

# Shiqi Wang

Ph.D. Candidate at Columbia University

Phone: (+1) 3474173258 Mail: tcwangshiqi@cs.columbia.edu

Personal Webpage: <https://www.cs.columbia.edu/tcwangshiqi/>

## EDUCATION

---

**Ph.D. in Computer Science** Columbia University (Sept 2017 — Mar 2022)

- **Advisor:** Prof. Suman Jana

**M.S. in Computer Science** Columbia University (Sept 2017 — May 2019)

- **Advisor:** Prof. Suman Jana

**B.S. in Computer Science** Shanghai Jiaotong University (Sept 2013 — June 2017)

## WORK EXPERIENCE

---

**Computer Vision Research Intern** June 2021 — Sept 2021  
Microsoft Research Redmond, US

- **Mentor:** Hamid Vaezi Joze
- **Research:** Autoencoder optimization, pruning, and acceleration for object detection, face enhancement, and super resolution.

**Security Research Intern** June 2020 — Sept 2020  
IBM Research Yorktown Height, US

- **Mentor:** Kevin Eykholt, Taesung Lee, Jiyong Jang, Ian Molloy
- **Research:** Neural network certifiable training accounting for label similarity (published in AAI 2021).

**Security Research Intern** June 2019 — Dec 2019  
Baidu Security X-Lab Sunnyvale, US

- **Mentor:** Yunhan Jia, Zhenyu Zhong
- **Research:** Pruning to accelerate neural network adversarial training; Incorporate neural network verification into Baidu adversarial robustness evaluation toolbox *Perceptron*.

## RESEARCH TOPICS

---

### Topic 1: Neural Network Verification [VNN]

- I am one of the pioneers in the field of neural network verification (i.e., guarantee the correctness of model predictions under arbitrary adversary). In 2018, I proposed a novel bound propagation method called symbolic interval analysis to efficiently analyze and verify various properties on targeted networks. Using this method, I implemented one of the first branch-and-bound complete verifiers, ReluVal and Neurify (published in Usenix Security 2018 and NeurIPS 2018) significantly outperformed existing verifiers.
- My teammates and I built so far the strongest neural network verifier  $\alpha, \beta$ -CROWN, winning the 2nd neural network verification competition in 2021 with the highest score. It arouses wide attention and was covered by several media press. We are also invited to give a 3-hr tutorial in AAI 2022. I am one of the main contributors to the core strategies in  $\alpha, \beta$ -CROWN including optimizable bound propagation method and branch and bound refinement accelerated on GPUs (published in ICLR 2021 and NeurIPS 2021).

### Topic 2: Model Pruning, Optimization, and Acceleration [Prune]

- During the internship at Microsoft Research in 2021, I worked on pruning to optimize and accelerate autoencoders for multiple vision tasks. With a novel pruning method using knowledge distillation, I can train models up to 100x smaller and over 10x faster. Also, their inference cost can be adaptive according to the computation resources available.
- I studied how to train networks that are both compact and (verifiably) robust. We proposed a novel pruning framework Hydra (published in NeurIPS 2020) that can guide prune by desired objectives, pruning up to 99% weights while getting desired robustness.

### Topic 3: Adversarial Machine Learning [AdvML]

- Strong adversarial attacks are important baselines for evaluating model robustness. I proposed interval attack relying on symbolic interval analysis which was once the strongest attack on MadryLab MNIST challenge leaderboard in 2019. Recently, we proposed BaB attack extended from our verifier  $\alpha, \beta$ -CROWN and achieved state-of-the-art attack performance.
- I have done several papers on accelerating adversarial training and certifiable training to robustify neural networks.

### Topic 4: Security [Sec]

- We designed new algorithms including certifiable training and robust tree ensembles to secure ML applications for security tasks like PDF malware classifications and malicious Twitter account detections (published in Usenix Security 2020 & 2021, CCS 2021).

