SHITONG ZHU

Redmond, WA (98052) ⋄ **zhu@shitong.me** ⋄ **shitong.me**

RESEARCH INTERESTS

I am broadly interested in (1) AI for security/privacy: designing SoTA ML/DL models tailored to empower non-ML security/privacy-enhancing technologies (e.g., program/traffic analysis) to scale with high efficacy; and (2) AI security: studying robustness (from both offensive and defensive lens) of SoTA/emerging (e.g., video/Web) ML models. I have published frequently at top-tier venues across security/privacy/ML/Web/networking communities, as a seasoned member in them by serving on PCs of prestigious conferences (e.g., IEEE S&P/INFOCOM).

EDUCATION

University of California, Riverside Sep 2016 - Dec 2021
Ph.D. in **Computer Science** - CGPA: 3.67/4.00 *Riverside, CA*
Advisors: Zhiyun Qian & Srikanth V. Krishnamurthy
Dissertation: Understanding and Taming Adversarial Actions Against Internet Content Blockers

Chongqing University of Posts and Telecommunications Sep 2012 - Jun 2016
BEng. in Telecoms Engineering (**with Honors**) - CGPA: 3.73/4.00 - Top 6% *Chongqing, China*

PROFESSIONAL SERVICES

TPC Member	*IEEE S&P 2024, EAI SecureComm 2023, IEEE INFOCOM 2023*
Reviewer	*IEEE TDSC, ACM CSCW 2022, ACM IMWUT 2022, Springer Machine Learning*

PUBLICATIONS & PRE-PRINTS

* indicates equal contributions. Click paper titles to download PDFs.
Top-tier publications: S&P×2, CCS×1, NDSS×2, NeurIPS×1, CoNEXT×1, WWW×1, ECCV×1

Eluding ML-based Adblockers With Actionable Adversarial Examples
Shitong Zhu, Zhongjie Wang, Xun Chen, Shasha Li, Keyu Man, Umar Iqbal, Zhiyun Qian, Kevin Chan, Srikanth Krishnamurthy, Zubair Shafiq,, Yu Hao, Guoren Li, Zheng Zhang, Xiaochen Zou
Annual Computer Security Applications Conference (**ACSAC '21**)

You Do (Not) Belong Here: Detecting DPI Evasion Attacks with Context Learning
Shitong Zhu, Shasha Li, Zhongjie Wang, Xun Chen, Zhiyun Qian, Srikanth V. Krishnamurthy, Kevin S. Chan, Ananthram Swami
Conference on emerging Networking EXperiments and Technologies (**CoNEXT '20**)

ShadowBlock: A Lightweight and Stealthy Adblocking Browser
Shitong Zhu, Umar Iqbal, Zhongjie Wang, Zhiyun Qian, Zubair Shafiq, and Weiteng Chen
The Web Conference (**WWW '19**)

Measuring and Disrupting Anti-Adblockers Using Differential Execution Analysis
Shitong Zhu, Xunchao Hu, Zhiyun Qian, Zubair Shafiq, and Heng Yin
Network & Distributed System Security Symposium (**NDSS '18**)

SyzDescribe: Principled, Automated, Static Generation of Syscall Descriptions for Kernel Drivers
Yu Hao, Guoren Li, Xiaochen Zou, Weiteng Chen, **Shitong Zhu**, Zhiyun Qian, Ardalan Amiri Sani
IEEE Symposium on Security & Privacy (**S&P '23**)

Adversarial Attacks on Black Box Video Classifiers: Leveraging the Power of Geometric Transformations
Shasha Li*, Abhishek Aich*, **Shitong Zhu**, Salman Asif, Chengyu Song, Amit Roy-Chowdhury, Srikanth Krishnamurthy
Advances in Neural Information Processing Systems (**NeurIPS '21**)

