

Junfeng Yang

Department of Computer Science
Columbia University
519 CSB, Mail Code 0401
New York, NY 10027

Phone: (212) 939-7012
Fax: (212) 666-0140
junfeng@cs.columbia.edu
<https://www.cs.columbia.edu/~junfeng/>

Research Interests

Trustworthy Software and AI Systems

Education

2008	PhD in Computer Science <i>Dissertation:</i> Automatically Finding Serious Storage System Errors <i>Advisor:</i> Dawson Engler	STANFORD UNIVERSITY
2002	MS in Computer Science	STANFORD UNIVERSITY
2000	BS in Computer Science <i>Dissertation:</i> Computer-aided Rapid Prototyping of Mechanical Parts Admitted without entrance examination based on an outstanding high-school record. One of only 60 students selected for a two-year special academic program (1995–1997).	TSINGHUA UNIVERSITY

Professional Experience

2021–now	Professor	COLUMBIA UNIVERSITY
2013–2021	Associate Professor	COLUMBIA UNIVERSITY
2014–2018	Co-founder	NIMBLEDROID, INC. (Acquired)
2008–2013	Assistant Professor	COLUMBIA UNIVERSITY
2011–now	Co-director, Software Systems Lab	COLUMBIA UNIVERSITY
2008–2011	Director, Reliable Computer Systems Lab	COLUMBIA UNIVERSITY
2008–2012	Technical Consultant	MICROSOFT RESEARCH, SILICON VALLEY
2007–2008	Postdoc Researcher	MICROSOFT RESEARCH, SILICON VALLEY
2001–2007	Research Assistant	STANFORD UNIVERSITY
06/2005–09/2005	Research Intern	MICROSOFT RESEARCH, SILICON VALLEY

Select Honors and Awards

2026	Asian American Scholar Forum (AASF) Fellow
2025	ACM Fellow, for leadership and contributions to trustworthy software and AI systems
2025	Distinguished Speaker, ACM SIGOPS Strategic Workshop
2025	ACM SIGOPS Mark Weiser Award
2025	IEEE S&P Test-of-Time Award [68]
2019	Inaugural Rock Star Award, Association of Chinese Scholars in Computing
2012	Sloan Research Fellowship
2012	AFOSR Young Investigator Program Award

2011 NSF CAREER Award

Major Publication Awards and Distinctions

2026 IEEE S&P Distinguished Paper Award [8]
2025 CCS Distinguished Artifact Award [11]
2025 NAACL Oral Presentation [15]
2024 ICML Spotlight (top 3.5%) [23]
2024 CIDR Best Paper Award [29]
2022 OSDI Best Paper Award [39]
2021 USENIX ATC Best Paper Award [45]
2020 ECCV Oral Presentation (top 2%) [48]
2019 CACM Research Highlight [112]
2017 SOSP Best Paper Award [62]
2004 OSDI Best Paper Award [84]

Research Awards

2026 Columbia Center for AI and Responsible Financial Innovation (CAIRFI) Research Award (with Capital One)
2025 Google Cyber NYC Faculty Research Award (three awards)
2025 Columbia Center for Digital Finance and Technologies (CDFT) Research Award
2024 Finalist, Amazon Nova AI Challenge: Trusted AI
2024 Columbia Center of Artificial Intelligence Technology (CAIT) Award (with Amazon)
2024 Columbia Center for Digital Finance and Technologies (CDFT) Research Award
2024 Google Cyber NYC Faculty Research Award (three awards)
2024 SEAS-KFAI Generative AI and Public Discourse Research Award
2023 Google Cyber NYC Faculty Research Award (three awards)
2023 Columbia Center for Digital Finance and Technologies (CDFT) Research Award
2022 Columbia Center of Artificial Intelligence Technology (CAIT) Award (with Amazon)
2020 Facebook Faculty Research Award
2019 J.P. Morgan Faculty Research Award
2014 Google Faculty Research Award

Research Grants and Contracts

10/2026–09/2030 PI, NSF CNS-2526620, \$1,220,000, Collaborative Research: SaTC 2.0: RES: AIGIS: Securing the Deep Learning Model Supply Chain, with co-PI: James Davis, Vasileios P. Kemerlis
07/2026–06/2027 sole PI, Columbia Center for AI and Responsible Financial Innovation (CAIRFI), with Capital One, \$110,000, Proactive Defenses for Large Language Models
08/2025–07/2026 sole PI, Samsung, \$230,769, Robust Neural Network for Analog Circuits
06/2025–12/2029 co-PI, DARPA ScAN, \$5,402,118, Large-Scale Analog Neural Network Processor and Machine-Learning Models, with Mingoo Seok (PI), Chris H. Kim

07/2025–06/2026 sole PI, Google Cyber NYC Research and Workforce Initiative, \$80,000, In-model Firewall: Flexible, Generalized Defense Against LLM Jailbreaks

07/2025–06/2026 co-PI, Google Cyber NYC Research and Workforce Initiative, \$80,000, KAgent: A Neuro-Symbolic Agent for Safely Resolving Kernel Vulnerabilities, with Baishakhi Ray (PI)

07/2025–06/2026 co-PI, Google Cyber NYC Research and Workforce Initiative, \$80,000, Fuzzing with Tool-assisted LLMs, with Suman Jana (PI)

07/2025–06/2026 sole PI, Google Cyber NYC Research and Workforce Initiative, \$70,000, Google Cloud Platform Credits for Research

07/2025–06/2026 PI, Columbia Center for Digital Finance and Technologies (CDFT), \$25,000, zkFuzz: Practical Fuzzing for Zero-Knowledge Circuits, with co-PI: Suman Jana

09/2025–08/2026 sole PI, Columbia Research Stability Funding, \$100,000, Neuro-symbolic Static Analysis for Effective Vulnerability Detection

11/2024–10/2025 sole PI, Amazon Nova AI Challenge: Trusted AI, \$250,000, SecCoder: A Secure AI Coding System

07/2024–06/2025 sole PI, Columbia Center for Digital Finance and Technologies (CDFT), \$25,000, SMARTTEST: a Multi-Modal LLM system for Effective Smart Contract Fuzzing

07/2024–06/2025 sole PI, Google Cyber NYC Research and Workforce Initiative, \$80,000, Securing the Software Supply Chain via Effective Protection of ML Model Deserialization

07/2024–06/2025 co-PI, Google Cyber NYC Research and Workforce Initiative, \$80,000, Towards Bug-oriented Fuzzing Exploration, with Suman Jana (PI)

07/2024–06/2025 co-PI, Google Cyber NYC Research and Workforce Initiative, \$80,000, Multi-modal, Semantic-aware AI Agents for Automated Vulnerability Detection and Repair, with Baishakhi Ray (PI)

07/2024–06/2025 sole PI, Google Cyber NYC Research and Workforce Initiative, \$59,000, Google Cloud Platform Credits for Research

09/2024–08/2025 sole PI, Columbia Center of AI Technology (with Amazon), \$135,000, EditGuardian: A Universal Multi-Agent LLM Framework for Semantic-Preserving Code Editing

01/2024–12/2024 PI, SEAS-KFAI Generative AI and Public Discourse Research Program, \$80,000, Detecting AI-Generated Content via Rewriting, with co-PI: Carl Vondrick

07/2023–06/2024 sole PI, Columbia Center for Digital Finance and Technologies (CDFT), \$40,000, Automating Smart Contract Verification and Monitoring with Learned Invariants

07/2023–06/2024 PI, Google Cyber NYC Research and Workforce Initiative, \$80,000, Securing the Software Supply Chain via Effective Detection of Rogue Updates, with co-PI: Yaniv David

07/2023–06/2024 co-PI, Google Cyber NYC Research and Workforce Initiative, \$80,000, Towards Statistically Rigorous Fuzzers, with Suman Jana (PI)

07/2023–06/2024 co-PI, Google Cyber NYC Research and Workforce Initiative, \$80,000, Multi-modal, Semantic-aware LLMs for Automated Bug Repair, with Baishakhi Ray (PI)

07/2023–06/2024 sole PI, Google Cyber NYC Research and Workforce Initiative, \$53,655, Google Cloud Platform Credits for Research

09/2022–08/2024 PI, Columbia SEAS/EVPR Stimulus Grants for Large Center Proposal Development, \$400,000, OASIS: Software Supply-chain Security, with co-PI: Rebecca Wright, Ronghui Gu, Suman Jana, Baishakhi Ray

03/2022–09/2022 PI, DARPA/GE, \$45,000, Reverse-Engineering Deceptions, with co-PI: Carl Vondrick

