

Linyi Li

Bellevue, WA 98004, USA

☎ (+1) 217 205 6216 • ✉ linyi2@illinois.edu • 🌐 linyil.com

Ph.D. in computer science with research interests in **machine learning**, **computer security** and **software engineering**, with a special focus on building **certifiably trustworthy** deep learning systems.

Education

University of Illinois Urbana-Champaign

○ *Ph.D. in Computer Science*

Aug 2018 – Aug 2023

- Advisor: *Prof. Bo Li* Co-advisor: *Prof. Tao Xie*
- Thesis proposal: Certifying trustworthy deep learning systems at scale

Tsinghua University

○ *Bachelor of Computer Science and Technology*

Beijing, China

Aug 2014 – Jul 2018

- GPA: Major: 91.6/100 Overall: 90.1/100
- Advisor: *Prof. Xiaoying Bai*
- Thesis: Model-Based Automated Web API Test Generation.
- Tsinghua University Outstanding Undergraduate, Class of 2018
- Excellent Undergraduate, Department of Computer Science and Technology

Publications

(* stands for equal contribution) (first / co-first / corresponding author publications highlighted)

28. **Linyi Li**. [Certifiably Trustworthy Deep Learning Systems at Scale](#). *Doctoral Thesis*.
27. Zhangheng Li, Tianlong Chen, **Linyi Li**, Bo Li, Zhangyang Wang. Can Pruning Improve Certified Robustness of Neural Networks? *Transactions on Machine Learning Research (TMLR)*, 2023.
26. **Linyi Li**, Tao Xie, Bo Li. [SoK: Certified Robustness for Deep Neural Networks](#). *IEEE Symposium on Security and Privacy (SP)* 2023.
25. **Linyi Li**, Yuhao Zhang, Luyao Ren, Yingfei Xiong, Tao Xie. [Reliability Assurance for Deep Neural Network Architectures Against Numerical Defects](#). *International Conference on Software Engineering (ICSE)* 2023.
24. Jiawei Zhang, **Linyi Li**, Ce Zhang, Bo Li. CARE: Certifiably Robust Learning with Reasoning via Variational Inference. *IEEE Conference on Secure and Trustworthy Machine Learning (SatML)* 2023.
23. Mintong Kang, **Linyi Li**, Bo Li. FaShapley: Fast and Approximated Shapley Based Model Pruning Towards Certifiably Robust DNNs. *IEEE Conference on Secure and Trustworthy Machine Learning (SatML)* 2023.
22. Mintong Kang*, **Linyi Li***, Maurice Weber, Yang Liu, Ce Zhang, Bo Li. [Certifying Some Distributional Fairness with Subpopulation Decomposition](#). *Advances in Neural Information Processing Systems 35 (NeurIPS 2022)*.
21. Xiaojun Xu, **Linyi Li**, Bo Li. LOT: Layer-wise Orthogonal Training on Improving ℓ_2 Certified Robustness. *Advances in Neural Information Processing Systems 35 (NeurIPS 2022)*.
20. Bhaskar Ray Chaudhury, **Linyi Li**, Mintong Kang, Bo Li, Ruta Mehta. Fairness in Federated Learning via Core-Stability. *Advances in Neural Information Processing Systems 35 (NeurIPS 2022)*.
19. Huan Zhang*, Shiqi Wang*, Kaidi Xu*, **Linyi Li**, Bo Li, Suman Jana, Cho-Jui Hsieh, J. Zico Kolter. General Cutting Planes for Bound-Propagation-Based Neural Network Verification. *Advances in Neural Information Processing Systems 35 (NeurIPS 2022)*.

