

# Yinfang Chen

## Research Interests

My research develops *computer systems* that can both leverage and advance *artificial intelligence*, with a focus on *system reliability*. My work sits at the intersection of Systems, AI, and Software Engineering, all toward building *reliable, intelligent, and secure* computer systems.

## Education

- 2021– **University of Illinois Urbana-Champaign (UIUC)**, *Urbana, IL*.  
Present PhD student in Computer Science  
Advisor: Tianyin Xu
- 2019–2020 **National University of Singapore**, *Singapore*.  
Master of Computing
- 2015–2019 **Huazhong University of Science and Technology (HUST)**, *Wuhan, China*.  
B.Eng (Honors) in Computer Science and Technology

## Work Experience

- 08/2021– **University of Illinois Urbana-Champaign**, *Urbana, IL*.  
Present Research Assistant/Teaching Assistant, Siebel School of Computing and Data Science
- 05/2024– **Microsoft Research**, *Seattle, WA*.  
08/2024 Research Intern
- 03/2023– **Microsoft Research Asia**, *China*.  
08/2023 Research Intern
- 08/2020– **National University of Singapore**, *Singapore*.  
08/2021 Research Assistant
- 07/2018– **National University of Singapore**, *Singapore*.  
02/2020 Research Assistant

## Awards and Honors

- 2025-2026 Mavis Future Faculty Fellows (MF3) Program, UIUC
- 2025 Spotlight poster & Oral presentation in ICML'25
- 2024 Student Travel Grant for MLSys'25
- 2024 Student Travel Grant for EuroSys'24
- 2022 Student Travel Grant for S&P'22
- 2017 Best Creative Award in China-US Young Maker Competition
- 2017 University Pacemaker to Merit Student Scholarship (Top 1%), HUST
- 2016 Culture and Sports Scholarship (Top 10%), HUST
- 2015 Culture and Sports Scholarship (Top 10%), HUST

## Publications

### Peer-Reviewed Conference Papers

- [NeurIPS'25] **Yinfang Chen\***, Jiaqi Pan\*, Jackson Clark\*, Yiming Su\*, Noah Zheutlin, Bhavya Bhavya, Rohan Arora, Yu Deng, Saurabh Jha, and Tianyin Xu. (**\*: Co-primary authors**) "STRATUS: A Multi-agent System for Autonomous Reliability Engineering of Modern Clouds." In *Proceedings of the Thirty-Ninth Annual Conference on Neural Information Processing Systems (NeurIPS)*, December 2025.
- [MLSys'25] **Yinfang Chen**, Manish Shetty, Gagan Somashekar, Minghua Ma, Yogesh Simmhan, Jonathan Mace, Chetan Bansal, Rujia Wang, and Saravan Rajmohan. "AIOpsLab: A Holistic Framework to Evaluate AI Agents for Enabling Autonomous Clouds." In *Proceedings of the Eighth Annual Conference on Machine Learning and Systems (MLSys)*, May 2025.
- [ICML'25] Saurabh Jha, Rohan Arora, Yuji Watanabe, Takumi Yanagawa, **Yinfang Chen**, Jackson Clark, Bhavya Bhavya, Mudit Verma, Harshit Kumar, Hirokuni Kitahara, Noah Zheutlin, Saki Takano, Divya Pathak, Felix George, Xinbo Wu, Bekir O. Turkkan, Gerard Vanloo, Michael Nidd, Ting Dai, Oishik Chatterjee, and Pranjal Gupta, Suranjana Samanta, Pooja Aggarwal, Rong Lee, Pavankumar Murali, Jae-wook Ahn, Debanjana Kar, Ameet Rahane, Carlos Fonseca, Amit Paradkar, Yu Deng, Pratibha Moogi, Prateeti Mohapatra, Naoki Abe, Chandrasekhar Narayanaswami, Tianyin Xu, Lav R. Varshney, Ruchi Mahindru, Anca Sailer, Laura Shwartz, Daby Sow, Nicholas C. M. Fuller, and Ruchir Puri. "IT-Bench: Evaluating AI Agents across Diverse Real-World IT Automation Tasks." In *Proceedings of the Forty-Second International Conference on Machine Learning (ICML)*, July 2025. **Spotlight & Oral.**
- [ISSRE'25] Haoran Yan\*, **Yinfang Chen\***, Minghua Ma, Ming Wen, Shan Lu, Shenglin Zhang, Tianyin Xu, Rujia Wang, Chetan Bansal, Saravan Rajmohan, Chaoyun Zhang, and Dongmei Zhang. (**\*: Equal contributions**) "An Empirical Study of Production Incidents in Generative AI Cloud Services." In *Proceedings of the 36th IEEE International Symposium on Software Reliability Engineering (ISSRE)*, October 2025.
- [ICSE'25] Xinyu Lian\*, **Yinfang Chen\***, Runxiang Cheng, Jie Huang, Parth Thakkar, and Tianyin Xu. (**\*: Co-primary authors**) "Large Language Models as Configuration Validators." In *Proceedings of the 47th International Conference on Software Engineering (ICSE)*, April 2025.
- [ICSE'25] Anna Mazhar, Saad Sher Alam, Xinze William Zheng, **Yinfang Chen**, Suman Nath, and Tianyin Xu. "Fidelity of Cloud Emulators: The Imitation Game of Testing Cloud-based Software." In *Proceedings of the 47th International Conference on Software Engineering (ICSE)*, April 2025.
- [EuroSys'24] **Yinfang Chen**, Huaibing Xie, Minghua Ma, Yu Kang, Xin Gao, Liu Shi, Yunjie Cao, Xuedong Gao, Hao Fan, Ming Wen, Jun Zeng, Supriyo Ghosh, Xuchao Zhang, Qingwei Lin, Saravan Rajmohan, Dongmei Zhang, and Tianyin Xu. "Automatic Root Cause Analysis via Large Language Models for Cloud Incidents." In *Proceedings of the 19th European Conference on Computer Systems (EuroSys)*, April 2024.
- [SoCC'24] Manish Shetty, **Yinfang Chen**, Gagan Somashekar, Minghua Ma, Yogesh Simmhan, Xuchao Zhang, Jonathan Mace, Dax Vandevor, Pedro Las-Casas, Shachee Mishra Gupta, Suman Nath, Chetan Bansal, and Saravan Rajmohan. "Building AI Agents for Autonomous Clouds: Challenges and Design Principles." In *Proceedings of the ACM Symposium on Cloud Computing (SoCC)*, November 2024.

