench for Industry 4.0** | Links: [Paper] [arXiv] [artifact] [poster] [slides] [video]
- **TMLR'25** | Yifan Zhang, Chen Huang, Yueke Zhang, Huajie (Jay) Shao, Kevin Leach, Yu Huang: *Pre-Training Representations of Binary Code Using Contrastive Learning.* The Transactions on Machine Learning Research, 2025.
  Links: [Paper] [arXiv] [artifact] [slides] [video]
- **ACL'25** | Xinwei Yang, Zhaofeng Liu, Chen Huang, Jiashuai Zhang, Tong Zhang, Yifan Zhang, Wenqiang Lei: *ELABORATION: A Comprehensive Benchmark on Human-LLM Competitive Programming.* In Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics, July 27-August 1, 2025.
  CORE A* | Links: [Paper] [arXiv] [artifact]
- **EuroS&P'25** | Jiliang (Eric) Li*, Yifan Zhang*†, Yu Huang, Kevin Leach: *MalMixer: Few-Shot Malware Classification with Retrieval-Augmented Semi-Supervised Learning.* In Proceedings of the 10th IEEE European Symposium on Security and Privacy, June 30-July 4, 2025.
  CORE A | Links: [Paper] [arXiv] [artifact] [poster] [slides] [video]
- **NAACL'25** | Yifan Zhang, Xue (Sharon) Yang: *CONSTRUCTA: Automating Commercial Construction Schedules in Fabrication Facilities with Large Language Models.* In Proceedings of the 2025 Annual Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics. April 29–May 4, 2025. Industry Track Paper.
  CORE A | **Intel Foundry Collaboration** | Links: [Paper] [arXiv] [poster] [slides] [video]
- **TMech'24** | Yichang He, Yifan Zhang, Yunpeng Fan, U-Xuan Tan: *Real-Time Vibration Compensation with Long Short-Term Memory Recurrent Neural Network and Adaptive Filter.* The IEEE/ASME Transactions on Mechatronics, 2024.
  JCR Q1 | Links: [Paper]
- **TOSEM'24** | Zachary Karas, Aakash Bansal, Yifan Zhang, Toby Jia-Jun Li, Collin McMillan, Yu Huang: *A Tale of Two Comprehensions? Studying Human Attention during Code Summarization.* The ACM Transactions on Software Engineering and Methodology, 2024.
  JCR Q1 | Presented at ICSE'25 Journal-First Track | Links: [Paper] [artifact] [slides]
- **FSE'24** | Yifan Zhang, Jiliang (Eric) Li, Zachary Karas, Aakash Bansal, Toby Jia-Jun Li, Collin McMillan, Kevin Leach, Yu Huang: *EyeTrans: Merging Human and Machine Attention for Neural Code Summarization.* In Proceedings of the ACM International Conference on the Foundations of Software Engineering, July 15-19, 2024.
  CORE A* | **The C. F. Chen Best Paper Award** | Links: [Paper] [arXiv] [artifact] [poster] [slides]
- **WWW'24** | Chen Huang, Haoyang Li, Yifan Zhang, Wenqiang Lei, Jiancheng Lv: *Cross-Space Adaptive Filter: Integrating Graph Topology and Node Attributes for Alleviating the Over-smoothing Problem.* In Proceedings of the ACM Web Conference, May 13-17, 2024.
  CORE A* | Links: [Paper] [arXiv] [poster] [slides] [video]

- **TMI'23** | Haoyu Dong, Yifan Zhang, Hanxue Gu, Nicholas Konz, Maciej Mazurowski: *SWSSL: Sliding-Window based Self-Supervised Learning Framework for Anomaly Detection.* The IEEE Transactions on Medical Imaging, 2023.
  JCR Q1 | Links: [Paper] [artifact]

