National Taiwan University
No. 1, Sec. 4, Roosevelt Rd., Taipei 10617, Taiwan

Shang-Tse Chen

stchen@csie.ntu.edu.tw
http://www.csie.ntu.edu.tw/~stchen

Research Interests

Machine Learning, Security, Algorithmic Game Theory, and Data Mining

Current Position

Feb. 2020 — present **Assistant Professor**
Department of Computer Science and Information Engineering
National Taiwan University, Taipei, Taiwan

Education

2013 — 2019 **Ph.D. in Computer Science**
Georgia Institute of Technology, Atlanta, GA
Committee: Polo Chau (advisor), Nina Balcan (co-advisor), Wenke Lee, Le Song, Kevin Roundy, and Cory Cornelius

2006 — 2010 **B.S. in Computer Science Information Engineering**
National Taiwan University, Taipei, Taiwan

Selected Honors and Awards

2020 ACM Trans. Interactive Intelligent Systems (TiiS) 2018 Best Paper, Honorable Mention

2018 — 2019 IBM PhD Fellowship
For my Ph.D. research on "AI-infused Security: Robust Defense by Bridging Theory and Practice"

2018 KDD'18 Audience Appreciation Award, Runner Up
For "Shield: Fast, Practical Defense and Vaccination for Deep Learning using JPEG Compression"

2016 Symantec Fellowship Runner-Up

2016 KDD'16 Best Student Paper Award, Runner-Up
F.or "Firebird: Predicting Fire Risk and Prioritizing Fire Inspections in Atlanta"

2010 National Science Council Research Creativity Award
For my undergraduate research on "Link Prediction in Heterogeneous Networks"

2009 National Science Council Undergraduate Research Fellowship
A grant that recognizes a small number of outstanding undergraduate researchers

2009 KDD Cup 2009 3rd Prize (slow track)
Out of 400+ submissions
KDD CUP is the most prestigious data mining contest

Industry Research Experience

Summer 2018 **Intel Labs**, Hillsboro, OR
Graduate Machine Learning Security Intern
Mentor: Cory Cornelius, Jason Martin
Explored regularization techniques as defense against adversarial attack.

Summer 2017 **Intel Labs**, Hillsboro, OR
Graduate Security Intern
Mentor: Cory Cornelius, Jason Martin
Developed *ShapeShifter*, the **1st physical adversarial attack** that fools Faster R-CNN object detectors

Summer 2016 **Symantec Research Labs**, Culver City, CA
Research Engineer Intern
Mentor: Kevin A. Roundy
Developed **patent-pending** *Virtual Product*, a novel framework for enterprise cyber threat detection.

Summer 2015 **Pindrop Security**, Atlanta, GA
Research Intern
Mentor: Raj Bandyopadhyay
Improved phone fraud detection system significantly by 10 absolute percentage.

Academic Research Experience

2013 — 2019 **Georgia Institute of Technology**, Atlanta, GA
Graduate Research Assistant, School of Computational Science and Engineering
Advisors: Polo Chau and Nina Balcan

2011 — 2013 **Academia Sinica**, Taipei, Taiwan
Graduate Research Assistant, Institute of Information Science
Advisors: Chi-Jen Lu and Hsuan-Tien Lin

2008 — 2010 **National Taiwan University**, Taipei, Taiwan
Undergraduate Research Assistant, Department of Computer Science and Information Engineering
Advisors: Shou-De Lin and Hsuan-Tien Lin

Publications

Refereed Conference Papers

ShapeShifter: Robust Physical Adversarial Attack on Faster R-CNN Object Detector
Shang-Tse Chen, Cory Cornelius, Jason Martin, and Duen Horng (Polo) Chau
European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML-PKDD). Dublin, Ireland. Sept. 2018.
Github | Video

Shield: Fast, Practical Defense and Vaccination for Deep Learning using JPEG Compression
Nilaksh Das, Madhuri Shanbhogue, Shang-Tse Chen, Fred Hohman, Siwei Li, Li Chen, Michael E. Kounavis, and Duen Horng (Polo) Chau

ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD). London, UK. Aug. 2018.
Github | Video | ***Audience Appreciation Award, Runner-Up***

Diversified Strategies for Mitigating Adversarial Attacks in Multiagent Systems
Maria-Florina Balcan, Avrim Blum, and Shang-Tse Chen (alphabetic order)
International Conference on Autonomous Agents and Multiagent Systems (AAMAS). Stockholm, Sweden. July 2018.