18. Zhuolin Yang*, Zhikuan Zhao*, Boxin Wang, Jiawei Zhang, **Linyi Li**, Hengzhi Pei, Bojan Karlaš, Ji Liu, Heng Guo, Ce Zhang, Bo Li. Improving Certified Robustness via Statistical Learning with Logical Reasoning. *Advances in Neural Information Processing Systems 35 (NeurIPS 2022)*.
17. Hanjiang Hu, Zuxin Liu, **Linyi Li**, Jiacheng Zhu, Ding Zhao. Robustness Certification of Visual Perception Models via Camera Motion Smoothing. *6th Annual Conference on Robot Learning (CoRL 2022)*.
16. **Linyi Li**, Jiawei Zhang, Tao Xie, Bo Li. [Double Sampling Randomized Smoothing](#). *International Conference on Machine Learning (ICML) 2022*. [\[Video\]](#)
15. Wenda Chu, **Linyi Li**, Bo Li. TPC: Transformation-Specific Smoothing for Point Cloud Models. *International Conference on Machine Learning (ICML) 2022*.
14. Maurice Weber, **Linyi Li**, Boxin Wang, Zhikuan Zhao, Bo Li, Ce Zhang. Certifying Out-of-Domain Generalization for Blackbox Functions. *International Conference on Machine Learning (ICML) 2022*.
13. Fan Wu*, **Linyi Li***, Chejian Xu, Huan Zhang, Bhavya Kailkhura, Krishnaram Kenthapadi, Ding Zhao, Bo Li. [COPA: Certifying Robust Policies for Offline Reinforcement Learning against Poisoning Attacks](#). *International Conference on Learning Representations (ICLR) 2022*.
12. Fan Wu, **Linyi Li**, Zijian Huang, Yevgeniy Vorobeychik, Ding Zhao, Bo Li. CROP: Certifying Robust Policies for Reinforcement Learning through Functional Smoothing. *International Conference on Learning Representations (ICLR) 2022*.
11. Zhuolin Yang*, **Linyi Li***, Xiaojun Xu, Bhavya Kailkhura, Tao Xie, Bo Li. [On the Certified Robustness for Ensemble Models and Beyond](#). *International Conference on Learning Representations (ICLR) 2022*.
10. Ripon Saha, Akira Ura, Sonal Mahajan, Chenguang Zhu, **Linyi Li**, Yang Hu, Hiroaki Yoshida, Sarfraz Khurshid, Mukul R. Prasad. SapientML: Synthesizing Machine Learning Pipelines by Learning from Human-Written Solutions. *International Conference on Software Engineering (ICSE) 2022*.
9. Zhuolin Yang*, **Linyi Li***, Xiaojun Xu*, Shiliang Zuo, Qian Chen, Benjamin Rubinstein, Ce Zhang, Bo Li. [TRS: Transferability Reduced Ensemble via Encouraging Gradient Diversity and Model Smoothness](#). *Advances in Neural Information Processing Systems 34 (NeurIPS 2021)*.
8. Jiawei Zhang*, **Linyi Li***, Huichen Li, Xiaolu Zhang, Shuang Yang, Bo Li. [Progressive-Scale Boundary Blackbox Attack via Projective Gradient Estimation](#). *International Conference on Machine Learning (ICML) 2021*. [\[Video\]](#)
7. **Linyi Li***, Maurice Weber*, Xiaojun Xu, Luka Rimanic, Bhavya Kailkhura, Tao Xie, Ce Zhang, Bo Li. [TSS: Transformation-Specific Smoothing for Robustness Certification](#). *ACM Conference on Computer and Communications Security (CCS) 2021*. [\[Video\]](#)
6. Huichen Li*, **Linyi Li***, Xiaojun Xu, Xiaolu Zhang, Shuang Yang, Bo Li. [Nonlinear Projection Based Gradient Estimation for Query Efficient Blackbox Attacks](#). *International Conference on Artificial Intelligence and Statistics (AISTATS) 2021*.
5. **Linyi Li**, Zhenwen Li, Weijie Zhang, Jun Zhou, Pengcheng Wang, Jing Wu, Guanghua He, Xia Zeng, Yuetang Deng, Tao Xie. [Clustering Test Steps in Natural Language toward Automating Test Automation](#). *ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering (ESEC/FSE) 2020, Industry Track*. [\[Video\]](#)
4. **Linyi Li***, Zexuan Zhong*, Bo Li, Tao Xie. [Robustra: Training Provable Robust Neural Networks over Reference Adversarial Space](#). *International Joint Conference on Artificial Intelligence (IJCAI) 2019*.
3. Klas Leino, Shayak Sen, Anupam Datta, Matt Fredrikson, **Linyi Li**. Influence-Directed Explanations for Deep Convolutional Networks. *International Test Conference (ITC) 2018*.
2. Junyi Wang, Xiaoying Bai, **Linyi Li**, Zhicheng Ji, Haoran Ma. A Model-Based Framework For Cloud API Testing. *Computer Software and Applications Conference (COMPSAC) 2017*.
1. Junyi Wang, Xiaoying Bai, Haoran Ma, **Linyi Li**, Zhicheng Ji. Cloud API Testing. *IEEE International Conference on Software Testing, Verification and Validation Workshops (ICSTW) 2017*.