- **IJCG'17** | Jing Li, Xiangfang Li, Keliu Wu, Dong Feng, Tao Zhang, Yifan Zhang: *Thickness and Stability of Water Film Confined inside Nanoslits and Nanocapillaries of Shale and Clay.* The International Journal of Coal Geology, 2017.
  JCR Q1 | Links: [Paper]

Refereed Short Conference Papers *4 papers*

Short-paper tracks at FSE, ICPC, and ASE *2023 – 2025*

- **FSE'25** | Yifan Zhang, Chen Huang, Zachary Karas, Thuy Dung Nguyen, Kevin Leach, Yu Huang: *Enhancing Code LLM Training with Programmer Attention.* In Companion Proceedings of the 33rd ACM International Conference on the Foundations of Software Engineering. June 23-27, 2025. IVR Track Paper.
  CORE A* | Links: [Paper] [arXiv] [slides]

- **ICPC'25** | Manish Acharya*, Yifan Zhang*†, Kevin Leach, Yu Huang: *Optimizing Code Runtime Performance through Context-Aware Retrieval-Augmented Generation.* In Proceedings of the 33rd IEEE/ACM International Conference on Program Comprehension. April 27-28, 2025. ERA Track Paper.
  CORE A | Links: [Paper] [arXiv] [artifact] [slides]

- **ICPC'24** | Jiliang (Eric) Li, Yifan Zhang†, Zachary Karas, Collin McMillan, Kevin Leach, Yu Huang: *Do Machines and Humans Focus on Similar Code? Exploring Explainability of Large Language Models in Code Summarization.* In Proceedings of the 32nd IEEE/ACM International Conference on Program Comprehension, April 15-16, 2024. RENE Track Paper.
  CORE A | Links: [Paper] [arXiv] [slides]

- **ASE'23** | Aakash Bansal, Chia-Yi Su, Zachary Karas, Yifan Zhang, Yu Huang, Toby Jia-Jun Li, Collin McMillan: *Modeling Programmer Attention as Scanpath Prediction.* In Proceedings of the 38th IEEE/ACM International Conference on Automated Software Engineering, September 11-15, 2023. NIER Track Paper.
  CORE A* | Links: [Paper] [arXiv]

Refereed Extended Abstract & Position Papers *3 papers*

Position and doctoral symposium papers at FSE and ASE *2022 – 2025*

- **FSE'25** | Yifan Zhang, Kevin Leach: *Training Large Language Models to Comprehend LLVM IR via Feedback-Driven Optimization.* In Companion Proceedings of the 33rd ACM International Conference on the Foundations of Software Engineering. June 23-27, 2025. HumanAISE'25 Position Paper.
  CORE A* | Links: [Paper] [slides]

- **FSE'25** | Yifan Zhang, Kevin Leach: *Leveraging Human Insights for Enhanced LLM-based Code Repair.* In Companion Proceedings of the 33rd ACM International Conference on the Foundations of Software Engineering. June 23-27, 2025. HumanAISE'25 Position Paper.
  CORE A* | Links: [Paper] [slides]

- **ASE'22** | Yifan Zhang: *Leveraging Artificial Intelligence on Binary Code Comprehension.* In Proceedings of the 37th IEEE/ACM International Conference on Automated Software Engineering, October 10-14, 2022. Doctoral Symposium Paper.
  CORE A* | Links: [Paper] [arXiv] [poster] [slides]

Refereed Workshop & Symposium Papers *2 papers*

IAEval@NeurIPS and HotSoS *2025 – 2026*

- **IAEval'26** | Yifan Zhang*, Xiaohan Wang*, Yueke Zhang, Yu Huang, Kevin Leach: *Constraint-Guided Multi-Agent Decompilation for Executable Binary Recovery.* To be presented at the Workshop on Evaluation of Interactive Agents (IAEval@NeurIPS'26).
  Links: [arXiv]