Predicting Cyber Threats with Virtual Security Products
Shang-Tse Chen, Yufei Han, Duen Horng (Polo) Chau, Christopher Gates, Michael Hart, and Kevin Roundy
Annual Computer Security Applications Conference (ACSAC). Orlando, FL. Dec. 2017.
Patented

Firebird: Predicting Fire Risk and Prioritizing Fire Inspections in Atlanta
Michael Madaio, Shang-Tse Chen, Oliver Haimson, Wenwen Zhang, Xiang Cheng, Matthew Hinds-Aldrich, Duen Horng (Polo) Chau, and Bistra Dilkina.
ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD). San Francisco, CA. Aug. 2016.
Site | Github | Video | ***Best Student Paper Award, Runner-Up***

Communication Efficient Distributed Agnostic Boosting
Shang-Tse Chen, Maria-Florina Balcan, and Duen Horng (Polo) Chau
International Conference on Artificial Intelligence and Statistics (AISTATS). Cadiz, Spain. May 2016.

Boosting with Online Binary Learners for the Multiclass Bandit Problem
Shang-Tse Chen, Hsuan-Tien Lin, and Chi-Jen Lu
International Conference on Machine Learning (ICML). Beijing, China. June 2014.

An Online Boosting Algorithm with Theoretical Justifications
Shang-Tse Chen, Hsuan-Tien Lin, and Chi-Jen Lu
International Conference on Machine Learning (ICML). Edinburgh, Scotland. June 2012.
Code

Journal Articles and Book Chapters

Chronodes: Interactive Multi-focus Exploration of Event Sequences
Peter J. Polack, Shang-Tse Chen, Minsuk Kahng, Kaya De Barbaro, Rahul Basole, Moushumi Sharmin, Duen Horng (Polo) Chau
ACM Transactions on Interactive Intelligent Systems (TiiS) Special Issue on Interactive Visual Analysis of Human and Crowd Behaviors, 2018.
Video | ***Best Paper, Honorable Mention***

Exploratory Visual Analytics of Mobile Health Data: Sensemaking Challenges and Opportunities
Peter J. Polack, Moushumi Sharmin, Kaya de Barbaro, Minsuk Kahng, Shang-Tse Chen, and Duen Horng (Polo) Chau
Mobile Health: Sensors, Analytic Methods, and Applications. Springer, 2017.

Constructing, Analyzing and Visualizing Social Networks: Exemplified by the Academia Social Network in Taiwan
Cheng-Te Li, Chun-Min Chang, Chien-Pang Liu, Shang-Tse Chen, and Shou-De Lin
Journal of Librarianship and Information Studies, 67:72-87, 2008.

Refereed Workshop, Poster, and Demo papers

Extracting Knowledge For Adversarial Detection and Defense in Deep Learning
Scott Freitas, Shang-Tse Chen, and Duen Horng (Polo) Chau
KDD Workshop on Learning and Mining for Cybersecurity (LEMINCS). Anchorage, AK. Aug. 2019.

Talk Proposal: Towards the Realistic Evaluation of Evasion Attacks using CARLA
Cory Cornelius, Shang-Tse Chen, Jason Martin, and Duen Horng (Polo) Chau
Dependable and Secure Machine Learning (DSML). Portland, OR, June 2019.

ADAGIO: Interactive Experimentation with Adversarial Attack and Defense for Audio
Nilaksh Das, Madhuri Shanbhogue, Shang-Tse Chen, Li Chen, Michael E. Kounavis, and Duen Horng (Polo) Chau
European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML-PKDD) (demo). Dublin, Ireland. Sept. 2018.
Video

Physical Adversarial Attack on Object Detectors
Shang-Tse Chen, Cory Cornelius, Jason Martin, and Duen Horng (Polo) Chau
ACM KDD Project Showcase. London, UK. Aug. 2018.