01/2022–06/2023	PI, Columbia Center of AI Technology (with Amazon), \$150,000, <u>A Tale of Two Models: Constructing and Defending Against Adversarial Deviation Attacks on Edge Models</u> , with co-PI: Asaf Cidon
02/2020–01/2021	co-PI, Facebook Faculty Research Award, \$50,000, <u>Concerto: Generating Embeddings with Systems-level Objectives</u> , with Asaf Cidon (PI), Suman Jana
07/2019–06/2020	sole PI, Google Cloud Platform Credits For Research, \$10,000, <u>Learning Versatile Models Improves Adversarial Robustness</u>
03/2019–02/2020	PI, J.P. Morgan Faculty Research Award, \$150,000, <u>Efficient Formal Safety Analysis of Neural Networks</u> , with co-PI: Jeannette M. Wing, Suman Jana
02/2019–01/2020	sole PI, Amazon Web Services Cloud Credits For Research, \$20,000, <u>Metric Learning for Adversarial Robustness</u>
05/2018–04/2019	sole PI, Accenture Research Gift, \$30,000, <u>Security and Robustness of Machine Learning</u>
09/2017–08/2020	co-PI, ONR N00014-17-1-2788, \$3,243,244, <u>ABIDES: Adaptive BINARY Debloating and Security</u> , with Georgios Portokalidis (PI), Vasileios P. Kemerlis
02/2017–01/2018	sole PI, Canon, \$75,000, <u>Reliability and Security of Cyber-Physical Systems</u>
09/2017–08/2018	sole PI, Amazon Web Services Cloud Credits For Research, \$25,000, <u>Grandet: A Unified, Economical Object Store for Web Applications</u>
09/2016–08/2022	PI, NSF CNS-1564055, \$1,200,000, <u>Efficient Repair of Learning Systems via Machine Unlearning</u> , with co-PI: Yinzhi Cao
03/2016–02/2019	PI, ONR N00014-16-1-2263, \$1,114,488, <u>Adapting Static and Dynamic Program Analysis to Effectively Harden Debloated Software</u> , with co-PI: Georgios Portokalidis
02/2016–01/2017	sole PI, Canon, \$75,000, <u>Reliability and Security of Cyber-Physical Systems</u>
08/2015–07/2018	co-PI, ONR, \$1,415,127, <u>The YOLO Approach to Resilient Cyber-Physical Systems</u> , with Simha Sethumadhavan (PI)
01/2015–12/2015	sole PI, Amazon Web Services Cloud Credits For Education, \$25,000, <u>Leveraging Cloud for Systems Education</u>
02/2014–01/2015	sole PI, Google, \$77,807, <u>Efficiently, Effectively Detecting Mobile App Bugs with AppDoctor</u>
06/2013–09/2013	sole PI, NSF CNS-1340511, \$16,000, <u>Research Experiences for Undergraduates (REU) for LOOM: a Language and System for Bypassing and Diagnosing Concurrency Errors</u>
06/2013–09/2013	sole PI, NSF CNS-1340506, \$16,000, <u>Research Experiences for Undergraduates (REU) for Guanyin: a Thousand hands with a Thousand eyes for Distributed Software Checking</u>
09/2012–08/2016	co-PI, NSF CCF-1162021, \$800,000, <u>SHF: Medium: RacePro: Automatically Detecting API Races in Deployed Systems</u> , with Jason Nieh (PI)
2012–2016	sole PI, Sloan Foundation, \$50,000, <u>Sloan Research Fellowship</u>
07/2012–06/2016	sole PI, AFOSR YIP, \$359,998, <u>Concurrency Attacks and Defenses</u>
07/2012–06/2016	PI, ONR N00014-12-1-0166, \$749,975, <u>Transparently Extending Programs at Compilation to Prevent Bugs</u> , with co-PI: Angelos Keromytis
09/2011–01/2016	co-PI, DARPA MRC, \$6,619,270, <u>MEERKATS: Maintaining Enterprise Resiliency via Kaleidoscopic Adaptation and Transformation of Software Services</u> , with Angelos Keromytis (PI), Roxana Geambasu, Simha Sethumadhavan, Sal Stolfo
09/2011–08/2014	sole PI, NSF CNS-1117805, \$250,000, <u>LOOM: a Language and System for Bypassing and Diagnosing Concurrency Errors</u>
02/2011–01/2017	sole PI, NSF CAREER CNS-1054906, \$450,000, <u>CAREER: Making Threads More Deterministic by Memoizing Schedules</u>

- 10/2010–08/2015 co-PI, DARPA CRASH, \$6,424,180, SPARCHS: Symbiotic, Polymorphic, Autonomic, Resilient, Clean-slate, Host Security, with Simha Sethumadhavan (PI), Sal Stolfo, Angelos Keromytis, David August
- 07/2010–06/2013 PI, NSF CNS-1012633, \$325,000, SemGrep: a System for Improving Software Reliability through Semantic Similarity Bug Search, with co-PI: Angelos Keromytis, Dawson Engler
- 08/2010–01/2015 co-PI, IARPA, \$7,530,113, MINESTRONE, with Angelos Keromytis (PI), Sal Stolfo
- 09/2009–08/2014 PI, NSF CNS-0905246, \$1,012,000, Guanyin: a Thousand hands with a Thousand eyes for Distributed Software Checking, with co-PI: Gail Kaiser, Jason Nieh

Publications

(All papers are available online at <https://www.cs.columbia.edu/~junfeng/>.)

Refereed Conference Publications

- [1] Penghui Li, Josef Sarfati Korich, Songchen Yao, Changhua Luo, Jianjia Yu, Yinzhi Cao, and Junfeng Yang. Automatically learning vulnerability patterns for scalable static analysis of web applications. In CCS, 2026.
- [2] Hideaki Takahashi, Suman Jana, and Junfeng Yang. Efficient branch-and-bound testing and verification of zkVMs. In CCS, 2026.
- [3] Weiliang Zhao, Daniel Ben-Levi, Jinjun Peng, Zhou Yu, and Junfeng Yang. Jailbreaking jailbreaks: A proactive defense for LLMs. In EMNLP, 2026.
- [4] Ira Ceka, Hailie Mitchell, Saurabh Pujar, Luca Buratti, Shyam Ramji, Junfeng Yang, Gail Kaiser, and Baishakhi Ray. Understanding automated program repair agents through the lens of traceability: An empirical study. In ISSTA, 2026.
- [5] Chenxi Huang, Alex Mathai, Feiyang Yu, Aleksandr Nogikh, Petros Maniatis, Franjo Ivančić, Eugene Wu, Kostis Kaffes, Junfeng Yang, and Baishakhi Ray. Outrunning LLM cutoffs: A live kernel crash resolution benchmark for all. In ICML, 2026.
- [6] Hideaki Takahashi, Jihwan Kim, Suman Jana, and Junfeng Yang. zkFuzz: Foundation and framework for effective fuzzing of zero-knowledge circuits. In IEEE S&P, 2026.
- [7] Penghui Li, Hong Yau Chong, Yinzhi Cao, and Junfeng Yang. Detecting privilege escalation in polyglot microservices via agentic program analysis. In IEEE S&P, 2026.
- [8] Simin Chen, Jinjun Peng, Yixin He, Junfeng Yang, and Baishakhi Ray. Your compiler is backdooring your model: Understanding and exploiting compilation inconsistency vulnerabilities in deep learning compilers. In IEEE S&P, 2026.

IEEE S&P Distinguished Paper Award

- [9] Haoda Wang, Steven Myint, Vandii Verma, Yonatan Winetraub, Junfeng Yang, and Asaf Cidon. Radshield: Software radiation protection for commodity hardware in space. In ASPLOS, 2026.
- [10] Ioannis Zarkadas, Kelly Kostopoulou, Thomas Graham, Junfeng Yang, Philip A. Bernstein, Asaf Cidon, and Tamer Eldeeb. Rosé: Flexible replication with strong semantics for partitioned databases. In CIDR, 2026.
- [11] Andreas D. Kellas, Neophytos Christou, Wenxin Jiang, Penghui Li, Laurent Simon, Yaniv David, Vasileios P. Kemerlis, James C. Davis, and Junfeng Yang. PickleBall: Secure deserialization of pickle-based machine learning models. In CCS, 2025.

CCS Distinguished Artifact Award

- [12] Wei Hao, Ran Li, Weiliang Zhao, Junfeng Yang, and Chengzhi Mao. Learning to rewrite: Generalized LLM-generated text detection. In ACL, 2025.

- [13] Weichen Li, Albert Jan, Baishakhi Ray, Junfeng Yang, Chengzhi Mao, and Kexin Pei. Ed-itlord: Learning code transformation rules for code editing. In ICML, 2025.
- [14] Wei Hao, Van Tran, Vincent Rideout, Zixi Wang, AnMei Dasbach-Prisk, M. H. Afifi, Junfeng Yang, Ethan Katz-Bassett, Grant Ho, and Asaf Cidon. Do spammers dream of electric sheep? characterizing the prevalence of LLM-generated malicious emails. In IMC, 2025.
- [15] Weiliang Zhao, Daniel Ben-Levi, Wei Hao, Junfeng Yang, and Chengzhi Mao. Diversity helps jailbreak large language models. In NAACL, 2025.

NAACL Oral Presentation

- [16] Zirui Zhang, Wei Hao, Aroon Sankoh, William Lin, Emanuel Mendiola-Ortiz, Junfeng Yang, and Chengzhi Mao. I can hear you: Selective robust training for deepfake audio detection. In ICLR, 2025.
- [17] Wei Hao, Zixi Wang, Lauren Hong, Lingxiao Li, Nader Karayanni, AnMei Dasbach-Prisk, Chengzhi Mao, Junfeng Yang, and Asaf Cidon. Nazar: Monitoring and adapting ML models on mobile devices. In ASPLOS, 2025.
- [18] Pierre Tholoniati, Kelly Kostopoulou, Mosharaf Chowdhury, Asaf Cidon, Roxana Geambasu, Mathias Lécuyer, and Junfeng Yang. DPack: Efficiency-oriented privacy budget scheduling. In EuroSys, 2025.
- [19] Alex Mathai, Chenxi Huang, Petros Maniatis, Aleksandr Nogikh, Franjo Ivančić, Junfeng Yang, and Baishakhi Ray. kGym: A platform and dataset to benchmark large language models on linux kernel crash resolution. In NeurIPS Datasets and Benchmarks Track, 2024.
- [20] Yangruibo Ding, Jinjun Peng, Marcus J. Min, Gail Kaiser, Junfeng Yang, and Baishakhi Ray. SemCoder: Training code language models with comprehensive semantics reasoning. In NeurIPS, 2024.
- [21] James Liyuan Wang, Ran Li, Junfeng Yang, and Chengzhi Mao. RAFT: Realistic attacks to fool text detectors. In EMNLP, 2024.
- [22] Tamer Eldeeb, Sebastian Burckhardt, Reuben Bond, Asaf Cidon, Junfeng Yang, and Philip A. Bernstein. Cloud actor-oriented database transactions in Orleans. VLDB, 17(12):3720–3730, 2024.
- [23] Kexin Pei, Weichen Li, Qirui Jin, Shuyang Liu, Scott Geng, Lorenzo Cavallaro, Junfeng Yang, and Suman Jana. Exploiting code symmetries for learning program semantics. In ICML, 2024.

