Scott Freitas

I work at the intersection of **applied** and **theoretical machine learning**, with a strong application focus on **cybersecurity**. My goal is to develop explainable next-generation defenses to protect systems against adversarial attacks.

At Georgia Tech I work with Polo Chau as a member of the Polo Club of Data Science. During this time, I have co-authored several winning research proposals, including a multi-million dollar DARPA grant.

I have been fortunate to work with amazing engineers and scientists at IBM Research, Amazon, Microsoft Advanced Threat Protection, Microsoft Research, Intel and Naval Air Warfare Center.

My research is generously supported by PhD fellowships from IBM Research, NSF GRFP and Raytheon.

scottfreitas.com
safreita@gatech.edu
Curriculum Vitae (PDF)
Github
Linkedin
Google Scholar

Education

Present — Aug. 2018 **Ph.D. in Machine Learning**
Georgia Institute of Technology, Atlanta, GA
Advisor: Duen Horng (Polo) Chau

May 2018 **M.S. in Computer Science**
Summer 2017 Arizona State University, Tempe, AZ
Advisor: Hanghang Tong, Thesis: "Mining Marked Nodes in Large Graphs"
Overall GPA: 4.00/4.00

May 2017 — Aug. 2015 **B.S. in Computer Science**
Arizona State University, Tempe, AZ
Overall GPA: 3.98/4.00

May 2014 — Aug. 2010 **B.S.E. in Electrical Engineering**
Arizona State University, Tempe, AZ
Overall GPA: 3.64/4.00

Honors and Awards

2021 IBM PhD Fellowship
One of sixteen fellows; awarded for my work in developing next-generation explainable defenses

2021 Nvidia Data Science Teaching Kit
Helped develop one of five Nvidia teaching kits used by educators around the world

2019 Raytheon Research Fellowship
Awarded for my PhD work in adversarial machine learning

2018 — 2023 NSF Graduate Research Fellowship
National Science Foundation recognizes and supports outstanding graduate students in STEM fields

2018 Outstanding Computer Science Masters Student (ASU)
Awarded to single master student demonstrating exemplary performance

2017 Best Demo Award, Runner Up at CIKM '17
For "Rapid Analysis of Network Connectivity"

2017 CIKM Travel Grant
Funding from NSF and SIGWEB to present at CIKM

2016 — 2017 FURI Grant
Undergraduate research grant awarded for work in network connectivity

2016 — 2017 Arizona Graduate Scholar Award
Merit scholarship awarded to select number of master students

2010 — 2014 Provost's Scholarship
Merit scholarship awarded to select number of incoming undergraduate students

Industry Research Experience

Fall 2021 **IBM Research**, Yorktown Heights, NY
Research Intern, Cyber Security Intelligence (CSI) Team
Mentor: Jiyong Jang

Summer 2021 **Amazon**, Seattle, WA
Applied Scientist Intern, Fraud Detection and Risk Transaction (CTPS)
Mentor: Hao Zheng, Yanni Lai
Developed unsupervised and semi-supervised methods to prevent fraudulent transactions across the Amazon marketplace

Summer 2020 **Microsoft**, Seattle, WA
Research Intern, Microsoft ATP + Microsoft Research
Mentor: Karishma Sanghvi, Yuxiao Dong
Developed graph neural network approach to detect malware.

Summer 2019 **Microsoft**, Seattle, WA
Research Intern, Microsoft Advanced Threat Protection (ATP)
Mentor: Andrew Wicker, Joshua Neil
Designed the first framework to model lateral attacks on enterprise networks, enabling IT admins to quantify and mitigate network vulnerability to lateral attacks

Summer 2013 **Naval Air Warfare Center**, Point Mugu, CA
Research Intern, Naval Research Entperprise Internship Program (NREIP)
Mentor: Balaji Iyer
Explored methods of preventing electromagentic interference from coupling into superconducting receivers

Academic Research Experience

Present — Aug. 2018 **Georgia Institute of Technology**, Atlanta, GA
Graduate Research Assistant, School of Computational Science and Engineering
Mentor: Duen Horng (Polo) Chau
Member of the Polo Club of Data Science where we innovate scalable, interactive, and interpretable tools that amplify human's ability to understand and interact with billion-scale data and machine learning models

May 2018 — Summer 2017 **Arizona State University**, Tempe, AZ
Graduate Research Assistant, School of Computing, Informatics, and Decision Systems Engineering
Mentor: Hanghang Tong
Conducted research in graph based connectivity analysis to improve local graph partitioning. Developed web-based prototype for explainable ranking in complex multi-layered networks.

Summer 2017 **Arizona State University**, Tempe, AZ
Summer Research Assistant, School of Computing, Informatics, and Decision Systems Engineering
Mentor: Ross Maciejewski
Developed interactive augmented reality (AR) graph models in the Microsoft Hololens.