- **HotSoS'25** | Yifan Zhang*, Michael Sandborn*, Stefan Larson, Yu Huang, Kevin Leach: *K-ASTRO: Structure-Aware Adaptation of LLMs for Code Vulnerability Detection.* Presented at the Hot Topics in the Science of Security Symposium 2025. April 2-4, 2025. WIP Paper.
  Links: [webpage] [WIP] [slides]

Preprint & Under Review Papers *6 papers*

Under review, available on arXiv *2025 – 2026*

- Yifan Zhang*, Yutong Dai*, Juntao Tan, Luyu Yang, Rishi Mullur, Thai Hoang, Zhiyuan Hu, James Zhu, Phil Mui, Silvio Savarese, Ran Xu, Zeyuan Chen: *DarwinX: Evolving Agent Harnesses Through Natural Selection.* **Featured in VentureBeat; #1 Hugging Face Paper of the Day** | Links: [arXiv] [news]
- Jixuan Chen, Jiaxin Zhang, Qinyuan Ye, Yada Pruksachatkun, Haoxiang Zhang, Jingming Zhuo, Yifan Zhang, Yutong Dai, Juntao Tan, Xiangyu Peng, Silvio Savarese, Zeyuan Chen, Lianhui Qin, Chien-Sheng Wu: *CoTrace: Data Recipes for Training Terminal Agents with Harness-Model Co-Evolution.*
  Links: [arXiv]
- Kai Mei, Zhiyuan Hu, Yutong Dai, Juntao Tan, Yifan Zhang, Dingjie Song, Dimitris N. Metaxas, Silvio Savarese, Ran Xu, Zeyuan Chen: *Opera: A Verbal Critic Framework for Long-horizon Coding Agents.*
  Links: [arXiv]
- Zihan Fang, Yifan Zhang†, Yueke Zhang, Kevin Leach, Yu Huang: *DPO-F+: Aligning Code Repair Feedback with Developer Preferences.*
  Links: [arXiv]
- Yueke Zhang, Yifan Zhang†, Kevin Leach, Yu Huang: *CodeGrad: Integrating Multi-Step Verification with Gradient-Based LLM Refinement.*
  Links: [arXiv]
- Yueke Zhang, Yifan Zhang, Zihan Fang, Kevin Leach, Wei Zhang, Yu Huang: *SCOPE: Leveraging Subgoal Critiques for Code Generation.*
  Links: [arXiv]

Industry Experience

Salesforce AI Research *Jan. 2026 – Present*

Senior Research Scientist, Supervisors: Dr. Zeyuan Chen & Dr. Ran Xu *Palo Alto, CA, USA*

- **Computer-Use Agents**: Developed *CLIFT*, training a Gemma 4 31B web agent with LoRA using conformally certified, URL-aware rewards and reusing the verifier for judge-free test-time selection. Raised WebArena Infinity success from 61.8% to 74.6% and transferred verification to VisualWebArena and Online Mind2Web, improving the latter by 11.3 points on GPT-5.5 (NeurIPS'26).
- **Coding Agents and Agent Harnesses**: Developed *DarwinX*, evolving agent harnesses around a frozen model through regression-aware selection, independent confirmation, and cross-lineage recombination. Improved official Terminal-Bench 2.1 and DeepSWE success by 7.2 and 9.4 percentage points, respectively, and achieved 91.2% audit-clean pass@1 on WebArena Infinity. Also contributed to verbal critiques for coding agents (*Opera*) and training data for harness–model co-evolution (*CoTrace*).

Amazon Web Services *Sep. 2025 – Dec. 2025*

Applied Scientist Intern, Supervisors: Dr. Brandon Paulsen & Dr. Joey Dodds *Arlington, VA, USA*

- **Stateful Cloud API Simulation**: Developed *CloudTwin*, an LLM-backed AWS API simulator integrated with LocalStack for Terraform testing without live cloud provisioning. Designed *OASIS*, a bounded operation-as-state memory for consistent replay of resource IDs and dependencies across long API traces. Achieved 91.4% project-level `terraform apply` success across 151 projects, versus 36.4% for LocalStack.