ICML Spotlight (top 3.5%)

- [24] Wei Hao, Daniel Mendoza, Rafael da Silva, Deepak Narayanan, Amar Phanishayee, Asaf Cidon, and Junfeng Yang. MGit: A model versioning and management system. In ICML, 2024.
- [25] Chengzhi Mao, Carl Vondrick, Hao Wang, and Junfeng Yang. Raidar: generative ai detection via rewriting. In ICLR, 2024.
- [26] Haozhe Chen, Junfeng Yang, Carl Vondrick, and Chengzhi Mao. INVITE: INterpret and control vision-language models with text explanations. In ICLR, 2024.
- [27] Yun-Yun Tsai, Fu-Chen Chen, Albert Y. C. Chen, Junfeng Yang, Che-Chun Su, Min Sun, and Cheng-Hao Kuo. GDA: Generalized diffusion for robust test-time adaptation. In CVPR, 2024.
- [28] Sally Junsong Wang, Kexin Pei, and Junfeng Yang. SmartInv: Multimodal learning for smart contract invariant inference. In IEEE S&P, 2024.
- [29] Tamer Eldeeb, Philip A. Bernstein, Asaf Cidon, and Junfeng Yang. Chablis: Fast and general transactions in geo-distributed systems. In CIDR, 2024.

CIDR Best Paper Award

- [30] Raphael J. Sofaer, Yaniv David, Mingqing Kang, Jianjia Yu, Yinzhi Cao, Junfeng Yang, and Jason Nieh. RogueOne: Detecting rogue updates via differential data-flow analysis using trust domains. In ICSE, 2024.
- [31] Yaniv David, Neophytos Christou, Andreas D. Kellas, Vasileios P. Kemerlis, and Junfeng Yang. QUACK: Hindering deserialization attacks via static duck typing. In NDSS, 2024.
- [32] Yun-Yun Tsai, Chengzhi Mao, and Junfeng Yang. Convolutional visual prompt for robust visual perception. In NeurIPS, 2023.
- [33] Tamer Eldeeb, Xincheng Xie, Philip A. Bernstein, Asaf Cidon, and Junfeng Yang. Chardonnay: Fast and General Datacenter Transactions for On-Disk Databases. In OSDI, 2023.
- [34] Lingmei Weng, Yigong Hu, Peng Huang, Jason Nieh, and Junfeng Yang. Effective Performance Issue Diagnosis with Value-Assisted Cost Profiling. In EuroSys, 2023.
- [35] Chengzhi Mao, Revant Teotia, Amrutha Sundar, Sachit Menon, Junfeng Yang, Xin Wang, and Carl Vondrick. Doubly right object recognition: A why prompt for visual rationales. In CVPR, 2023.
- [36] Chengzhi Mao, Scott Geng, Junfeng Yang, Xin Wang, and Carl Vondrick. Understanding zero-shot adversarial robustness for large-scale models. In ICLR, 2023.
- [37] Chengzhi Mao, Lingyu Zhang, Abhishek Vaibhav Joshi, Junfeng Yang, Hao Wang, and Carl Vondrick. Robust Perception through Equivariance. In ICML, 2023.
- [38] Yaniv David, Xudong Sun, Raphael J. Sofaer, Aditya Senthilnathan, Junfeng Yang, Zhiqiang Zuo, Guoqing Harry Xu, Jason Nieh, and Ronghui Gu. UPGRAVISOR: Early adopting dependency updates using hybrid program analysis and hardware tracing. In OSDI, 2022.
- [39] Yuhong Zhong, Haoyu Li, Yu Jian Wu, Ioannis Zarkadas, Jeffrey Tao, Evan Mesterhazy, Michael Makris, Junfeng Yang, Amy Tai, Ryan Stutsman, and Asaf Cidon. XRP: In-Kernel Storage Functions with eBPF. In OSDI, 2022.

OSDI Best Paper Award

- [40] Kexin Pei, Dongdong She, Michael Wang, Scott Geng, Zhou Xuan, Yaniv David, Junfeng Yang, Suman Jana, and Baishakhi Ray. NeuDep: Neural Binary Memory Dependence Analysis. In ESEC/FSE, 2022.
- [41] Wei Hao, Aahil Awatramani, Jiayang Hu, Chengzhi Mao, Pin-Chun Chen, Eyal Cidon, Asaf Cidon, and Junfeng Yang. A tale of two models: Constructing evasive attacks on edge models. In MLSys, 2022.
- [42] Chengzhi Mao, Kevin Xia, James Wang, Hao Wang, Junfeng Yang, Elias Bareinboim, and Carl Vondrick. Causal transportability for visual recognition. In CVPR, 2022.
- [43] Kexin Pei, Jonas Guan, Matthew Broughton, Zhongtian Chen, Songchen Yao, David Williams-King, Vikas Ummadisetty, Junfeng Yang, Baishakhi Ray, and Suman Jana. StateFormer: Fine-grained type recovery from binaries using generative state modeling. In ESEC/FSE, 2021.
- [44] Chengzhi Mao, Mia Chiquier, Hao Wang, Junfeng Yang, and Carl Vondrick. Adversarial attacks are reversible with natural supervision. In ICCV, 2021.
- [45] Lingmei Weng, Peng Huang, Jason Nieh, and Junfeng Yang. Argus: Debugging performance issues in modern desktop applications with annotated causal tracing. In USENIX ATC, 2021.

USENIX ATC Best Paper Award

- [46] Chengzhi Mao, Augustine Cha, Amogh Gupta, Hao Wang, Junfeng Yang, and Carl Vondrick. Generative interventions for causal learning. In CVPR, 2021.
- [47] Kexin Pei, Jonas Guan, David Williams-King, Junfeng Yang, and Suman Jana. XDA: Accurate, robust disassembly with transfer learning. In NDSS, 2021.
- [48] Chengzhi Mao, Amogh Gupta, Vikram Nitin, Baishakhi Ray, Shuran Song, Junfeng Yang, and Carl Vondrick. Multitask learning strengthens adversarial robustness. In ECCV, 2020.

ECCV Oral Presentation (top 2%)

- [49] Yang Tang and Junfeng Yang. Lambdata: Optimizing serverless computing by making data intents explicit. In IEEE CLOUD, 2020.

IEEE TCSVC YESC Student Paper Award

- [50] Xinhao Yuan and Junfeng Yang. Effective concurrency testing for distributed systems. In ASPLOS, 2020.
- [51] David Williams-King, Hidenori Kobayashi, Kent Williams-King, Graham Patterson, Frank Spano, Yu Jian Wu, Junfeng Yang, and Vasileios P. Kemerlis. Egalito: Layout-agnostic binary recompilation. In ASPLOS, 2020.
- [52] Chengzhi Mao, Ziyuan Zhong, Junfeng Yang, Carl Vondrick, and Baishakhi Ray. Metric learning for adversarial robustness. In NeurIPS, 2019.
- [53] Dongdong She, Kexin Pei, Dave Epstein, Junfeng Yang, Baishakhi Ray, and Suman Jana. NEUZZ: Efficient fuzzing with neural program smoothing. In IEEE S&P, 2019.

CSAW 2019 Applied Research Finalist

- [54] Miguel A. Arroyo, M. Tarek Ibn Ziad, Hidenori Kobayashi, Junfeng Yang, and Simha Sethumadhavan. YOLO: Frequently resetting cyber-physical systems for security. In SPIE, 2019.
- [55] Shiqi Wang, Kexin Pei, Justin Whitehouse, Junfeng Yang, and Suman Jana. Efficient formal safety analysis of neural networks. In NeurIPS, 2018.
- [56] Gang Hu, Linjie Zhu, and Junfeng Yang. Appflow: Using machine learning to synthesize robust, reusable UI tests. In ESEC/FSE, 2018.
- [57] Hao Zhou, Ming Chen, Qian Lin, Yong Wang, Xiaobin She, Sifan Liu, Rui Gu, Beng Chin Ooi, and Junfeng Yang. Overload control for scaling wechat microservices. In SoCC, 2018.
- [58] Shiqi Wang, Kexin Pei, Justin Whitehouse, Junfeng Yang, and Suman Jana. Formal security analysis of neural networks using symbolic intervals. In USENIX Security, 2018.
- [59] Xinhao Yuan, Junfeng Yang, and Ronghui Gu. Partial order aware concurrency sampling. In CAV, 2018.
- [60] Shixiong Zhao, Rui Gu, Haoran Qiu, Tsz On Li, Yuexuan Wang, Heming Cui, and Junfeng Yang. OWL: Understanding and detecting concurrency attacks. In DSN, 2018.
- [61] Yinzhi Cao, Alexander Fangxiao Yu, Andrew Aday, Eric Stahl, Jon Merwine, and Junfeng Yang. Efficient repair of polluted machine learning systems via causal unlearning. In AsiaCCS, 2018.
- [62] Kexin Pei, Yinzhi Cao, Junfeng Yang, and Suman Jana. DeepXplore: Automated whitebox testing of deep learning systems. In SOSP, 2017.