**Skills:** Programming (Python, C/C++); ML (Tensorflow, PyTorch); Latex; Over 10 technical talks in top-tier ML/Security conferences.

## PUBLICATIONS (\* INDICATES EQUAL CONTRIBUTION)

---

Google scholar profile (900+ citations): [https://scholar.google.com/citations?user=u\\_MzXeMAAAAJ](https://scholar.google.com/citations?user=u_MzXeMAAAAJ)

- [VNN] Beta-CROWN: Efficient Bound Propagation with Per-neuron Split Constraints for Complete and Incomplete Neural Network Verification, **Shiqi Wang\***, Huan Zhang\*, Kaidi Xu\*, Xue Lin, Suman Jana, Cho-Jui Hsieh, J. Zico Kolter, *Conference on Neural Information Processing Systems (NeurIPS)*, 2021.
- [Sec] Learning Security Classifiers with Verified Global Robustness Properties, Yizheng Chen, **Shiqi Wang**, Yue Qin, Xiaojing Liao, Suman Jana, and David Wagner, *Conference on Computer and Communications Security (CCS)*, 2021. [Best Paper Award Runner-Up].
- [VNN] Fast and Complete: Enabling Complete Neural Network Verification with Rapid and Massively Parallel Incomplete Verifiers, Kaidi Xu\*, Huan Zhang\*, **Shiqi Wang**, Yihan Wang, Suman Jana, Xue Lin, Cho-Jui Hsieh, *International Conference on Learning Representations (ICLR)*, 2021.
- [AdvML] Adaptive Verifiable Training Using Pairwise Class Similarity, **Shiqi Wang**, Kevin Eykholt, Taesung Lee, Jiyong Jang, Ian Mollo, *AAAI Conference on Artificial Intelligence (AAAI)*, 2021.
- [Sec] Cost-Aware Robust Tree Ensembles for Security Applications, Yizheng Chen, **Shiqi Wang**, Weifan Jiang, Asaf Cidon, Suman Jana, *USENIX Security Symposium (Usenix Security)*, 2021.
- [Prune] On Pruning Adversarially Robust Neural Networks, Vikash Sehwal, **Shiqi Wang**, Prateek Mittal, Suman Jana, *Conference on Neural Information Processing Systems (NeurIPS)*, 2020.
- [Sec] On Training Robust PDF Malware Classifiers, Yizheng Chen, **Shiqi Wang**, Dongdong She, Suman Jana, *USENIX Security Symposium (Usenix Security)*, 2020.
- [AdvML] Enhancing Gradient-based Attacks with Symbolic Intervals, **Shiqi Wang**, Yizheng Chen, Ahmed Abdou, Suman Jana, *ICML 2019 Workshop on Security and Privacy of Machine Learning (SPML)*, 2019.
- [VNN] Bringing engineering rigor to Deep Learning, Kexin Pei, **Shiqi Wang**, Yuchi Tian, Justin Whitehouse, Carl Vondrick, Yinzhi Cao, Baishakhi Ray, Suman Jana, Junfeng Yang, *ACM SIGOPS Operating Systems*, 2019
- [VNN] Efficient Formal Safety Analysis of Neural Networks, **Shiqi Wang**, Kexin Pei, Justin Whitehouse, Junfeng Yang, Suman Jana, *Conference on Neural Information Processing Systems (NeurIPS)*, 2018.
- [VNN] Formal Security Analysis of Neural Networks using Symbolic Intervals, **Shiqi Wang**, Kexin Pei, Justin Whitehouse, Junfeng Yang, Suman Jana, *USENIX Security Symposium (USENIX Security)*, 2018.
- [Sec] Contextlot: Towards Providing Contextual Integrity to Applied IoT platforms, Yunhan Jack Jia, Qi Alfred Chen, **Shiqi Wang**, Amir Rahmati, Earlene Fernandes, Z. Morley Mao, Atul Prakash, *Network and Distributed System Security Symposium (NDSS)*, 2017.
- [Sec] Defense against impersonating attackers: An efficient RFID mutual authentication protocol based on standard, **Shiqi Wang**, Linsen Li, Gaosheng Chen, Tao Chen, Zeming Wang, *International Conference on Information and Communication Systems (ICICS)*, 2017.
- [Sec] Improved Group Management Protocol of RFID password Method, Tao Chen, Linsen Li, **Shiqi Wang**, Gaosheng Chen, Zeming Wang, *International Conference on Internet of Things, Data and Cloud Computing (ICC)*, 2017.