- [NSDI'23] **Yinfang Chen**, Xudong Sun, Suman Nath, Ze Yang, and Tianyin Xu. "Push-Button Reliability Testing for Cloud-Backed Applications with Rainmaker." In *Proceedings of the 20th USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, April 2023.
- [Oakland'23] Muhammad Adil Inam, **Yinfang Chen**, Akul Goyal, Jason Liu, Jaron Mink, Noor Michael, Sneha Gaur, Adam Bates, and Wajih Ul Hassan. "SoK: History is a Vast Early Warning System: Auditing the Provenance of System Intrusions." In *Proceedings of the 44th IEEE Symposium on Security and Privacy (Oakland)*, May 2023.
- [Oakland'22] Jun Zeng, Xiang Wang, Jiahao Liu, **Yinfang Chen**, Zhenkai Liang, Tat-Seng Chua, and Zheng Leong Chua. "ShadeWatcher: Recommendation-guided Cyber Threat Analysis using System Audit Records." In *Proceedings of the 43rd IEEE Symposium on Security and Privacy (Oakland)*, May 2022.
- [NDSS'21] Jun Zeng, Zheng Leong Chua, **Yinfang Chen**, Kaihang Ji, Zhenkai Liang, and Jian Mao. "WATSON: Abstracting Behaviors from Audit Logs via Aggregation of Contextual Semantics." In *Proceedings of the 28th Annual Network and Distributed System Security Symposium (NDSS)*, February 2021.

## Teaching

- Spring 2026 Teaching Assistant, UIUC CS 523: Advanced Operating Systems (forthcoming), IL, USA
- Spring 2024 Teaching Assistant, UIUC CS 598 AI Efficiency: Systems & Algorithms, IL, USA
- Fall 2023 Teaching Assistant, UIUC CS/STAT 107: Data Science Discovery, IL, USA
- Fall 2021 Teaching Assistant, UIUC CS 241: System Programming, IL, USA
- Summer 2019 & 2018 Teaching Assistant, School of Computing Summer Workshop: Computer Graphics, National University of Singapore, Singapore

## Service

- Reviewer Machine Learning for Systems Workshop 2025 (MLForSys)
- Computing Surveys (CSUR)
- IEEE Transactions on Information Forensics & Security (TIFS)
- Future Generation Computer Systems (FGCS)
- Information Fusion
- The Journal of Open Source Software (JOSS)
- The Computer Journal (Oxford)
- Machine Learning for Systems Workshop 2024 (MLForSys)
- Organizer Co-organizer of 2022 Brett Daniel Software Engineering Seminar, UIUC
- Slack Co-Chair (Student Volunteer) at SOSP'21

