

SHITONG ZHU

Seattle, WA (98109)

zhu@shitong.me \diamond shitong.me

EDUCATION

- University of California, Riverside** Sep 2016 - Nov 2021 (expected)
Ph.D. in Computer Science - CGPA: 3.67/4.00 *Riverside, CA*
Advisors: Zhiyun Qian & Srikanth V. Krishnamurthy
- 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*

PUBLICATIONS

* indicates equal contributions.

- [C11] **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
To appear in *Advances in Neural Information Processing Systems (NeurIPS '21)*
- [C10] **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
To appear in *Annual Computer Security Applications Conference (ACSAC '21)*
- [C9] **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
To appear in *ACM Conference on Computer and Communications Security (CCS '21)*
- [C8] **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)
- [C7] **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)
- [C6] **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)
- [C5] **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)
- [C4] **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)

[C3] **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)

Before 2016 (undergraduate work)

[C2] **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)

[C1] **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)

[J2] **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

[J1] **Propagation of Active Worms in P2P Networks: Modeling and Analysis**

Haokun Tang, Yukui Lu, **Shitong Zhu**, Jun Huang
Journal of Computers, 2514-2524, Sep 2014

WORK EXPERIENCE

Summer Research Intern @ IBM Research

Thomas J. Watson Research Center (Host: Supriyo Chakraborty)

Jun 2021 - Sep 2021

Remote

- Model interpretability/explainability
- Deep learning for program analysis

PhD Machine Learning Intern @ Facebook

Business Integrity Team (Host: Abdel Baligh)

Jul 2020 - Sep 2020

Remote

- Designed and implemented ML models for detecting bad advertisers through effective and efficient neural web modeling, with a blend of graphical and NLP models
- Improved classification accuracies for different applications by a significant margin (>30%)

Research Intern @ Samsung Research America

KNOX Security Team (Host: Xun Chen)

Jan 2020 - Mar 2020

Remote

- ML-based cyber-security infrastructure

Research Intern @ Samsung Research America

KNOX Security Team (Host: Xun Chen)

Jun 2019 - Sep 2019

Mountain View, CA

- Adversarial machine learning in restricted domains

Graduate Student Researcher @ UCR CSE

UCR SecLab

Sep 2016 - Present

Riverside, CA

- Pursued research on computer security and fulfilled the entire cycle of projects
- Published/co-authored papers accepted by or submitted to top-tier venues

Consulting Intern @ Deloitte TTL

ERS - Technology Risk

Jan 2016 - Mar 2016

Shenzhen, China

- Advised tech organizations to avoid being victim of a security breach through big data analytics
- Assisted team of a major bank in designing reliable identity and access management framework

Software Engineering Intern @ Douban Inc.

Research & Development Center

Jul 2015 - Sep 2015

Beijing, China

- Implemented new functionalities on server side, conducted web development in Python
- Designed and tuned Spark/Hadoop scripts processing data on distributed clusters
- Prototyped, implemented and tailored algorithmic details of "Douban NewMov Chart"

SELECTED PROJECTS

Explaining Graph-based Code Models

[WIP]

- Non-empirical gradient-based interpretation strategies
- Achieved significantly improved attribution accuracies (in faithfulness tests etc.)

Semantic-aware Symbolic Execution

[WIP]

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

ML-based Solution for Detecting DPI Evasion Attacks

[C8]

- 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

[C7]

- 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 in most cases, a >20% improvement over state-of-the-art context-agnostic methods

Adversarial Examples in Web Domain

[C10]

- 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 $\approx 60\%$, surpassing the state-of-the-art attack by a significant margin of 84.3%

ML-based Automatic and Effective Adblocking

[C6]

- 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 Deep Packet Inspection Systems Using Symbolic Execution

[C5]

- 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 middleboxes that target Zeek (formerly Bro), Snort and GFW within an hour

Stealthy Adblcoking

[C4]

- Built invisible adblocker that evades current generation of anti-adblockers with 100% of success rate in manual evaluation
- Replicated 98.2% of ad coverage achieved by popular adblocking extensions, while causing minor visual breakage on less than 0.6% of Alexa top 1K websites. In the meantime, page loads are sped up by over 5% on average

SKILLS

Languages Python, C/C++, JavaScript, MATLAB

Others PyTorch, Chromium, Selenium/Puppeteer, NodeJS, Hadoop/Spark, Git

PROFESSIONAL SERVICES

Reviewer

PeerJ Computer Science

Sub-reviewer *ISOC NDSS 2019/2020, ACM CCS 2019, IEEE S&P 2019/2020, ICML 2021, NeurIPS 2021, Journal of Systems and Software*

Artifact Evaluation Committee

USENIX Security 2022

INVITED TALKS

Eluding ML-based Adblockers With Actionable Adversarial Examples <i>Cyber Security Collaborative Research Alliance (Webinar)</i>	Online Oct 2021
You Do (Not) Belong Here: Detecting DPI Evasion Attacks with Context Learning <i>Cyber Security Collaborative Research Alliance (Webinar)</i>	Online Dec 2020
Adblocking: A Slient Online Arms Race <i>XJTU InForSec Event</i>	Xi'an, China Dec 2019
Arms Race between Adblockers and Anti-adblockers <i>Mozilla Security Research Summit</i>	San Francisco, CA May 2019
Detection and Circumvention of Ad-Block Detectors <i>Data Transparency Lab Conference</i>	Barcelona, Spain Dec 2017

HONORS & AWARDS

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 <i>Everett and Imogene Ross Associate Professor</i> · Department of Computer Science and Engineering · University of California, Riverside Contact: zhiyunq@cs.ucr.edu	Riverside, CA <i>Co-advisor</i>
Srikanth V. Krishnamurthy <i>Professor, IEEE Fellow</i> · Department of Computer Science and Engineering · University of California, Riverside Contact: krish@cs.ucr.edu	Riverside, CA <i>Co-advisor</i>
Xun Chen <i>Director</i> · Knox Advanced Research and Development · Samsung Research America Contact: xun.chen@samsung.com	Mountain View, CA <i>Intern Mentor</i>