Themis: Ambiguity-Aware Network Intrusion Detection based on Symbolic Model Comparison
Zhongjie Wang, **Shitong Zhu**, Keyu Man, Pengxiong Zhu, Yu Hao, Zhiyun Qian, Srikanth V. Krishnamurthy, Tom La Porta, Michael J. De Lucia
ACM Conference on Computer and Communications Security (**CCS '21**)

Connecting the Dots: Detecting Adversarial Perturbations Using Context Inconsistency
Shasha Li, **Shitong Zhu**, Sudipta Paul, Amit Roy-chowdhury, Chengyu Song, Srikanth V. Krishnamurthy, Ananthram Swami, Kevin S Chan
European Conference on Computer Vision (**ECCV '20**)

AdGraph: A Graph-Based Approach to Ad and Tracker Blocking
Umar Iqbal, Peter Snyder, **Shitong Zhu**, Benjamin Livshits, Zhiyun Qian and Zubair Shafiq
IEEE Symposium on Security & Privacy (**S&P '20**)

SymTCP: Eluding Stateful Deep Packet Inspection with Automated Discrepancy Discovery
Zhongjie Wang, **Shitong Zhu**, Yue Cao, Zhiyun Qian, Chengyu Song, Srikanth V. Krishnamurthy, Tracy D. Braun and Kevin S. Chan
Network & Distributed System Security Symposium (**NDSS '20**)

Generating Practical Adversarial Network Traffic Flows Using NIDSGAN
Bolor-Erdene Zolbayar, Ryan Sheatsley, Patrick McDaniel, Michael J Weisman, Sencun Zhu, **Shitong Zhu**, Srikanth Krishnamurthy
arXiv preprint arXiv:2203.06694

Before 2016 (undergraduate work)

Source-location Privacy Protection Strategy via Pseudo Normal Distribution-based Phantom Routing in WSNs
Jun Huang, Meisong Sun, **Shitong Zhu**, Yi Sun, Cong-cong Xing, and Qiang Duan
Annual ACM Symposium on Applied Computing (**SAC '15**)
On Selecting Composite Network-Cloud Services: A Quality-of-Service Based Approach
Minkailu Mohamed Jalloh, **Shitong Zhu**, Fang Fang, and Jun Huang
Conference on Research in Adaptive and Convergent Systems (**RACS '15**)
A Defense Model of Reactive Worms Based on Dynamic Time
Haokun Tang, **Shitong Zhu**, Jun Huang, and Hong Liu
Journal of Software, 2778-2788, Sep 2014

WORK EXPERIENCE

Research Scientist @ Meta Dec 2021 - Present
Privacy AI *Bellevue, WA*

- Developing various SoTA ML models[1] to detect and mitigate privacy risks in various artifacts (e.g., code changes)
- Training/enhancing LLMs to improve privacy-aware code authoring[2] and project reviewing

Summer Research Intern @ IBM Research Jun 2021 - Sep 2021
Thomas J. Watson Research Center (Host: Supriyo Chakraborty) *Remote*

- Model interpretability/explainablity

[1] Detecting privacy-sensitive code changes: Paper
[2] AI-assisted code authoring: Paper / Meta Post

- Deep learning for program analysis

Research Intern @ Samsung Research America Jan-Mar 2020/Jun-Sep 2019
KNOX Security Team (Host: Xun Chen) *Mountain View, CA/Remote*

- ML-based cyber-security infrastructure
- Adversarial machine learning in restricted domains

SELECTED PROJECTS

Privacy Understanding Using Large Language and Graph Models **[WIP]**

- LM/graph-based modeling over Meta-internal artifacts; achieved SoTA performance and integrated for various downstream privacy-critical tasks to detect/contextualize risks
- Trained (from scratch) and aligned foundational LLM models to privacy/security-oriented applications

Explaining Graph-based Code Models **[WIP, to be submitted]**

- Non-empirical gradient-based interpretation strategies with graph-structural guidance
- Achieved significantly improved attribution accuracy in multiple metrics compared to current SoTAs

Context-aware Symbolic Execution **[WIP, to be submitted]**

- Learning-based execution strategy that speeds up symbolic execution engines via intelligent decision making