CSAW 2018 Applied Research Second Place; CACM Research Highlight; SOSP Best Paper Award

- [63] David Williams-King, Graham Gobieski, Kent Williams-King, James P Blake, Xinhao Yuan, Patrick Colp, Michelle Zheng, Vasileios P Kemerlis, Junfeng Yang, and William Aiello. Shuffler: Fast and deployable continuous code re-randomization. In OSDI, 2016.
- [64] Yang Tang, Gang Hu, Xinhao Yuan, Lingmei Weng, and Junfeng Yang. Grandet: A unified, economical object store for web applications. In SoCC, 2016.
- [65] Eric Koskinen and Junfeng Yang. Reducing crash recoverability to reachability. In POPL, 2016.
- [66] Heming Cui, Rui Gu, Cheng Liu, Tianyu Chen, and Junfeng Yang. Paxos made transparent. In SOSP, 2015.
- [67] Yang Tang and Junfeng Yang. Secure deduplication of general computations. In USENIX ATC, 2015.
- [68] Yinzhi Cao and Junfeng Yang. Towards making systems forget with machine unlearning. In IEEE S&P, 2015.

IEEE S&P Test-of-Time Award; ICBS Frontiers of Science Award

- [69] Gang Hu, Xinhao Yuan, Yang Tang, and Junfeng Yang. Efficiently, effectively detecting mobile app bugs with appdoctor. In EuroSys, 2014.
- [70] Heming Cui, Jiri Simsa, Yi-Hong Lin, Hao Li, Ben Blum, Xinan Xu, Junfeng Yang, Garth A. Gibson, and Randal E. Bryant. Parrot: a practical runtime for deterministic, stable, and reliable threads. In SOSP, 2013.
- [71] Jingyue Wu, Gang Hu, Yang Tang, and Junfeng Yang. Effective dynamic detection of alias analysis errors. In ESEC/FSE, 2013.
- [72] Heming Cui, Gang Hu, Jingyue Wu, and Junfeng Yang. Verifying systems rules using rule-directed symbolic execution. In ASPLOS, 2013.
- [73] Jingyue Wu, Yang Tang, Gang Hu, Heming Cui, and Junfeng Yang. Sound and precise analysis of parallel programs through schedule specialization. In PLDI, 2012.
- [74] Heming Cui, Jingyue Wu, John Gallagher, Huayang Guo, and Junfeng Yang. Efficient deterministic multithreading through schedule relaxation. In SOSP, 2011.
- [75] Oren Laadan, Nicolas Viennot, Chia-Che Tsai, Chris Blinn, Junfeng Yang, and Jason Nieh. Pervasive detection of process races in deployed systems. In SOSP, 2011.
- [76] Huayang Guo, Ming Wu, Lidong Zhou, Gang Hu, Junfeng Yang, and Lintao Zhang. Practical software model checking via dynamic interface reduction. In SOSP, 2011.
- [77] Ding Yuan, Yinglian Xie, Rina Panigrahy, Junfeng Yang, Chad Verbowski, and Arunvijay Kumar. Context-based online configuration-error detection. In USENIX ATC, 2011.
- [78] Heming Cui, Jingyue Wu, Chia-Che Tsai, and Junfeng Yang. Stable deterministic multithreading through schedule memoization. In OSDI, 2010.
- [79] Jingyue Wu, Heming Cui, and Junfeng Yang. Bypassing races in live applications with execution filters. In OSDI, 2010.
- [80] Mark Gabel, Junfeng Yang, Yuan Yu, Moises Goldszmidt, and Zhendong Su. Scalable and systematic detection of buggy inconsistencies in source code. In OOPSLA, 2010.
- [81] Junfeng Yang, Tisheng Chen, Ming Wu, Zhilei Xu, Xuezheng Liu, Haoxiang Lin, Mao Yang, Fan Long, Lintao Zhang, and Lidong Zhou. MODIST: Transparent model checking of unmodified distributed systems. In NSDI, 2009.
- [82] Junfeng Yang, Can Sar, and Dawson Engler. Explode: a lightweight, general system for finding serious storage system errors. In OSDI, 2006.
- [83] Junfeng Yang, Can Sar, Paul Twohey, Cristian Cadar, and Dawson Engler. Automatically generating malicious disks using symbolic execution. In IEEE S&P, 2006.

- [84] Junfeng Yang, Paul Twohey, Dawson Engler, and Madanlal Musuvathi. Using model checking to find serious file system errors. In OSDI, 2004.

OSDI Best Paper Award

- [85] Ted Kremenek, Ken Ashcraft, Junfeng Yang, and Dawson Engler. Correlation exploitation in error ranking. In FSE, 2004.
- [86] Junfeng Yang, Ted Kremenek, Yichen Xie, and Dawson Engler. MECA: an extensible, expressive system and language for statically checking security properties. In CCS, 2003.
- [87] Andy Chou, Junfeng Yang, Benjamin Chelf, Seth Hallem, and Dawson Engler. An empirical study of operating systems errors. In SOSP, 2001.

Refereed Workshop Publications

- [88] Alex Mathai, Chenxi Huang, Suwei Ma, Jihwan Kim, Hailie Mitchell, Aleksandr Nogikh, Petros Maniatis, Franjo Ivančić, Junfeng Yang, and Baishakhi Ray. kAgent: An execution-guided crash resolution agent for the linux kernel. DL4C; non-archival workshop paper, 2026.
- [89] Haoda Wang, Robert Leo Pendergrast, Kristófer Fannar Björnsson, Anika Somaia, Ezra Landa, Junfeng Yang, and Asaf Cidon. Above the clouds: New software challenges in space computing. In HotNets, 2025.
- [90] Jinjun Peng, Leyi Cui, Kele Huang, Junfeng Yang, and Baishakhi Ray. CWEval: Outcome-driven evaluation on functionality and security of llm code generation. In LLM4Code, 2025.
- [91] Yun-Yun Tsai, Ruize Xu, Chengzhi Mao, and Junfeng Yang. From detection to deception: Are ai-generated image detectors adversarially robust? ReGenAI; non-archival workshop paper, 2024.
- [92] Qingyuan Liu, Pengyuan Shi, Yun-Yun Tsai, Chengzhi Mao, and Junfeng Yang. Turns out i'm not real: Towards robust detection of ai-generated videos. GCV; non-archival workshop paper, 2024.
- [93] Haoda Wang, Steven Myint, Vandí Verma, Yonatan Winetraub, Junfeng Yang, and Asaf Cidon. Mars attacks! software protection against space radiation. In HotNets, 2023.
- [94] Andreas D. Kellas, Alan Cao, Peter Goodman, and Junfeng Yang. Divergent representations: When compiler optimizations enable exploitation. In WOOT, 2023.
- [95] Yun-Yun Tsai, Ju-Chin Chao, Albert Wen, Zhaoyuan Yang, Chengzhi Mao, Tapan Shah, and Junfeng Yang. Test-time detection and repair of adversarial samples via masked autoencoder. AdvML; non-archival workshop paper, 2023.
- [96] Tamer Eldeeb, Zhengneng Chen, Asaf Cidon, and Junfeng Yang. Neuroshard: Towards automatic multi-objective sharding with deep reinforcement learning. In aiDM, 2022.
- [97] Matthew Lawhon, Chengzhi Mao, and Junfeng Yang. Using multiple self-supervised tasks improves model robustness. PAIR2Struct; non-archival workshop paper, 2022.
- [98] Yuhong Zhong, Hongyi Wang, Yu Jian Wu, Asaf Cidon, Ryan Stutsman, Amy Tai, and Junfeng Yang. BPF for storage: An exokernel-inspired approach. In HotOS, 2021.
- [99] Jake Lee, Junfeng Yang, and Zhangyang Wang. What does cnn shift invariance look like? a visualization study. In ECCV Workshops, 2020.
- [100] Guangyu Shen, Chengzhi Mao, Junfeng Yang, and Baishakhi Ray. Fooling semantic segmentation in one step via manipulating nuisance factors. AROW; non-archival workshop paper, 2020.
- [101] Robby Costales, Chengzhi Mao, Raphael Norwitz, Bryan Kim, and Junfeng Yang. Live trojan attacks on deep neural networks. In CVPR Workshops, 2020.