IBM T.J. Watson Research *May 2025 – Aug. 2025*

AI Research Intern, Supervisors: Dr. Dhaval Patel & Dr. Achille Fokoue *Yorktown Heights, NY, USA*

- **Grounded and Reflective Agent Planning**: Developed *SPIRAL*, an LLM planning framework that embeds Planner, Simulator, and Critic roles in Monte Carlo Tree Search. Used predicted action outcomes and dense reflective feedback to guide multi-step tool-use planning without model fine-tuning. Achieved 83.6% cascaded accuracy on TaskBench DailyLifeAPIs with Llama 4 Maverick, versus 67.4% for the LATS baseline (AAAI'26).

Google Research *Jan. 2025 – May 2025*

Student Researcher, Supervisors: Dr. Alekh Agarwal & Dr. Mircea Trofin *Sunnyvale, CA, USA*

- **LLVM IR Benchmark Automation**: Designed an automated pipeline with Google CORE/Google Research, using LLMs to generate C++ code that simulates LLVM IR profiles, enabling benchmarking and performance analysis of production code. The system supported AI models in learning compiler behavior and optimization strategies across diverse hardware platforms, with safeguards in place to preserve IP integrity.

Intel Corporation *Jun. 2024 – Dec. 2024*

GenAI Research Intern, Supervisors: Dr. Jun Li & Dr. Xue (Sharon) Yang *Santa Clara, CA, USA*

- **Construction Schedule Automation**: Developed *CONSTRUCTA* with Intel's Incubation and Disruptive Innovation group to automate scheduling for semiconductor fabrication facilities, combining construction-specific retrieval, dependency-aware context sampling, and DPO-based alignment with expert preferences. Evaluated on 4,340 construction activities, the system improved automated-planning accuracy by 28.9 percentage points over a GPT-4o basic-context baseline (NAACL'25 Industry).

TikTok/ByteDance *May 2023 – Aug. 2023*

Research Scientist Intern, Supervisors: Dr. Zhibing Zhao & Dr. Tieying Zhang *San Jose, CA, USA*

- **SQL Hint Recommendation**: Enhanced ByteBrain's SQL pipeline by developing a representation-learning-based system for recommending optimizer hints and reranking execution plans across diverse table spaces and schemas. Designed a two-Transformer workflow and pretrained a generalizable plan reranker, addressing the query-optimization stage downstream of SQL generation. Gave a technical talk at TikTok InfraLab and contributed to the team's LLM literature reviews and conference-paper reviews.

JD.com *Dec. 2018 – Mar. 2021*

Machine Learning Engineer II, Supervisors: Dr. Hu Wang & Dr. Chen Huang *Beijing, CN*

- **Action Model**: Built Bi-GRU and DeepFM models on user behavior features to predict the credit use rate and overall profit of every user in cash loan and consumer debt. The model recommends credit-limit increases to maximize income and achieved a 21.4% increase in overall profit.
- **Credit Score Propagation**: Built a heterogeneous graph over different types of user connections and applied GNN models to propagate credit scores and improve risk prediction, raising the overall accuracy of the XGB model by 5% in user classification.
- **Privacy-Preserving Collaboration**: Invented a GAN-style model using differential privacy to improve the efficiency and security of federated learning. Granted 5 CN patents based on this work; one was awarded 1st Runner-up in the 3rd JD Discovery Cup Patent Competition.

Academic Experience

Duke University *Jul. 2021 – Jun. 2022*

Research Associate, Advisor: Prof. Maciej A. Mazurowski *Durham, NC, USA*

- **Medical Image Analysis**: Developed tumor-insertion augmentation and a lightweight ResNet–ViT detector for breast tomosynthesis, improving lesion-detection AP50 by 13.1%. Co-developed *SWSSL*, a sliding-window self-supervised method for high-resolution anomaly detection that improved image-level AUC on breast tomosynthesis and chest X-rays (TMI'23).