## ARXIV PREPRINTS (\* INDICATES EQUAL CONTRIBUTION)

---

- [AdvML] A Branch and Bound Framework for Stronger Adversarial Attacks of ReLU Networks, Huan Zhang\*, **Shiqi Wang\***, Kaidi Xu\*, Yihan Wang, Suman Jana, Cho-Jui Hsieh, J. Zico Kolter.
- [AdvML] Towards Understanding Fast Adversarial Training, Bai Li, **Shiqi Wang**, Suman Jana, Lawrence Carin.
- [VNN] MixTrain: Scalable Training of Verifiably Robust Neural Networks, **Shiqi Wang**, Yizheng Chen, Ahmed Abdou, Suman Jana.
- [Prune] Towards Practical Lottery Ticket Hypothesis for Adversarial Training, Bai Li\*, **Shiqi Wang\***, Yunhan Jia, Yantao Lu, Zhenyu Zhong, Lawrence Carin, Suman Jana.
- [Prune] Towards Compact and Robust Deep Neural Networks, Vikash Sehwal\*, **Shiqi Wang\***, Prateek Mittal, Suman Jana.

## PATENTS

---

- Computation Aware Adaptive Pruning for Autoencoders, **Shiqi Wang**, Hamid Vaezi Joze (under review).
- Book management method based on color rectangular code and color rectangular code label, Linsen Li, **Shiqi Wang**, Junhua Tang, Yue Wu, Jianhua Li (CN106919966A).

## OPEN-SOURCE TOOLS

---

- [\$\alpha, \beta\$ -CROWN, the strongest neural network verification toolbox so far](#) [Github link]

I am one of the main contributors to  $\alpha, \beta$ -CROWN, a complete neural network verifier that won the highest total score in 2nd International Verification of Neural Network Competition (VNN-COMP 2021). It can provide rigorous guarantees on the predictions of neural networks under safety and security considerations. One main contribution to its strength comes from the novel algorithms I proposed in optimizing approximated bounds of the models fully supported and accelerated on GPUs, allowing it to outperform traditional linear solver-based algorithms by over 3 orders of magnitude. Overall,  $\alpha, \beta$ -CROWN is written in Python and PyTorch which has over 10,000 lines of code.

- [ReluVal, one of the earliest complete neural network verification toolboxes in 2018](#) [Github link]

I developed ReluVal, one of the earliest complete neural network verifiers. In ReluVal, I mainly proposed and implemented two novel algorithms: (1) symbolic interval analysis for approximating neural network predictions under given input domains and (2) input neuron splits to refine approximations and guarantee completeness. The insights of ReluVal arouse wide attention (**ReluVal paper has >240 citations**) in the neural network verification community and are still widely used in many recent verifiers. I wrote ReluVal in C which has over 5,000 lines of code.

- [Neurify, a follow-up complete neural network verification toolbox based on ReluVal](#) [Github link]

I developed Neurify which is a follow-up complete neural network verifier based on ReluVal. It can support more general safety-/security-properties and modern-size neural networks. Instead of using input neuron splits to refine, I proposed and implemented to use hidden neuron splits with a linear solver. This insight is widely used in most recent state-of-the-art verifiers (**Neurify paper has >200 citations**). The core part of Neurify has over 7,000 lines of C code.