- [102] Kexin Pei, Yinzhi Cao, Junfeng Yang, and Suman Jana. Towards practical verification of machine learning: The case of computer vision systems. In DeepTest, 2019.
- [103] David Williams-King and Junfeng Yang. CodeMason: Binary-level profile-guided optimization. In FEAST, 2019.
- [104] Xinhao Yuan, David Williams-King, Junfeng Yang, and Simha Sethumadhavan. Making lock-free data structures verifiable with artificial transactions. In PLOS, 2015.
- [105] Heming Cui, Rui Gu, Cheng Liu, and Junfeng Yang. Repframe: An efficient and transparent framework for dynamic program analysis. In APSys, 2015.
- [106] Junfeng Yang, Heming Cui, and Jingyue Wu. Determinism is overrated: What really makes multithreaded programs hard to get right and what can be done about it? In HotPar, 2013.
- [107] Junfeng Yang, Ang Cui, Sal Stolfo, and Simha Sethumadhavan. Concurrency attacks. In HotPar, 2012.
- [108] Oren Laadan, Chia-Che Tsai, Nicolas Viennot, Chris Blinn, Peter Senyao Du, Junfeng Yang, and Jason Nieh. Finding concurrency errors in sequential code—os-level, in-vivo model checking of process races. In HotOS, 2011.
- [109] Qifa Ke, Vijayan Prabhakaran, Yinglian Xie, Yuan Yu, Jingyue Wu, and Junfeng Yang. Optimizing data partitioning for data-parallel computing. In HotOS, 2011.
- [110] Junfeng Yang, Paul Twohey, Ben Pfaff, Can Sar, and Dawson Engler. eXplode: A lightweight, general approach to finding serious errors in storage systems. In BUGS, 2005.

Refereed or Invited Journal Publications

- [111] Kexin Pei, Zhou Xuan, Junfeng Yang, Suman Jana, and Baishakhi Ray. Learning approximate execution semantics from traces for binary function similarity. IEEE TSE, 49(4):2776–2790, 2023.
- [112] Kexin Pei, Yinzhi Cao, Junfeng Yang, and Suman Jana. DeepXplore: Automated whitebox testing of deep learning systems. CACM, 62(11):137–145, 2019.

CACM Research Highlight

- [113] Kexin Pei, Shiqi Wang, Yuchi Tian, Justin Whitehouse, Carl Vondrick, Yinzhi Cao, Baishakhi Ray, Suman Jana, and Junfeng Yang. Bringing engineering rigor to deep learning. SIGOPS OSR, 53(1):59–67, 2019.

Invited Article

- [114] Kexin Pei, Yinzhi Cao, Junfeng Yang, and Suman Jana. DeepXplore: Automated whitebox testing of deep learning systems. GetMobile, 22(3):36–38, 2019.

Invited Article

- [115] Junfeng Yang, Heming Cui, Jingyue Wu, Yang Tang, and Gang Hu. Making parallel programs reliable with stable multithreading. CACM, 57(3):58–69, 2014.

Invited Article

- [116] John MacCormick, Nicholas Murphy, Venugopalan Ramasubramanian, Udi Wieder, Junfeng Yang, and Lidong Zhou. Kinesis: A new approach to replica placement in distributed storage systems. ACM TOS, 4(4):1–28, 2009.
- [117] Junfeng Yang, Paul Twohey, Dawson Engler, and Madanlal Musuvathi. Using model checking to find serious file system errors. ACM TOCS, 24(4):393–423, 2006.

Invited Article

Preprints and Technical Reports

- [118] Alex Mathai, Shobini Iyer, Aleksandr Nogikh, Petros Maniatis, Franjo Ivančić, Junfeng Yang, and Baishakhi Ray. TRIM: Reducing AI-generated CodeSlop via agent trajectory minimization. arXiv:2607.18161, 2026.
- [119] Yun-Yun Tsai, Jia Hao Liang, Chuan Guo, Junfeng Yang, and Laurens van der Maaten. RAFF: Identifying LLM lineages via rare-region fingerprints. arXiv:2505.12682, 2026.
Revised July 2026; first released in May 2025 as RoFL: Robust Fingerprinting of Language Models.
- [120] Jinjun Peng, Magnus Saebo, Tianjun Zhong, Yi-Jie Cheng, Junfeng Yang, Baishakhi Ray, Simin Chen, and Yangruibo Ding. SWE-Spot: Building small repo-experts with repository-centric learning. arXiv:2601.21649, 2026.
- [121] Qingyuan Liu, Yun-Yun Tsai, Ruijian Zha, Victoria Li, Pengyuan Shi, Chengzhi Mao, and Junfeng Yang. LAVID: An agentic LVLM framework for diffusion-generated video detection. arXiv:2502.14994, 2025.
- [122] Leon Zhou, Junfeng Yang, and Chengzhi Mao. SPIN: Self-supervised prompt INjection. arXiv:2410.13236, 2024.
- [123] Sally Junsong Wang, Jianan Yao, Kexin Pei, Hidedaki Takahashi, and Junfeng Yang. Detecting buggy contracts via smart testing. arXiv:2409.04597, 2024.
- [124] Nageswar Keetha, Leon Wu, Gail Kaiser, and Junfeng Yang. A software checking framework using distributed model checking and checkpoint/resume of virtualized process domains. Technical Report CUCS-032-09, CS Department, Columbia University, 2009.
- [125] Nageswar Keetha, Leon Wu, Gail Kaiser, and Junfeng Yang. Distributed eXplode: A high-performance model checking engine to scale up state-space coverage. Technical Report CUCS-051-08, CS Department, Columbia University, 2008.

Patents

Junfeng Yang, Lintao Zhang, Lidong Zhou, Zhenyu Guo, Xuezheng Liu, Jian Tang, Mao Yang, Distributed system checker, US Patent 7,984,332

Rina Panigrahy, Chad Verbowski, Yinglian Xie, Junfeng Yang, Ding Yuan, Fingerprinting event logs for system management troubleshooting, US Patent 8,069,374

Qifa Ke, Vijayan Prabhakaran, Yinglian Xie, Yuan Yu, Jingyue Wu, Junfeng Yang, Optimizing data partitioning for data-parallel computing, US Patent 9,235,396

Junfeng Yang, Heming Cui, Jingyue Wu, Methods, systems, and media for providing determinism in multithreaded programs, US Patent 9,454,460

Junfeng Yang, Heming Cui, Jingyue Wu, Methods, systems, and media for protecting applications from races, US Patent 9,038,048

Selected Talks

Invited Research Talks and Panels

- 06/2026 “Verification and Validation.” Panelist, Trustworthy AI for Code Industry Roundtable, IBM, New York, NY. Moderator: Rajdeep Mukherjee.
- 10/2025 “Towards Trustworthy AI Systems.” Distinguished speaker, ACM SIGOPS Strategic Workshop, co-located with SOSP 2025.
- 10/2025 “Jailbreaks for Good: Automated, Effective Red-Teaming of AI Before It Breaks in the Wild.” Columbia Center for AI and Responsible Financial Innovation Symposium. Hosts: Agostino Capponi, Ali Hirsu.

02/2025	"Solving Hard Software Engineering Challenges using LLM Agents." Amazon. With Kexin Pei. Host: Serdar Tasiran.
05/2023	"Large Language Models Bring Unprecedented Supply-Chain Risks." Presenter and Q&A panelist, Congressional cybersecurity briefing.
08/2022	"Blessings and Curses of Software Dependencies." Keynote speaker, APSys. Hosts: Marco Serafini, Harry Xu.
05/2022	"Debugging Performance Issues in Modern Desktop Applications." Keynote speaker, ChinaSys. Host: Yu Huang.
03/2018	"Tools 2.0: Towards Effective Tools to Build Safer and More Secure Machine Learning Systems." Distinguished lecture, Purdue. Host: Yiyang Zhang.
10/2017	"Practical Model Checking of Real-world Distributed Systems." Speaker and panelist, Cloud Reliability Workshop. Host: Nemat Bidokhti.
06/2013	"How Useful Is Determinism for Reliability?" Panelist, Invited panel "Determinism: Blessing or Curse" at Fifth USENIX Workshop on Hot Topics in Parallelism. Hosts: Emery Berger, Kim Hazelwood.

Outreach and Mentoring

03/2021	"Careers in Data Science and AI." Panelist, Speyer Middle School Speaker Series. Moderator: Michelle Cristella.
11/2018	"Why Pursue a Career as a Professor." Panelist, SEAS Path to the Professorship Workshop, Columbia University, New York, NY. Host: Tiffany Simon. Moderator: Jay Sethuraman.

Professional Activities

Award Committees

2026	Selection Committee Member, ACM SIGOPS Mark Weiser Award
------	--

Conference Program Chairmanship

2027	Program Committee Co-chair, <u>21st USENIX Symposium on Operating Systems Design and Implementation (OSDI '27)</u>
2013	Program Committee Co-chair, <u>4th ACM SIGOPS Asia-Pacific Workshop on Systems (APSys '13)</u>

Conference Program Committees

SOSP	(7 times) 2023–2026, 2021, 2019, 2017
OSDI	(9 times) 2023–2026, 2021 (Best Paper Committee Chair), 2020, 2014, 2012, 2010
MLSys	2023
ASPLOS	2022
EuroSys	2021
PLDI	2017
IEEE S&P	2016
FAST	2012
APSys	2011
USENIX ATC	2011

SSV	2010
ESSoS	2010
StorageSS	2007

Conference Organization Committees

2014–2018	Steering Committee Member, <u>ACM SIGOPS Asia-Pacific Workshop on Systems (APSys)</u>
2013	Poster-Demo Session Program Committee Member, <u>10th USENIX Symposium on Networked Systems Design and Implementation (NSDI '13)</u>
2012	Poster Session Co-chair, <u>10th USENIX Symposium on Operating Systems Design and Implementation (OSDI '12)</u>
2011	Publicity Co-chair, <u>23rd ACM Symposium on Operating Systems Principles (SOSP '11)</u>

Journal and Conference Refereeing (in addition to conference program committees)

