





Jingxuan He 何静轩

 PhD Candidate, CS @ ETH Zurich
 jingxuan.he@inf.ethz.ch
 www.sri.inf.ethz.ch/people/jingxuan

 +41 78 671 55 68
 CAB H66, Universitätstrasse 6
8092 Zurich, Switzerland

Research Summary

My research focuses on the synergy of machine learning and programming, as well as its implication to security and reliability. Specifically, I worked on the following topics:

- Trustworthiness of large language models, in secure code generation [CCS'23, R2FM'24] and hallucination mitigation [ICLR'24].
- Machine learning models for bug detection [PLDI'21, ICML'22], program repair [ICML'21], and reverse engineering of binary code [CCS'18].
- Integration of machine learning into symbolic program analysis [CCS'19, PLDI'20, CCS'21].

Education

Ph.D. in Computer Science , ETH Zurich, Switzerland	2018 - 2024 (Expected)
Advisor: Martin Vechev	
M.S. in Computer Science , ETH Zurich, Switzerland	2016 - 2018
B.E. in Computer Science and Technology , Zhejiang University, China	2012 - 2016

Research Papers

- [R2FM'24] *Instruction Tuning for Secure Code Generation*
[pdf](#) **Jingxuan He***, Mark Vero*, Gabriela Krasnopolska, and Martin Vechev
Workshop on Reliable and Responsible Foundation Models, ICLR 2024
In Submission to ICML, 2024
- [ICLR'24] *Self-contradictory Hallucinations of LLMs: Evaluation, Detection and Mitigation*
[pdf](#), [web](#), [code](#) Niels Mündler, **Jingxuan He**, Slobodan Jenko, and Martin Vechev
International Conference on Learning Representations, 2024
- [CCS'23] *Large Language Models for Code: Security Hardening and Adversarial Testing*
[pdf](#), [code](#) **Jingxuan He** and Martin Vechev
ACM SIGSAC Conference on Computer and Communications Security, 2023
ICML Workshop on Challenges in Deploying Generative AI, 2023
Awarded: OpenAI Cybersecurity Grant, ACM CCS 2023 Distinguished Paper
- [ICML'22] *On Distribution Shift in Learning-based Bug Detectors*
[pdf](#), [code](#) **Jingxuan He**, Luca Beurer-Kellner, and Martin Vechev
International Conference on Machine Learning, 2022

- [ICML'21] *TFix: Learning to Fix Coding Errors with a Text-to-Text Transformer*
pdf, code Berkay Berabi, **Jingxuan He**, Veselin Raychev, and Martin Vechev
International Conference on Machine Learning, 2021
- [PLDI'21] *Learning to Find Naming Issues with Big Code and Small Supervision*
pdf **Jingxuan He**, Cheng-Chun Lee, Veselin Raychev, and Martin Vechev
ACM SIGPLAN Conference on Programming Language Design and Implementation, 2021
- [CCS'21] *Learning to Explore Paths for Symbolic Execution*
pdf, code **Jingxuan He**, Gishor Sivanrupan, Petar Tsankov, and Martin Vechev
ACM SIGSAC Conference on Computer and Communications Security, 2021
- [PLDI'20] *Learning Fast and Precise Numerical Analysis*
pdf, code **Jingxuan He**, Gagandeep Singh, Markus Püschel, and Martin Vechev
ACM SIGPLAN Conference on Programming Language Design and Implementation, 2020
- [CCS'19] *Learning to Fuzz from Symbolic Execution with Application to Smart Contracts*
pdf, code **Jingxuan He**, Mislav Balunović, Nodar Ambroladze, Petar Tsankov, and Martin Vechev
ACM SIGSAC Conference on Computer and Communications Security, 2019
- [CCS'18] *DeBin: Predicting Debug Information in Stripped Binaries*
pdf, web, code **Jingxuan He**, Pesho Ivanov, Petar Tsankov, Veselin Raychev, and Martin Vechev
ACM SIGSAC Conference on Computer and Communications Security, 2018

Research Impact

Adoption by Industry Partners

[ICML'21]: developed at [Snyk](#) as a product for suggesting code fixes

[CCS'19]: used by [ChainSecurity](#) for security audits of smart contracts: [Polkadot](#), [Ren](#), [Paxos](#), [POA](#)

Push-button Tools

[ICLR'24]: [chatprotect.ai](#), for detecting and mitigating LLM hallucinations

[CCS'18]: [debin.ai](#), for reverse engineering binaries, hundreds of active users per month

Fixed Bugs for Important Software Projects

[CCS'23]: one false negative and one false positive for GitHub CodeQL

[ICML'22]: [CPython](#), [TensorFlow](#), [Pillow](#), [PyParsing](#), [CuPy](#), [digitalbuildings](#), [Pyro](#), [ERPNext](#), etc.