Selected Awards

- Rising Star in Data Science, the University of Chicago (among 32 awardees) 2022
- NeurIPS 2022 Scholar Award 2022
- AdvML Rising Star Award (among 2 awardees) 2022
- 1st Place, 3rd International Verification of Neural Networks Competition (VNN-COMP'22) 2022
- Qualcomm Innovation Fellowship Finalist (among 44 in North America) 2022
- Two Sigma PhD Fellowship Finalist (among 13 worldwide) 2022
- ACM CCS Conference Travel Award 2021
- 2nd Place, ICPC Mid-Central USA Regional Contest 2019
- 3rd Place, ICPC Mid-Central USA Regional Contest 2018
- Wing Kai Cheng Fellowship 2018
- Tsinghua University Outstanding Undergraduate, Class of 2018 (301 of 3555) 2018
- Excellent Undergraduate, Department of Computer Science and Technology at Tsinghua 2018
- Academic Excellence Award with HUAWEI Scholarship 2017
- “Sogou Cup” Artificial Intelligence Programming Contest Top 16 2015
- Top 0.03% in the National College Entrance Exam 2014
- National Olympiad in Informatics, Bronze Medal 2013
- National Olympiad in Informatics in Provinces, First Prize 2012

Selected Talks

Certifiable Deep Learning at Scale towards Trustworthy Machine Learning

- Talk at SE4SafeML Workshop, FSE 2023 *Dec 2023*
- Invited talk at AI Metacognitive Workshop, Arizona State University *Nov 2023*
- Invited virtual talk at Meta Central Applied Science *Sept 2023*
- Invited talk at Nanyang Technological University *Jul 2023*
- Virtual talk at Google Deepmind *Jul 2023*
- Invited talk at Washington University in St. Louis *Mar 2023*
- Invited talk at Georgia Institute of Technology *Mar 2023*
- Invited talk at Simon Fraser University *Mar 2023*
- Invited talk at Rochester Institute of Technology *Feb 2023*
- Invited virtual talk at Microsoft Research New England *Feb 2023*
- Virtual talk at Machine Learning & Security Seminar, Purdue University *Dec 2022*
- Invited talk at Data Science Institute, the University of Chicago *Nov 2022*

Large-Scale Certifiably Trustworthy Machine Learning

- Invited virtual talk at Lockheed Martin Corporation *Nov 2022*
- Virtual talk at Machine Heart platform *Nov 2022*
- Webinar at TrustML Young Scientist Seminars, RIKEN AIP *Aug 2022*
- Talk at 4th Workshop on Adversarial Learning Methods for Machine Learning and Data Mining (AdvML), KDD 2022 *Aug 2022*

Double Sampling Randomized Smoothing

- Virtual talk at AI Time platform *Aug 2022*

Boosting Certified Robustness of Deep Neural Networks Classifiers against Semantic Transformations

- Virtual talk at ICRA 2022 Workshop on Trustworthy Autonomy and Robotics May 2022
- **Certified Robustness for Deep Neural Networks: Overview and Outlook**
- Webinar at Jiangmen platform Feb 2022
- Virtual talk at Visual Informatics Group, University of Texas at Austin Oct 2021
- Webinar at Safe AI, Bilibili Mar 2021
- Virtual talk at Safe AI Lab, Carnegie Mellon University Mar 2021
- Virtual talk at Workshop on Robust Artificial Intelligence, Lorentz Center Jan 2021

Teaching and Mentorship Experiences

- **Logic and AI (Graduate Level)** **Lead Teaching Assistant**
University of Illinois Urbana-Champaign *Aug 2021 – Dec 2021*
 - Co-designed the *first version* of the course from scratch.
 - Lead the course project design and grading.
 - Setup infrastructure and help the lecture design of the new course.
- **Data Structure (Undergraduate Level)** **Teaching Assistant**
Tsinghua University *Sept 2015 – Jan 2016*
 - Host two seminars for homework problem discussions.
 - Contribute several original problems for assignments and exams.
- **Undergraduate Research Intern Co-Mentorship**
Mentored students:
 - Mintong Kang *Nov 2021 – May 2022*
Paper published at NeurIPS 2022 on certified fairness. Now a PhD student at UIUC.
 - Chenhui Zhang *Dec 2021 – May 2022*
Paper submitted on ensemble pruning for certified robustness. Now an undergraduate student at UIUC.
 - Wenda Chu *Nov 2021 – Feb 2022*
Paper published at ICML 2022 on certification of point cloud models. Now an undergraduate student at Tsinghua University.
 - Jiawei Zhang *Aug 2020 – Feb 2021*
Paper published at ICML 2021 on black-box neural network attacks. Now a master student at UIUC.