Annual Conference of the Special Interest Group on Data Communication (SIGCOMM)
USENIX Symposium on Operating Systems Design and Implementation (OSDI)
USENIX Conference on File and Storage Technologies (FAST)
USENIX Symposium on Networked Systems Design and Implementation (NSDI)
ACM Transactions on Computer Systems (TOCS)
ACM Transactions on Storage (TOS)
International Symposium on Distributed Computing (DISC)
IEEE Symposium on Security and Privacy (IEEE S&P)
IEEE Pervasive Computing
IEEE Workshop on Mobile Cloud Computing (MobiCC)

Funding Agency

Air Force Office of Scientific Research (AFOSR) Proposal Reviewer
 Department of Energy (DOE) Panelist
 National Science Foundation (NSF) Panelist; Proposal Reviewer

Professional Societies

	Fellow, Asian American Scholar Forum (AASF)
	Fellow, Association for Computing Machinery (ACM)
2019	Lifetime Member, Association of Chinese Scholars in Computing (ACSIC)
	Senior Member, Institute of Electrical and Electronics Engineers (IEEE)
	Full Member, Sigma Xi
	Member, Advanced Computing Systems Association (USENIX)

Advising and Research Supervision

(all activities at Columbia University unless otherwise noted)

Current PhD Students

2025–2031 (est.)	Weiliang Zhao (co-advised with Zhou Yu; Infosys PhD Fellow)
------------------	---

2025–2031 (est.)	Chenxi Huang (co-advised with Baishakhi Ray, Eugene Wu, and Kostis Kaffes)
2024–2030 (est.)	Hailie Mitchell (Presidential Fellow)
2024–2030 (est.)	Hideaki Takahashi (Funai Foundation Overseas Scholarship, Columbia–Ethereum Center Fellow, Capital One Fellow (expected Spring 2027))
2024–2030 (est.)	Jihwan Kim (Samsung Fellow)
2023–2029 (est.)	Alex Mathai (co-advised with Baishakhi Ray)
2023–2029 (est.)	Jinjun Peng (co-advised with Baishakhi Ray; Capital One Fellow)
2022–2028 (est.)	Harry Haoda Wang (co-advised with Asaf Cidon; NDSEG Fellow, NSF Graduate Research Fellow)
2021–2027 (est.)	Andreas Kellas (NDSEG Fellow)
2021–2027 (est.)	Raphael Jedidiah Sofaer (co-advised with Jason Nieh and Ronghui Gu; NDSEG Fellow)
2021–2027 (est.)	Yun-Yun Tsai (Dean’s Fellow)

PhD Graduates

2020–2026	Wei Hao (co-advised with Asaf Cidon; CAIT / Amazon Fellow) <i>Dissertation:</i> “Securing Machine Learning Ecosystems: Automated Debugging and Versioning for Model Derivatives” <i>Current Position:</i> Co-founder, Metlas
2019–2025	Tamer Eldeeb (co-advised with Asaf Cidon) <i>Dissertation:</i> “Optimizing Distributed Transactions via Modern AI, Storage and Networking Technologies” <i>Current Position:</i> Co-founder, Atomix DB
2018–2023	Chengzhi Mao (co-advised with Carl Vondrick) <i>Dissertation:</i> “Robust Machine Learning by Integrating Context” <i>Current Position:</i> Assistant Professor, Rutgers University
2016–2023	Kexin Pei (co-advised with Suman Jana) <i>Dissertation:</i> “Analyzing and Securing Software via Robust and Generalizable Learning” <i>Current Position:</i> Neubauer Family Assistant Professor, University of Chicago
2014–2023	Lingmei Weng (co-advised with Jason Nieh) <i>Dissertation:</i> “Utilizing Runtime Information for Accurate Root Cause Identification in Performance Diagnosis” <i>Current Position:</i> Google
2014–2020	David Williams-King <i>Dissertation:</i> “Improving Security Through Egalitarian Binary Recompilation” <i>Current Position:</i> Chief Technology Officer, Elpha Secure
2012–2019	Xinhao Yuan <i>Dissertation:</i> “Effective Randomized Concurrency Testing with Partial Order Methods” <i>Current Position:</i> Google
2011–2019	Yang Tang <i>Dissertation:</i> “Making Data Storage Efficient in the Era of Cloud Computing” <i>Current Position:</i> Clinical Associate Professor of Computer Science, New York University
2011–2018	Gang Hu <i>Dissertation:</i> “Techniques for Efficient and Effective Mobile Testing” <i>Current Position:</i> Google
2009–2015	Heming Cui <i>Dissertation:</i> “Stable Multithreading: A New Paradigm for Reliable and Secure Threads” <i>Current Position:</i> Associate Professor, University of Hong Kong

2008–2014 Jingyue Wu
Dissertation: “Sound and Precise Analysis of Multithreaded Programs through Schedule Specialization”
Current Position: NVIDIA

Current Postdoctoral Researchers

2026–now Christian Cianfarani, DSI Cybersecurity Center (co-advised with Jason Healey)
 2026–now Yaniv Nemcovsky, Infosys Topaz-Columbia University Enterprise AI Center

Former Postdoctoral Researchers

2024–2026 Penghui Li
Current Position: Assistant Professor, Nanyang Technological University

2023–2024 Chengzhi Mao (co-advised with Carl Vondrick)
Current Position: Assistant Professor, Rutgers University

2023–2024 Kexin Pei (co-advised with Baishakhi Ray and Suman Jana)
Current Position: Neubauer Family Assistant Professor, University of Chicago

2021–2024 Yaniv David
Current Position: Assistant Professor, Technion

2014–2015 Yinzhi Cao
Current Position: Professor, Johns Hopkins University

2013–2015 Yan Cui
Current Position: Staff Software Engineer, Facebook

2012–2013 Hao Li
Current Position: Industry

2010–2011 Oren Laadan (co-advised with Jason Nieh)
Current Position: Founder and Chief Technology Officer, Cellrox

Current MS and Undergraduate Researchers

2026–now Jiahao Liang, MIT undergraduate researcher (began at Columbia) (coauthored [119])

2026–now Pulak Mehrotra, MS (co-advised with Kostis Kaffes)

2025–now Chunyi “Ryan” Wang, MS

2025–now Hong Yau Chong, undergraduate (coauthored [7])

2025–now Chris Chang, undergraduate, Cyber NYC REU

2024–now Daniel Ben-Levi, undergraduate (coauthored [3, 15])

Former MS and Undergraduate Researchers

Total 27 went on to PhD programs

2024–2025 Weiliang Zhao, MS (coauthored [12, 15])
 → PhD at Columbia

2023–2025 Chenxi Huang, undergraduate at University of Minnesota (co-advised with Baishakhi Ray; coauthored [19, 88])
 → PhD at Columbia

2023–2025 Qingyuan Liu, MS (coauthored [92, 121])
 → PhD at NTU Singapore

2024–2025 Ruize Xu, MS (coauthored [91])
 → PhD at Dartmouth

2024–2025	Zirui Zhang, MS (coauthored [16]) → PhD at Rutgers
2024–2025	Suwei Ma, undergraduate, Cyber NYC REU (co-advised with Baishakhi Ray; coauthored [88]) → PhD at Penn
2024–2025	Leyi Cui, MS (co-advised with Baishakhi Ray; coauthored [90]) → PhD at CMU
2022–2024	Sally Junsong Wang, MS (coauthored [28, 123]) → PhD at Stanford (Stanford Graduate Fellowship)
2022–2024	Weichen Li, MS (coauthored [13, 23]) → PhD at UChicago
2022–2023	Xincheng Xie, MS (coauthored [33]) → PhD at UW–Madison
2021–2023	Scott Kun Geng, undergraduate, Rabi Scholar (co-advised with Carl Vondrick; coauthored [23, 36, 40]) → PhD at the University of Washington (NSF Graduate Research Fellow)
2021–2022	Lingyu Zhang, MS (coauthored [37]) → PhD at Duke
2021–2022	Aditya Senthilnathan, undergraduate at IIT-Delhi (coauthored [38]) → PhD at Cornell
2020–2022	Tal Zussman, undergraduate → PhD at Columbia
2020–2022	Zhou Xuan, undergraduate researcher at UCR (continued after graduation; coauthored [40, 111]) → PhD at Purdue
2022	Lauren Hong, MS (coauthored [17]) → PhD at NUS
2022	Lingxiao Li, MS (coauthored [17]) → PhD at Boston University
2019–2020	Jonas Guan, MS (coauthored [43, 47]) → PhD at Toronto
2019–2020	Robby Costales, undergraduate (coauthored [101]) → PhD at USC
2019	Guangyu Shen, MS at Purdue, Summer research (co-advised with Baishakhi Ray; coauthored [100]) → PhD at Purdue
2014–2016	Graham Gobieski, undergraduate, Rabi Scholar (coauthored [63]) → PhD at CMU
2015–2016	Andrew Aday, undergraduate (coauthored [61]) → PhD at Stanford
2015	Tianyu Chen, undergraduate at Tsinghua, Summer research (coauthored [66]) → PhD at Indiana
2011–2013	Younghoon Jeon, MS → PhD at Columbia
2010–2011	Chia-che Tsai, MS (coauthored [75, 78]) → PhD at Stony Brook