[CCS'21]: 3 bugs for GNU make, 2 bugs for findutils, 4 bugs for binutils, and 2 bugs for coreutils

Popular Open-source Repositories

[CCS'18]: [debin](#), 398 stars

[CCS'19]: [ilf](#), 141 stars

[ICML'21]: [TFix](#), 63 stars

Honors and Awards

ACM CCS 2023 Distinguished Paper	2023
OpenAI Cybersecurity Grant	2023
NeurIPS 2023 Top Reviewer	2023
Birkigt Scholarship, ETH Zurich	2018
Undergraduate Research Fellowship, The Hong Kong Polytechnic University	Fall 2015

Undergraduate Research Fellowship, The University of Hong Kong
Scholarships for Outstanding Merits, Zhejiang University

Summer 2015
2013 - 2016

Invited Talks

Large Language Models for Code: Security Hardening and Adversarial Testing

Deep Learning-aided Verification Workshop @ CAV 2023	July 2023
PLSE Seminar @ National University of Singapore	June 2023
Peking University	June 2023
Zhejiang University	June 2023
LLMs for Code Seminar	May 2023
Privacy and Security in ML Seminar	April 2023
Dagstuhl Seminar on Programming Language Processing	February 2023

Machine Learning for Program Analysis

AISEC Team @ Huawei Research Munich	May 2022
BINSEC Team @ University of Paris-Saclay	March 2022
Symposium on High Confidence Software @ Peking University	December 2021
Democratizing Software Verification Workshop @ CAV 2020	July 2020

Learning to Explore Paths for Symbolic Execution

KLEE Workshop 2022	September 2022
--------------------	----------------

Learning to Detect and Fix Issues in Code

Facebook	October 2021
----------	--------------

Teaching

Program Analysis for System Security and Reliability, ETH CS Master's Course, Spring Semesters

Giving guest lectures	2020 - 2022 (3 times)
Organizing the course project	2020 - 2022 (3 times)
Teaching exercises, designing homeworks and exam questions	2020 - 2022 (3 times)

Reliable and Interpretable Artificial Intelligence, ETH CS Master's Course, Fall Semesters

Organizing for the course project	2019 - 2022 (4 times)
-----------------------------------	-----------------------

Rigorous Software Engineering, ETH CS Bachelor's Course, Spring Semesters

Giving guest lectures	2021 - 2023 (3 times)
Teaching exercises, designing homeworks and exam questions	2019 and 2023 (2 times)

ETH CS Seminar Courses: ML for Code, Blockchain Security, and Software Engineering

Co-organizing the entire course and co-examining students	2021 - 2023 (4 times)
Advising student presentations	2019 - 2023 (5 times)

Mentoring

I guided students in project definition, problem solving, paper or thesis writing, and publication cycle.

Finished

Gabriela Krasnopolkska: Master's Thesis [*R2FM'24*] → Machine Learning Engineer at Norvatis
Niels Mündler: Master's Thesis [*ICLR'24*] → Co-founder at OpenSwap Tech AG
Luca Beurer-Kellner: PhD Research [*ICML'22*] → Ongoing PhD Student at ETH
Berkay Berabi: Master's Thesis [*ICML'21*] → Software Engineer at Snyk
Gishor Sivanrupan: Semester Project [*CCS'21*] → Software Engineer at Snyk
Jiacheng Shen: Master's Thesis → Security Engineer at Tencent Keen Lab
Aurélia Autem: Master's Thesis → Security Engineer at Pictet Group
Axel Pohl: Bachelor's Thesis → ETH Master's Program

Ongoing

PhD Research: Mark Vero
Master's Thesis: Daniel Frey, Omkar Zade, Shuang Luo
Semester Project: Slobodan Jenko, Ivan Milev, Samuel Simko

Service

Program Committee

PLDI 2022 Artifact Evaluation
Machine Learning for Program Analysis Workshop 2020

Reviewing

NeurIPS 2023 (Top Reviewer)
IEEE Transactions on Software Engineering 2023
AISTATS 2023
Neural Conversational AI Workshop @ ICML 2023
Challenges of Deploying Generative AI Workshop @ ICML 2023
ICML 2022
ACM Transactions on Software Engineering and Methodology 2022
IEEE Transactions on Computers 2022