Internship & Industry Experiences

- **ByteDance Inc.** **Bellevue, WA**
Research Scientist *Aug 2023 – ??? 2024*
 - Research on data-centric methods and systematic evaluation for large language models.
- **Microsoft Research Lab - New England** **Cambridge, MA**
Research Intern mentored by Dr. Adam Kalai *May 2022 – Aug 2022*
 - Program synthesis by finetuning from large language models with a handcrafted distributed training framework and a novel generative adversarial RL-inspired training paradigm.
- **Fujitsu Laboratories of America** **Remote**
Research Intern mentored by Dr. Mukul Prasad *May 2021 – Aug 2021*
 - Program Synthesis for AutoML based on learning from mined corpus and static analysis based data augmentation.
 - Lead to a paper accepted by ICSE 2022.

- **Microsoft** **Redmond, WA**
Data Scientist Intern mentored by Dr. Neel Sundaresan *Jun 2019 – Aug 2019*
 - Build an efficient search engine for PR comments and commits.
 - Utilize transformer models for unsupervised commit classification and code change pattern extraction.
- **Carnegie Mellon University** **Pittsburgh, PA**
Undergrad Research Intern mentored by Prof. Matt Fredrikson *Jun 2017 – Sept 2017*
 - Apply integrated gradients to explain and visualize convolutional neural networks.
 - Develop an automatic method to capture vital lesions for diabetic retinopathy diagnosis, leading to a paper accepted by ITC 2018.
- **Sogou Inc.** **Beijing, China**
Back-end Engineer Intern *Aug 2015 – Oct 2015*
 - Design the interfaces between back-end and front-end for a tutor ordering platform.
 - Implement an efficient and advanced tutor search module that supported multiple keys.

Selected Open-Source Projects

- (Ongoing Work) Infocoder — powerful large language model for code.
 - Evaluation framework & leaderboard: <https://infi-coder.github.io/inficoder-eval/>
- Developer of leaderboard and toolbox on provable training and verification approaches for DNNs.
 - Leaderboard:
 (new version) <https://github.com/sokcertifiedrobustness/sokcertifiedrobustness.github.io>
 (old version) <https://github.com/AI-secure/Certified-Robustness-SoK-Oldver>
 - Toolbox: <https://github.com/AI-secure/VeriGauge>
- Developer of TSS: transformation-specific smoothing-based robustness certification against geometric perturbations.
<https://github.com/AI-secure/semantic-randomized-smoothing>
 - State-of-the-art verification approach for robustness against geometric perturbations.
 - Accompanying paper published at CCS 2021.
- Key contributor of α - β -CROWN (alpha-beta-CROWN), a scalable neural network verifier.
<https://github.com/huanzhang12/alpha-beta-CROWN>
 - 2x winner of International Verification of Neural Networks Competition (VNN-COMP'21, '22).
 - Accompanying paper published at NeurIPS 2022.

Services

- NeurIPS 2022, Workshop on Trustworthy and Socially Responsible Machine Learning Organizer
- NeurIPS (2021, 2022, 2023) PC Member
- ICML (2022, 2023) PC Member
- ICLR (2021, 2022, 2023) PC Member
- KDD (2023) PC Member
- AAAI (2022) PC Member
- UAI (2021, 2022, 2023) PC Member
- AISTATS (2021) PC Member
- TPAMI Reviewer

- TMLR Reviewer
- Neurocomputing Reviewer
- NeurIPS 2022, Workshop on ML Safety PC Member
- ICML 2022, Workshop on Formal Verification of Machine Learning PC Member
- KDD (2020-), Workshop on Adversarial Learning Methods for Machine Learning and Data Mining PC Member
- ICML 2019, Workshop on the Security and Privacy of Machine Learning (SPML) PC Member
- CVPR 2019, Workshop on Adversarial Machine Learning in Real-World Computer Vision Systems (AdvMLCV) PC Member
- ICLR 2023, Backdoor Attacks and Defenses in Machine Learning (BANDS) PC Member

Social Engagement

- Student mentor for new PhD students in computer science at UIUC. 2022
- Graduate ambassador for prospective PhD students in computer science at UIUC. 2021
- ACM student member. 2021 - present