2010	John Gallagher, MS (coauthored [74]) → PhD at Columbia
2009	Heming Cui, MS (coauthored [78, 79]) → PhD at Columbia
<i>Other student coauthors</i>	
2024–2026	Songchen “Sophia” Yao, MS (coauthored [1])
2025–2026	Kristófer Fannar Björnsson, MS (coauthored [89])
2025–2026	Robert Leo Pendergrast, undergraduate (coauthored [89])
2025	Josef Sarfati Korich, undergraduate (coauthored [1])
2024–2025	Ruijian Zha, MS (coauthored [121])
2024–2025	Albert Jan, undergraduate (coauthored [13])
2024–2025	Ran “Henry” Li, undergraduate (coauthored [12, 21])
2022–2025	Zixi “Claire” Wang, MS (co-advised with Asaf Cidon; coauthored [14, 17])
2023–2024	Leon Zhou, MS (coauthored [122])
2024	William Lin, undergraduate at NYU, Cyber NYC REU (coauthored [16])
2024	Emanuel Mendiola-Ortiz, undergraduate at Penn State, SURE (coauthored [16])
2024	Aroon Sankoh, undergraduate at WashU, SURE (coauthored [16])
2021–2024	James Liyuan Wang, Columbia undergraduate and MS researcher (continued after graduation; coauthored [21, 42])
2023–2024	Pengyuan “Frank” Shi, MS (coauthored [92])
2023	Ju-Chin Chao, MS (coauthored [95])
2022–2023	Albert Wen, UC Berkeley undergraduate and Columbia MS researcher (coauthored [95])
2022–2023	Amrutha Sundar, MS (coauthored [35])
2021–2022	Jiayang Hu, undergraduate (coauthored [41])
2021–2022	Matthew Lawhon, MS (coauthored [97])
2021–2022	Abhishek Vaibhav Joshi, MS (coauthored [37])
2020–2021	Zhengneng Chen, MS at CMU (coauthored [96])
2019–2021	Yu Jian “Eugene” Wu, MS (co-advised with Asaf Cidon; coauthored [39, 51, 98])
2021	Aahil Awatramani, undergraduate at Cornell (coauthored [41])
2020	Matthew Broughton, MS (coauthored [43])
2020	Zhongtian Chen, undergraduate (coauthored [43])
2020	Songchen “Sophia” Yao, undergraduate (coauthored [43])
2020	Augustine Cha, MS (coauthored [46])
2019–2020	Amogh Gupta, MS (coauthored [46, 48])
2019–2020	Jake Lee, MS (coauthored [99])
2018–2019	Graham Philip Patterson, undergraduate (coauthored [51])
2016–2019	Linjie Zhu, MS (coauthored [56])
2019	Frank Spano, undergraduate (coauthored [51])
2018	Raphael Shai Norwitz, MS (coauthored [101])

2018	Bryan Sanghyuk Kim, undergraduate (coauthored [101])
2015–2017	Rui Gu, MS (coauthored [57, 60, 66, 105])
2014–2016	James P. Blake, MS (coauthored [63])
2014–2016	Michelle Zheng, undergraduate (coauthored [63])
2016	Jason Zhao, undergraduate
2015–2016	Bo Gan, MS
2014–2015	Cheng Liu, MS (coauthored [66, 105])
2015	Alexander Fangxiao Yu, undergraduate, Summer research (coauthored [61])
2013–2014	Xinan Xu, MS (coauthored [70])
2011–2013	YiHong Lin, MS (coauthored [70])
2010–2011	Huayang Guo, MS (coauthored [74, 76])
2010–2011	Chris Blinn, undergraduate (coauthored [75])
2008–2009	Nageswar Keetha, MS (co-advised with Gail Kaiser; coauthored [124, 125])
2004–2006	Can Sar, undergraduate at Stanford (coauthored [82, 83, 110])

High School Researchers

2025–now	Ezra Landa, gap-year researcher (with us since high school; coauthored [89]) → Yale
2020–2021	Vikas Ummadisetty, DUBLIN HIGH SCHOOL (coauthored [43]) → UC Berkeley (Regent’s Scholarship)
2014	Andrew Aday, GUNN HIGH SCHOOL → Columbia (Egleston Scholar)

Visiting and Other Researchers

2015–2018	Hidenori Kobayashi, Visiting Associate Research Scientist (coauthored [51, 54])
-----------	---

Thesis and Examination Record

PhD Dissertation Committees

01/23/2026	Wei Hao, “Securing Machine Learning Ecosystems: Automated Debugging and Versioning for Model Derivatives”
07/23/2025	Haoyu Li, “Unlocking Storage Performance: A Systems Approach to Kernel Bypass, Replication, and Caching”
07/14/2025	Yangruibo (Robin) Ding, “Advancing Code Intelligence with Language Models”
06/17/2025	Ioannis Zarkadas, “Optimizing Memory and Storage Performance in Cloud Datacenters”
04/30/2025	Tamer Eldeeb, “Optimizing Distributed Transactions via Modern AI, Storage and Networking Technologies”
11/14/2024	Xupeng Li, “Scale Deductive Verification to Real World Systems”
09/04/2024	Luoyao Hao, “Towards Large-Scale IoT Management: A Metadata-Oriented Approach”
06/06/2024	Jianan Yao, “Automated Verification of Safety and Liveness Properties for Distributed Protocols”
05/17/2024	Abhishek Shah, “Adaptive Sampling for Targeted Software Testing”

07/19/2023	Chengzhi Mao, "Robust Machine Learning by Integrating Context"
06/12/2023	Dongdong She, "Adaptive and Effective Fuzzing: a Data-Driven Approach"
05/10/2023	Kexin Pei, "Analyzing and Securing Software via Robust and Generalizable Learning"
04/04/2023	Lingmei Weng, "Utilizing Runtime Information for Accurate Root Cause Identification in Performance Diagnosis"
07/26/2022	Saikat Chakraborty, "Learning to Edit Code"
03/31/2022	Shiqi Wang, "Efficient Neural Network Verification using Branch and Bound"
03/10/2022	Alex Van't Hof, "New Container Architectures for Mobile, Drone, and Cloud Computing"
07/21/2021	Yuchi Tian, "Detect and Repair Errors for DNN-based Software"
03/17/2021	Shih-Wei Li, "A Secure and Formally Verified Commodity Multiprocessor Hypervisor"
08/20/2020	David Williams-King, "Improving Security Through Egalitarian Binary Recompilation"
02/27/2020	Jintack Lim, "The Design, Implementation, and Evaluation of Software and Architectural Support for Nested Virtualization on Modern Architectures"
12/09/2019	Naser AlDuaij, "Multi-mobile Computing"
08/19/2019	Marios Pomonis, "Preventing Code Reuse Attacks On Modern Operating Systems"
06/11/2019	Yang Tang, "Making Data Storage Efficient in the Era of Cloud Computing"
05/23/2019	Xinhao Yuan, "Effective Randomized Concurrency Testing with Partial Order Methods"
10/18/2018	Georgios Argyros, "Symbolic Model Learning: New Algorithms and Applications"
03/13/2018	Theofilos Petsios, "Compiler-assisted Adaptive Software Testing"
01/16/2018	Gang Hu, "Techniques for Efficient and Effective Mobile Testing"
09/26/2017	Nipun Arora, "Live Debugging: Debugging Applications On-the-fly"
05/03/2016	Jonathan Bell, "Making Software More Reliable by Uncovering Hidden Dependencies"
01/26/2016	Nicolas Viennot, "Record-Replay Mechanisms Applied to Operating Systems and Distributed Systems"
07/17/2015	Vasileios P. Kemerlis, "Protecting Commodity Operating Systems through Strong Kernel Isolation"
11/25/2014	Heming Cui, "Stable Multithreading: A New Paradigm for Reliable and Secure Threads"
05/15/2014	Jingyue Wu, "Sound and Precise Analysis of Multithreaded Programs through Schedule Specialization"
12/16/2013	Jiri Simsa, "Systematic Testing of Concurrent Programs", CARNEGIE MELLON UNIVERSITY
07/12/2011	Omer Boyaci, "High-Performance Multimedia Collaboration Tools for Application Sharing, Measuring Capture-to-display Latency, and User Created Services"
07/01/2011	Dinesh Subhraveti, "Record and vPlay: Problem Determination with Virtual Replay Across Heterogeneous Systems"
04/12/2010	Chris Murphy, "Metamorphic Testing Techniques to Detect Defects in Applications without Test Oracles"
08/04/2009	Wei Xu, "Program Transformation Techniques for Automatic Runtime Detection of Software Exploits", STONY BROOK UNIVERSITY