University of Hong Kong *Mar. 2021 – Sep. 2021*

Senior Research Assistant, Advisor: Prof. Qingchen Wang *Hong Kong SAR*

- **Sequential Decision Systems**: Developed an intelligent debt-collection system that uses Transformer-based representations of customer interaction histories and offline policy-gradient learning to recommend long-horizon collection strategies. Modeled evolving account states to support context-aware decisions across multi-step workflows.

Academic Services

Area Chair & Organizing Committee Member

- 2027 International World Wide Web Conference | WWW'27 | Area Chair
- 2025 Workshop on Human-Centered AI for Software Engineering | HumanAISE@FSE'25 | Publicity Chair
- 2025 Workshop on Human-Centered AI for Software Engineering | HumanAISE@FSE'25 | Web Chair

Program Committee Member & Conference Reviewer

- 2027 IEEE/ACM International Conference on Software Engineering | ICSE'27 | Demonstrations Track
- 2027 & 2026 International Workshop on Large Language Models for Code | LLM4Code@ICSE'27/'26
- 2026 Conference on Empirical Methods in Natural Language Processing | EMNLP'26 | Industry Track
- 2026 Annual Meeting of the Association for Computational Linguistics Rolling Review | ACL ARR'26
- 2026 & 2024 Annual Conference on Neural Information Processing Systems | NeurIPS'26/'24 | Links: [webpage]
- 2026 & 2024 & 2023 AAAI Conference on Artificial Intelligence | AAAI'26/'24/'23

- 2025 Annual Meeting of the Association for Computational Linguistics | ACL'25 | Industry Track
- 2025 International Conference on Machine Learning | ICML'25
- 2025 International Conference on Learning Representations | ICLR'25
- 2025 & 2024 Annual Conference on Neural Information Processing Systems | NeurIPS'25/'24 | Datasets & Benchmarks Track
- 2024 ACM Conference on Computer and Communications Security | CCS'24 | Artifact Evaluation Track | Links: [webpage]
- 2024 Data-Centric Machine Learning Research Workshop at ICLR | DMLR@ICLR'24 | **Exceptional Reviewer Award** (Top 9.43% venue-wide) | Links: [webpage]
- 2024 & 2023 International Conference on Medical Image Computing and Computer Assisted Intervention | MICCAI'24/'23
- 2023 Annual Conference on Machine Learning and Systems | MLSys'23 | Artifact Evaluation Track
- 2023 International Conference on Mining Software Repositories | MSR'23
- 2023 AAAI Workshop on DL-Hardware Co-Design for AI Acceleration | DCAA@AAAI'23
- 2022 MICCAI Workshop on Cancer Prevention through Early Detection | CaPTion@MICCAI'22

External Conference Reviewer

- 2025 The annual IEEE International Conference on Data Engineering | ICDE'25 | Industry Track
- 2025 International Conference on Extending Database Technology | EDBT'25 | Industry Track
- 2025 Network and Distributed System Security | NDSS'25
- 2024 International Conference on Very Large Data Bases | VLDB'24 | Scalable Data Science Track
- 2023 USENIX Security Symposium | USENIX Security'23
- 2022 IEEE/CVF Computer Vision and Pattern Recognition Conference | CVPR'22