- [Symbolic Interval Analysis Library, one of the earliest fast linear approximations for neural network verification](#) [Github link]

I implemented my proposed symbolic interval analysis into a general library, having novel algorithms to approximate neural network predictions for verification based on linear relaxation. The library is written in Python and PyTorch (>1,000 lines of code) enjoying the GPU acceleration. The library is used in various tasks and many of my research papers like certifiable training for vision models (MixTrain; Hydra), robust PDF malware classifiers (verifiably robust PDF classifiers), and adversarial attacks (interval attack).

## COMPETITION AND AWARDS

---

(\* indicates equal contribution)

- **Best Paper Award Runner-Up:** Learning Security Classifiers with Verified Global Robustness Properties, Yizheng Chen, **Shiqi Wang**, Yue Qin, Xiaojing Liao, Suman Jana, and David Wagner, Conference on Computer and Communications Security (CCS), 2021.
- $\alpha, \beta$ -CROWN: Huan Zhang\*, Kaidi Xu\*, **Shiqi Wang\***, Zhouxing Shi, Yihan Wang, Xue Lin, Suman Jana, Cho-Jui Hsieh, J. Zico Kolter, *The global winner of 2nd International Verification of Neural Networks Competition (VNN-Comp 2021)*, 2021.
- **Interval Attack:** **Shiqi Wang**, Yizheng Chen, Ahmed Abdou, Suman Jana, rank #1 white-box adversarial attack on *MadryLab MNIST Adversarial Examples Challenge* in 2019 and rank #7 in 2021.

## TEACHING EXPERIENCE

---

- **Teaching Assistant for COMS W4181: Security 1 (Fall 2019)**, Columbia University, Instructor: Prof. Suman Jana.

## INVITED TALKS AND GUEST LECTURES

---

- **AAAI 2022 3-hr Tutorial:** "Formal Verification of Deep Neural Networks: Theory and Practice" by Huan Zhang, Kaidi Xu, **Shiqi Wang**, Cho-Jui Hsieh, February 23rd 2022.
- **Guest Lecture:** Columbia University, title "Efficient Formal Safety Analysis of Neural Networks", E6998 Robustness and Security in ML Systems, instructor: Prof. Junfeng Yang, March 9th 2021.
- **Guest Lecture:** University of Nebraska at Lincoln, title "Efficient Formal Safety Analysis of Neural Networks & MixTrain: Scalable Training of Verifiably Robust Neural Networks", CSCE-990-Fall: Deep Learning and Assured Autonomy Analysis, instructor: Prof. Tran Hoang Dung, October 10th 2020.
- **Guest Lecture:** Columbia University, title "Efficient Formal Safety Analysis of Neural Networks", E6998 Robustness and Security in ML Systems, instructor: Prof. Junfeng Yang, March 24th 2020.
- **Invited Speaker:** SRI, title "Efficient Formal Safety Analysis of Neural Networks", host: Natarajan Shankar, June 6th 2019.
- **Invited Speaker:** Baidu X-Lab, title "Efficient Formal Safety Analysis of Neural Networks", host: Zhenyu Zhong, May 23rd 2019.

- **Guest Lecture:** Columbia University, title "Formal Security Analysis of Neural Networks using Symbolic Intervals", E6998 Robustness and Security in ML Systems, instructor: Prof. Junfeng Yang, April 8th 2019.
- **Guest Lecture:** Columbia University, title "Formal Security Analysis of Neural Networks using Symbolic Intervals", E6998 Robustness and Security in ML Systems, instructor: Prof. Junfeng Yang, April 9th 2018.

## ACADEMIA SERVICES

---

- **Conference PC/reviewer:** NeurIPS 2020, NeurIPS 2021, ICML 2021, AISEC 2021, ICLR 2022, AAAI 2022.
- **Journal reviewer:** IEEE Transactions on Neural Networks and Learning Systems.
- **Workshop Organizer:** ATVA 2021 Workshop on Security and Reliability of Machine Learning (SRML).