Compression to the Rescue: Defending from Adversarial Attacks Across Modalities
Nilaksh Das, Madhuri Shanbhogue, Shang-Tse Chen, Fred Hohman, Siwei Li, Li Chen, Michael E. Kounavis, and Duen Horng (Polo) Chau
ACM KDD Project Showcase. London, UK. Aug. 2018.

TimeStitch: Interactive Multi-focus Cohort Discovery and Comparison
Peter J. Polack, Shang-Tse Chen, Minsuk Kahng, Moushumi Sharmin, and Duen Horng (Polo) Chau
IEEE Conference on Visual Analytics Science and Technology (VAST'15) (Poster). Chicago, IL, Oct. 2015.
Video

Spotting Suspicious Reviews via (Quasi-)clique Extraction
Paras Jain, Shang-Tse Chen, Mozhgan Azimpourkivi, Duen Horng (Polo) Chau, and Bogdan Carbunar
IEEE Symposium on Security and Privacy (Oakland) (poster). SAN JOSE, CA. May 2015.

An Ensemble of Three Classifiers for KDD Cup 2009: Expanded Linear Model, Heterogeneous Boosting, and Selective Naive Bayes
Hung-Yi Lo, Kai-Wei Chang, Shang-Tse Chen, Tsung-Hsien Chiang, Chun-Sung Ferng, Cho-Jui Hsieh, Yi-Kuang Ko, Tsung-Ting Kuo, Hung-Che Lai, Ken-Yi Lin, Chia-Hsuan Wang, Hsiang-Fu Yu, Chih-Jen Lin, Hsuan-Tien Lin, and Shou-de Lin
JMLR Workshop and Conference Proceedings, V.7, 57-64, 2009.
3rd Place of the KDD Cup'09 Slow Track

Technical Reports

UnMask: Adversarial Detection and Defense Through Robust Feature Alignment
Scott Freitas, Shang-Tse Chen, Zijie Wang, and Duen Horng (Plo) Chau
arXiv:2002.09576, Feb. 2020.

Keeping the Bad Guys Out: Protecting and Vaccinating Deep Learning with JPEG Compression
Nilaksh Das, Madhuri Shanbhogue, Shang-Tse Chen, Fred Hohman, Li Chen, Michael E. Kounavis, and Duen Horng (Polo) Chau
arXiv:1705.02900, May 2017.

An Ensemble Ranking Solution to the Yahoo! Learning to Rank Challenge
Ming-Feng Tsai, Shang-Tse Chen, Yao-Nan Chen, Chun-Sung Ferng, Chia-Hsuan Wang, Tzay-Yeu Wen, and Hsuan-Tien Lin
National Taiwan University, Technical Report, Sept. 2010.

Teaching Experience

Courses

Spring 2020 **Security and Privacy of Machine Learning**
National Taiwan University, Taipei, Taiwan

Teaching Assistant

Fall 2015 **CSE-6040: Computing for Data Analytics**
Instructor: Richard Vuduc
Georgia Institute of Technology, Atlanta, GA
Introductory data analytics course for MS in Analytics students. (37 students)

Spring 2015 **CS-7545: Machine Learning Theory**
Instructor: Santosh Vempala
Georgia Institute of Technology, Atlanta, GA
Advanced ML theory course primarily taken by PhD students. (26 students)

Guest Lectures

Firebird: Predicting Fire Risk and Prioritizing Fire Inspections in Atlanta
Georgia Institute of Technology, Atlanta, GA
CSE-6242 Georgia Tech. Instructor: Polo Chau.
Fall 2018 245 students (32 undergrads)
Spring 2018 200 students (45 undergrads)
Fall 2017 273 students (40 undergrads)

Fall 2015 **Data analysis and visualization**
Georgia Institute of Technology, Atlanta, GA
CSE-6040 Georgia Tech. Instructor: Richard Vuduc. 37 students

Spring 2015 **Practical ML for Big Data**
Georgia Institute of Technology, Atlanta, GA
CS-7545 Georgia Tech. Instructor: Santosh Vempala. 26 students

Course Design

Fall 2016 **Big Data Bootcamp**

Georgia Institute of Technology, Atlanta, GA

Co-led the design of a two-day intense hands-on bootcamp on big data tools, that is now offered yearly to students in the MS in Analytics program (~ 50 students each year).