## Talks

### Agentic AI for Cloud System Operations

- Invited Talk, IBM Hybrid Cloud & AI Thrust Interlock Series, Virtual, October 2025
- Invited Talk, IBM AI Automation Reading Group, Virtual, October 2025
- Invited Talk, ByteDance, Virtual, August 2025
- Invited Talk, Autonomus IT Agents - IBM UIUC, Virtual, August 2025
- Guest Lecture, UIUC CS 523: Advanced Operating Systems, IL, USA
- Invited Talk, The 4th SJTU AI International Young Scholars Forum, Virtual, July 2025
- Conference Talk, MLSys'25, Santa Clara, CA, May 2025

### Automatic Root Cause Analysis via Large Language Models

- Conference Talk, EuroSys'24, Athens, Greece, April 2024
- Invited Talk, MLSys Seminar @Singapore, November 2023
- Invited Talk, UIUC Systems Reading Group, IL, November 2023

### Push-Button Reliability Testing for Cloud-Backed Applications

- Invited Talk, School of Computing, NUS, Singapore, May 2023
- Invited Talk, Sys-Net Area Retreat at Allerton Park, IL, May 2022

### Presentations of External Research

- S&P conference debrief, UIUC Systems Reading Group, IL, Fall 2022
- Brett Daniel Software Engineering Seminar, IL, March 2022
- CS 598 Advanced Software Testing and Debugging, IL, April 2022
- UIUC Systems Reading Group, IL, September 2021

---

## Selected Press

- Nov. 2025 IBM Research Blog. *Ctrl+Z for agents.*
- Nov. 2025 TipRanks. *IBM Builds an AI 'Undo Button' to Fix Cloud Outages Faster.*
- Oct. 2025 LinkedIn post. *Operate, Fail, Observe — Train Agents for Autonomous Clouds.*
- Sep. 2025 Medium article. *Transactional No-Regression (TNR): Safe-Rollback Strategies in Autonomous Agentic SRE Systems.*
- May. 2025 The New Stack. *Generative AI Incidents Hit Different.*
- May 2025 CIO. *IBM aims to set industry standard for enterprise AI with ITBench SaaS launch.*
- Feb. 2025 IBM Research Blog. *AI agents have the potential to revolutionize work — and now you can measure if they actually are.*
- Dec. 2024 JUTEQ. *2024's Most Powerful AI Agent Papers.*
- Dec. 2024 MarkTechPost. *An Open-Source Comprehensive AI Framework for AIOps Agents.*
- Dec. 2024 Medium article. *Revolutionizing Cloud Operations with an Open-Source AI Framework.*
- Nov. 2024 Microsoft Research Blog. *Research Focus: Week of November 11, 2024.*
- July 2023 The Weekend Read. *Reliability Testing for Cloud-Backed Applications.*

---

## Open-Source Artifacts

- AIOpsLab <https://github.com/microsoft/AIOpsLab> (Star: 741, Fork: 125)
- ITBench <https://github.com/itbench-hub/ITBench> (Star: 249, Fork: 23)  
ITBench Scenarios can be found here: [ITBench-Scenarios](#)
- Stratus <https://github.com/xlab-uiuc/stratus>
- SREGym <https://github.com/SREGym/SREGym>
- Ciri <https://github.com/xlab-uiuc/ciri>
- Cloudtest <https://github.com/xlab-uiuc/cloudtest>
- Rainmaker <https://github.com/xlab-uiuc/rainmaker>  
We detected 73 bugs (55 confirmed and 51 fixed). All bugs can be found in [bugs.md](#)
- Literature <https://github.com/yinfangchen/cloud-incident-lit>
- ShadeWatcher <https://github.com/jun-zeng/ShadeWatcher>

## ■ Mentoring

I mentored the research of six graduate students and two undergraduate students. All of them have co-authored at least one publication to a top Systems, AI, or Software Engineering venue with me.

PhD Student	Jackson Clark, UIUC	<a href="https://hacksonclark.github.io/">https://hacksonclark.github.io/</a>
PhD Student	Yiming Su, UIUC	<a href="https://yimingsu01.github.io/">https://yimingsu01.github.io/</a>
PhD Student	Xinyu Lian, UIUC	<a href="https://xinyulian.tech/">https://xinyulian.tech/</a>
MS Student	Anna Mazhar, UIUC	<a href="https://anna-mazhar.github.io/">https://anna-mazhar.github.io/</a>
MS Student	Huaibing Xie, Peking University	
MS Student	Haoran Yan, HUST	
Undergrad	Jiaqi Pan, Tsinghua University	
Undergrad	Ze Yang, UIUC	