Research and Instructional Mentorship

Research Mentorship Experience

- Zichen (Ryan) Zhu, Bachelor of Applied Math & CS at Vanderbilt University. From 2025 Spring to Present. Research Topic: Dynamic Sparsity in LLM Compiler Optimization.
- Suad Mohamed, Bachelor of Psychology & CS at Belmont University. *Former Software Engineer Intern at Microsoft.* Co-Advised with Mr. Zachary Karas. From 2024 Spring to Present. Research Topic: Eye-Tracking for Human-Centered AI in Code Summarization. First Placement: Software Engineer at Microsoft.
- Manish Acharya, Bachelor of Math & CS at Vanderbilt University. *Vanderbilt Chancellor's Scholar.* From 2024 Spring to Present. Research Topic: Software Patch In-Context Learning.
- Jieyu Li, Bachelor of EEE at Shanghai Jiaotong University (SJTU), Master student of ECE at Vanderbilt University. From 2023 Fall to Present. Research Topic: AI Agent for Automated Program Repair. First Placement: Ph.D. in ECE at University of Waterloo. Advisor: Prof. Weiyi (Ian) Shang.
- Luka Mushkudiani, Bachelor of Math & CS at Vanderbilt University. *Former Intern at Meta's PyTorch Compiler Team.* From 2023 Fall to Present. Research Topic: Neural Type Systems for Python/JavaScript.
- Jiliang (Eric) Li, Bachelor of Math & CS at Vanderbilt University. From 2022 Summer to Present. Research Topic: Few-Shot Malware Classification. First Placement: Master of Science in Computer Science (MSCS) at Stanford University.

Teaching and Instructional Experience

- Graduate Teaching Assistant for CS3892: Projects in Computing for Sustainability at Vanderbilt University, *Undergraduate Capstone Project*, 2024 Spring. Instructor: Prof. Douglas H. Fisher.
- Graduate Teaching Assistant for CS3276/CS5276: Compiler Design at Vanderbilt University, 2023 Fall. Instructor: Prof. Kevin J. Leach.
- Community Teaching Assistant for Data Structure and Algorithm Training Camp (Part I and Part II) at Tsinghua University (MOOC), 2019 Spring. Instructor: Prof. Junhui Deng | Links: [certificate]
- Online Course Mentor for Data Analysis Nanodegree at Udacity China, 2017 Spring. Instructors: Dr. Josh Magee, Ria Cheruvu & Matt Maybeno. **Outstanding Mentor Award** (Top 5% world-wide) | Links: [certificate]

Funds and Awards

Scholarships and Fellowships

- 2025 Student Travel Support from National Science Foundation (NSF) for FSE'25/ISSTA'25 | $3000 Funding
- 2025 The C.F. Chen Best Paper Award from Vanderbilt University | Top 1 university-wide | $5000 Scholarship | Links: [certificate] [news]
- 2021 Research Fellowship from National Institutes of Health (NIH) | $36000/year
- 2018 & 2017 & 2016 & 2015 First-class Scholarship at China University of Petroleum (CUP)
- 2017 Roberto Roca Education Scholarship | Top 10 nationwide
- 2015 Chinese National Fellowship for Overseas Studies | Full tuition fees & CA$6000/4 months
- 2014 Schlumberger Engineering Scholarship | Top 8 university-wide
- 2013 Chinese National Scholarship for Outstanding Merits | Top 0.2% nationwide

Honors and Awards

- 2026 Featured Research Coverage in VentureBeat | #1 Hugging Face Paper of the Day | Links: [webpage] [news]
- 2026 Featured Research Recognition from Vanderbilt Institute Newsletter for ACL'26 Oral Presentation | Top 4% venue-wide | Links: [news]
- 2025 Travel Award from ACM SIGSOFT CAPS for FSE'25/ISSTA'25 Conference Attendance
- 2020 First Runner-up in the 3rd JD.com Discovery Cup Patent Competition | Top 2 company-wide
- 2020 Silver Medal Award of Distinguished Technical Recruiter at JD.com | Top 5% company-wide
- 2020 Bronze Medal Award of Certified Technical Instructor at JD.com
- 2019 Beijing Outstanding Graduate Award | Top 0.1% nationwide
- 2017 & 2015 Meritorious Winner of American Mathematical Contest in Modeling | Top 5% worldwide
- 2016 *Summa Cum Laude* for the Best Undergraduate Students at CUP | Top 1% university-wide
- 2015 & 2013 Third Prizes of Chinese National Petroleum Engineering Design Competition | Top 5% nationwide