PhD Dissertation Proposal Committees

05/05/2026	Xuheng Li, "Scalable Verification and Synthesis of Arm System Software in the LLM Era"
12/01/2025	Yun-Yun Tsai, "Test-time Defenses with Reasoning for Safe and Robust Modern AI Systems"
04/29/2025	Haoyu Li, "Unlocking Storage Performance: A Systems Approach to Kernel Bypass, Replication, and Caching"
04/29/2025	Ioannis Zarkadas, "Enhancing Storage and Memory Performance in the Modern Data-center"
01/21/2025	Yangruibo (Robin) Ding, "Advancing Code Intelligence with Language Models"
09/13/2024	Wei Hao, "Securing Machine Learning Ecosystems: Automated Debugging and Versioning for Model Derivatives"
02/13/2024	Abhishek Shah, "Adaptive Sampling for Targeted Software Testing"
01/18/2024	Xupeng Li, "Scaling Machine-Checkable Systems Verification in Coq"
12/14/2023	Luoyao Hao, "Towards Large-Scale IoT Management: A Metadata-Oriented Approach"
02/28/2023	Tamer Eldeeb, "Optimizing Distributed Transactions via Modern AI, Storage and Networking Technologies"
03/31/2022	Chengzhi Mao, "Towards Robust Visual Perception"
12/16/2021	Dongdong She, "A Data-Driven Approach to Efficient Fuzzing"
12/07/2021	Saikat Chakraborty, "Learning To Edit Code"
11/29/2021	Kexin Pei, "Binary Analysis via Transfer Learning Trace Modeling"
11/19/2021	Shiqi Wang, "Efficient Neural Network Verification Using Branch and Bound"
12/18/2020	Yuchi Tian, "Detect and Repair Bugs/Errors in DNN-based Software"
02/25/2020	David Williams-King, "Improving Security Through Egalitarian Binary Transformation"
02/11/2020	Lingmei Weng, "Effective Diagnosis of Performance Issues with Inaccurate Causal Tracing"
05/23/2019	Jintack Lim, "The Design, Implementation, and Evaluation of Software and Architectural Support for Nested Virtualization on Modern Architectures"
01/31/2019	Alex Van't Hof, "Leveraging Higher-Level Device Abstractions for Mobile, Drone, and Desktop Computing"
12/06/2018	Marios Pomonis, "Preventing Code Reuse Attacks On Modern Operating Systems"
06/08/2018	Yang Tang, "Making Data Storage Efficient in the Era of Cloud Computing"
04/10/2018	Xinhao Yuan, "Effective Detection and Prevention of Concurrency Bugs"
06/13/2017	Gang Hu, "Automating UI testing by synthesizing reusable, robust tests"
06/13/2017	Naser AlDuaij, "Multi-Mobile Computing"
11/24/2015	Nipun Arora, "Live Debugging: Debugging Applications On-the-fly"
11/10/2015	Jonathan Bell, "Faster, More Reliable Software Builds"
06/08/2015	Vasileios P. Kemerlis, "Building Secure Operating Systems"
01/20/2015	Nicolas Viennot, "Record-Replay Mechanisms Applied to Operating Systems and Distributed Systems"
09/18/2014	Heming Cui, "Stable Multithreading: A New Paradigm for Reliable and Secure Threads"

09/06/2013	Jingyue Wu, "Sound and Precise Analysis of Multithreaded Programs through Schedule Specialization"
04/28/2011	Jiri Simsa, "Systematic Testing of Concurrent Programs", CARNEGIE MELLON UNIVERSITY
03/16/2010	Omer Boyaci, "Multimedia Tools for Application Sharing, Measuring Capture-to-display Latency, and User Created Services"
01/15/2009	Dinesh Subhraveti, "Record and Transplay: Partial Checkpointing for Fault Determination"
12/18/2008	Chris Murphy, "Using Metamorphic Testing at Runtime to Detect Defects in Applications without Test Oracles"

PhD Candidacy Exam Committees

08/20/2026	Hadleigh Schwartz
04/28/2026	Tal Zussman
04/21/2026	Raphael Jedidiah Sofaer
02/10/2026	Yuhong Zhong
05/29/2025	Andreas Kellas
08/05/2024	Harry Haoda Wang
04/24/2024	Yun-Yun Tsai
04/22/2024	Wei Hao
02/28/2024	Ioannis Zarkadas
06/27/2023	Abhishek Shah
05/31/2023	Yangruibo (Robin) Ding
05/30/2023	Kalliopi (Kelly) Kostopoulou
02/15/2023	Xupeng Li
02/10/2022	Tamer Eldeeb
11/29/2021	Chengzhi Mao
04/28/2021	Dongdong She
04/27/2021	Luoyao Hao
12/10/2020	Saikat Chakraborty
10/06/2020	Shiqi Wang
09/29/2020	Anthony Saieva Narin
05/01/2020	Kexin Pei
12/05/2019	Yuchi Tian
03/27/2018	David Williams-King
12/05/2017	Jintack Lim
06/13/2017	Lingmei Weng
12/07/2016	Xinhao Yuan
06/30/2016	Gang Hu
04/26/2016	Marios Pomonis

03/29/2016	Georgios Argyros
12/16/2015	Theofilos Petsios
05/05/2015	Naser AlDuaij
05/05/2015	David Tagatac
11/12/2013	Jonathan Bell
10/31/2013	Nipun Arora
07/23/2013	Heming Cui
05/28/2013	Nicolas Viennot
12/18/2012	Vasileios P. Kemerlis
05/06/2011	Jingyue Wu
02/03/2010	Jong Yul Kim
04/08/2009	Leon Wu

MS Dissertation Supervision

12/03/2025	Victoria Li, "Look Who's Talking Lightweight Spectrogram-Based Detection of Voice Impersonation Deepfakes"
11/24/2025	Ruijian Zha, "Optical-Flow-Guided Hybrid Detection and Thinking-Based Reasoning for AI Generated Video Forensics"
05/16/2024	Weichen Li, "A Universal LLM Framework for Semantic-Preserving Code Editing"
04/24/2024	Leon Zhou, "SPI: Self-Supervised Prompt Injections to Defend Large Language Models against Jailbreak Attacks"
04/04/2017	Rui Gu, "Reliable Synchronization on Multithreaded Servers"

MS Thesis Committees

12/19/2023	Wei Qiang, "Scaling Machine-Checkable Systems Verification in Coq"
12/12/2023	Nader Karayanni, "Towards Edge ML Ops: Monitoring and Adapting to Data Drift for Edge ML Deployments"
04/21/2021	Weifan Jiang, "Benefitting Internet Users with Secure Machine Learning and Internet Measurement Techniques"

Undergraduate Dissertation Supervision

2004–2006	Informally advised an undergraduate honors thesis by Can Sar, STANFORD UNIVERSITY
-----------	---

University Activities

University Research Leadership

2025–present	Research Lead, Secure & Responsible AI, Infosys Topaz Columbia University Enterprise AI Center
--------------	--

Departmental Service

2025–present	Undergraduate Academic Advisor
2025–present	MS Track Advisor, Computer Security
2022–2024	Adjunct Faculty Coordinator

2019–2024	MS Track Advisor, Software Systems
2019–2024	MS Admissions Committee
2009–2016	Facilities Committee
2011	Faculty Search Committee
2011	PhD Admissions Committee
2009–2024	Student Nominations Committee
2009–2013	PhD Comprehensive Examiner, Operating Systems
2012–2013	Academic Committee
2011–2012	MS Track Advisor, Software Systems
2008–present	MS Application Reviewer

Teaching

(all activities at Columbia University unless otherwise noted)

Teaching Experience

Fall 2026	COMS W4152, Engineering Software-as-a-Service. Faculty mentor. Taught by Teaching Fellow Raphael Sofaer.
Fall 2026	COMS E6998, Build an Agent Startup.
Fall 2025	COMS W4152, Engineering Software-as-a-Service.
Fall 2025	COMS E6113, Agent for Work.
Fall 2023	COMS W4152, Engineering Software-as-a-Service.
Fall 2023	COMS E6998, Engineering Blockchain and Web3 Apps.
Fall 2022	COMS W4152, Engineering Software-as-a-Service.
Fall 2022	COMS E6998, Engineering Blockchain and Web3 Apps.
Fall 2021	COMS W4995, Engineering Software-as-a-Service.
Fall 2021	COMS E6998, Security and Robustness of ML systems.
Spring 2021	COMS W4156, Advanced Software Engineering.
Spring 2021	COMS E6998, Security and Robustness of ML systems.
Spring 2020	COMS W4156, Advanced Software Engineering.
Spring 2020	COMS E6998, Security and Robustness of ML systems.
Spring 2019	COMS E6121, Reliable Software.
Spring 2019	COMS E6998, Security and Robustness of ML systems.
Spring 2018	COMS E6121, Reliable Software.
Spring 2018	COMS E6998, Security and Robustness of ML systems.
Spring 2017	COMS E6121, Reliable Software.
Spring 2016	COMS W4118, Operating Systems I.
Summer 2015	COMS W4118, Operating Systems I.
Fall 2014	COMS E6121, Reliable Software.
Fall 2013	COMS W4118, Operating Systems I.

Fall 2012	COMS E6121, Reliable Software.	
Spring 2012	COMS W4118, Operating Systems I.	
Fall 2011	COMS E6121, Reliable Software.	
Spring 2011	COMS W4118, Operating Systems I.	
Fall 2010	COMS E6998, Reliable Software.	
Spring 2010	COMS W4118, Operating Systems I.	
Fall 2009	COMS E6998, Reliable Software.	
Spring 2009	COMS W4118, Operating Systems I.	
Fall 2008	COMS E6998, How to Make Reliable Software.	
2004–2005	Teaching Assistant, CS 240 Advanced Operating Systems	STANFORD UNIVERSITY

Selected Press

- **Forbes** (2025). AI Is Behind 50% of Spam — And Now It's Hacking Your Accounts.
- **Science News Explores** (2024). A New Test Could Help Weed Out AI-Generated Text.
- **Scientific American** (2018). When AI Steers Us Astray.
- **IEEE Spectrum** (2017). A New Way to Find Bugs in Self-Driving AI Could Save Lives.
- **Network World** (2016). Solution to JIT-ROP Cyber Attacks: Scramble Code Quickly.
- **The Atlantic** (2016). Machine Unlearning.
- **The Register** (2011). Better Multithreading Offered by Columbia U Researchers.
- **LWN.net** (2006). KHB: Automating Bug Hunting.