Course material: http://www.sunlab.org/teaching/cse8803/fall2016/lab/

Mentoring

2016 — 2017 **Madhuri Shanbhogue**

M.S. CS, Georgia Tech

Adversarial Machine Learning

Now: Software Engineer at Facebook

2015 - 2016 **Peter Polack**

M.S. CS, Georgia Tech

Interactive visualization of health sensor data

Now: PhD student at UCLA

2014 — 2017 **Paras Jain**

B.S. CS, Georgia Tech

Fraud Detection on Yelp

Now: PhD student at UC Berkeley

Grants and Funding

2018 **IBM PhD Fellowship**

$95,000 over 2 years, covering full Tuition + $35,000 stipend for 2 years

2017 **SaTC: CORE: Medium: Understanding and Fortifying Machine Learning Based Security Analytics**

NSF CNS 1704701

PI: Polo Chau Co-PIs: Taesoo Kim, Wenke Lee, Le Song

Funded: $1,200,000, 8/1/2017 - 7/31/2021

Co-authored winning proposal, contributing a theory-guided decision-making and defense framework

2016 **Intel Science & Technology Center for Adversary-Resilient Security Analytics (ISTC-ARSA)**

PI: Wenke Lee

Co-PIs: Polo Chau, Taesoo Kim, Le Song

Gift Funding: $1,500,000, 2016 - 2019

Co-authored winning proposal, contributing robust, adaptive algorithms for efficient defenses

Invited Talks

Communication Efficient Distributed Agnostic Boosting

Apr. 2016 HotCSE Seminar, Georgia Tech, Atlanta, GA

Boosting with Online Binary Learners for the Multiclass Bandit Problem

June 2014 Appier Inc., Taipei, Taiwan.

Press

Apr. 2020 "Intel Joins Georgia Tech in DARPA Program to Mitigate Machine Learning Deception Attacks" Business Wire.

Apr. 2020 "Georgia Tech and Intel Awarded Multimillion-Dollar Program to Defend Against Attacks on AI" Georgia Tech CSE.

Oct. 2018 "Study reveals new vulnerability in self-driving cars" Tech HQ.

Sept. 2018 "Erasing Stop Signs: ShapeShifter Shows Self-Driving Cars Can Still Be Manipulated" Georgia Tech, College of Computing.

June 2018 "Georgia Tech Teams up with Intel to Protect Artificial Intelligence from Malicious Attacks Using SHIELD." Georgia Tech, College of Computing.

Apr. 2018 "CSE Ph.D. Students Claim Three Prestigious Fellowships." Georgia Tech, College of Computing.

Professional Activities

Program Committee

International Conference on Machine Learning (**ICML**) 2018 - 2020

Annual Conference on Neural Information Processing Systems (**NeurIPS**) 2017 - 2020

AAAI Conference on Artificial Intelligence (**AAAI**) 2018 - 2020

Uncertainty in Artificial Intelligence (**UAI**) 2015 - 2019

SIAM International Conference on Data Mining (**SDM**) 2019

Asian Conference on Machine Learning (**ACML**) 2017 - 2019

Conference on Technologies and Applications of Artificial Intelligence (**TAAI**) 2014

Deep Learning and Security Workshop @ IEEE S&P (**DLS**) 2019

Reviewer

International Conference on Artificial Intelligence and Statistics (**AISTATS**) 2019

Deep Learning and Security Workshop @ IEEE S&P (**DLS**) 2018

SIAM International Conference on Data Mining (**SDM**) 2016 - 2017

ACM SIGKDD Conference on Knowledge Discovery and Data Mining (**KDD**) 2017

Annual Network and Distributed System Security Symposium (**NDSS**) 2017

USENIX Security Symposium (**USENIX Security**) 2017

ACM Conference on Computer and Communications Security (**CCS**) 2017