ML-based Solution for Detecting DPI Evasion Attacks **[CoNEXT '20]**

- First ML-based solution that only relies on clean traffic traces for detecting and localizing 73 state-of-the-art evasion attacks against Deep Packet Inspection (DPI) systems
- Achieved a ROC-AUC of 0.963, an EER of 0.061 in detection, and an accuracy of 96.4% in localization, by constructing semantic representations for network traffic with *packet context* considered

Detecting Adversarial Perturbations Using Context Consistency **[ECCV '20]**

- Defined, extracted and formulated context information from clean images to detect adversarially perturbed samples against state-of-the-art object detectors
- Achieved a ROC-AUC of over 0.95, a >20% improvement over state-of-the-art context-agnostic methods

Adversarial Examples in Web Domain **[ACSAC '21]**

- First effort in generating *actionable* (i.e. non-disruptive and concretizable) adversarial examples in web domain against non-perceptual ML-based adblockers
- Achieved a success rate of ≈60%, surpassing the state-of-the-art attack by a significant margin of 84.3%

ML-based Automatic and Effective Adblocking **[S&P '20]**

- Leveraged multiple layers of the web stack (HTML/HTTP/JavaScript) to train a classifier for blocking ads/trackers
- Replicated state-of-the-art filter lists with high accuracy (97.7%)
- Enhanced filter lists by automatically correcting their errors

Evading & Defending DPI Systems Using Symbolic Execution **[NDSS '20 & CCS '21]**

- Used symbolic execution to guide the generation of insertion and evasion packets at the TCP level for automated testing against DPI middleboxes
- Discovered over 20 strategies to elude DPI middleboxses that target Zeek (formerly Bro), Snort and GFW within an hour

Measuring & Defending Anti-adblocking **[NDSS '18 & WWW '19]**

- First large-scale measurements of anti-blocking in the wild, revealing >3x prevalence than prior work
- Built invisible adblocker that evades current generation of anti-adblockers with 100% of success rate in manual evaluation

SKILLS

Languages/Tools Python, C/C++, JavaScript / PyTorch, KLEE, Angr, Chromium

INVITED TALKS

Eluding ML-based Adblockers With Actionable Adversarial Examples Online
Cyber Security Collaborative Research Alliance (Webinar) *Oct 2021*

You Do (Not) Belong Here: Detecting DPI Evasion Attacks with Context Learning Online
Cyber Security Collaborative Research Alliance (Webinar) *Dec 2020*

Adblocking: A Slient Online Arms Race Xi'an, China
XJTU InForSec Event *Dec 2019*

Arms Race between Adblockers and Anti-adblockers San Francisco, CA
Mozilla Security Research Summit *May 2019*

Detection and Circumvention of Ad-Block Detectors Barcelona, Spain
Data Transparency Lab Conference *Dec 2017*

HONORS & AWARDS

NYU CSAW Applied Research Competition 3rd Place	US-Canada, 2020
Dissertation Year Program (DYP) Award	UC Riverside CSE, 2020-2021
Dean's Distinguished Fellowship (full scholarship)	UC Riverside CSE, 2016-2017
2nd Class University Scholarship	CUPT, 2015-2016
National 2nd Prize @ National Olympiad in Informatics	China Computer Federation, 2009

REFERENCES

Zhiyun Qian Riverside, CA
Everett and Imogene Ross (Full) Professor *Co-advisor*

· Department of Computer Science and Engineering @ University of California, Riverside

Contact: zhiyunq@cs.ucr.edu

Srikanth V. Krishnamurthy Riverside, CA
Full Professor, AAAS/IEEE Fellow *Co-advisor*

· Department of Computer Science and Engineering @ University of California, Riverside

Contact: krish@cs.ucr.edu

Xun Chen Mountain View, CA
Director *Intern Mentor*

· Knox Advanced Research and Development @ Samsung Research America

Contact: xun.chen@samsung.com