May 2017 — Jan. 2016 **Arizona State University**, Tempe, AZ
Undergraduate Research Assistant, School of Computing, Informatics, and Decision Systems Engineering
Mentor: Hanghang Tong
Developed fast graph mining algorithms for network connectivity analysis, and award winning web platform for visualization and analysis.

Publications

Graph Vulnerability and Robustness: A Survey
Scott Freitas, Diyi Yang, Srijan Kumar, Hanghang Tong, Duen Horng (Polo) Chau
arXiv (arXiv). Online, 2021.
PDF BibTeX

EnergyVis: Interactively Tracking and Exploring Energy Consumption for ML Models
Omar Shaikh, Jon Saad-Falcon, Austin P Wright, Nilaksh Das, Scott Freitas, Omar Isaac Asensio, Duen Horng Chau
ACM Conference on Human Factors in Computing Systems (CHI). Online, 2021.
Demo PDF BibTeX

MalNet: A Large-Scale Cybersecurity Image Database of Malicious Software
Scott Freitas, Rahul Duggal, Duen Horng (Polo) Chau
arXiv (arXiv). Online, 2021.
Demo PDF Code BibTeX

A Large-Scale Database for Graph Representation Learning
Scott Freitas, Yuxiao Dong, Joshua Neil, Duen Horng (Polo) Chau
Neural Information Processing Systems Datasets and Benchmarks (NeurIPS). Online, 2021.
Project Demo PDF Blog Code BibTeX

UnMask: Adversarial Detection and Defense Through Robust Feature Alignment
Scott Freitas, Shang-Tse Chen, Zijie J. Wang, Duen Horng (Polo) Chau
IEEE International Conference on Big Data (Big Data). Atlanta, GA, 2020.
Project PDF Blog Code BibTeX

Evaluating Graph Vulnerability and Robustness using TIGER
Scott Freitas, Diyi Yang, Srijan Kumar, Hanghang Tong, Duen Horng (Polo) Chau
ACM International Conference on Information and Knowledge Management (CIKM). Online, 2021.
PDF Blog Code BibTeX Featured in Nvidia Data Science Toolkit

ELF: An Early-Exiting Framework for Long-Tailed Classification
Rahul Duggal, Scott Freitas, Sunny Dhamnani, Duen Horng (Polo) Chau, Jimeng Sun
arXiv (arXiv). Online, 2020.
PDF BibTeX

Argo Lite: Open-Source Interactive Graph Exploration and Visualization in Browsers
Siwei Li, Zhiyan Zhou, Anish Upadhayay, Omar Shaikh, Scott Freitas, Haekyu Park, Zijie J. Wang, Susanta Routray, Matthew Hull, Duen Horng (Polo) Chau
ACM International Conference on Information and Knowledge Management (CIKM). Online, 2020.
Demo PDF Code BibTeX

REST: Robust and Efficient Neural Networks for Sleep Monitoring in the Wild
Rahul Duggal*, Scott Freitas*, Cao Xiao, Duen Horng (Polo) Chau, Jimeng Sun
The Web Conference (WWW). Taipei, Taiwan, 2020.
Project PDF Blog Code BibTeX * Authors contributed equally

D^{2}M: Dynamic Defense and Modeling of Adversarial Movement in Networks
Scott Freitas, Andrew Wicker, Duen Horng (Polo) Chau, Joshua Neil
SIAM International Conference on Data Mining (SDM). Cincinnati, Ohio, 2020.
Project PDF Blog BibTeX

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

Local Partition in Rich Graphs
Scott Freitas, Nan Cao, Yinglong Xia, Duen Horng (Polo) Chau, Hanghang Tong
IEEE International Conference on Big Data (Big Data). Seattle, Washington, 2018.
Project PDF BibTeX

X-Rank: Explainable Ranking in Complex Multi-Layered Networks
Jian Kang*, Scott Freitas*, Haichao Yu, Yinglong Xia, Hanghang Tong
ACM International Conference on Information and Knowledge Management (CIKM). Turin, Italy, 2018.
Project PDF BibTeX * Authors contributed equally

Rapid Analysis of Network Connectivity
Scott Freitas, Hanghang Tong, Nan Cao, Yinglong Xia
ACM International Conference on Information and Knowledge Management (CIKM). Singapore, 2017.
Project PDF Video Code BibTeX Best Demo Paper, Runner up

Talks

Detecting Financial Fraud in Online Marketplaces
August 2021 Amazon

Developing Robust Models, Algorithms, Databases and Tools with Applications to Cybersecurity and Healthcare
May 2021 Georgia Institute of Technology

Exploring Graph Neural Networks for Malware Detection
July 2020 Microsoft Advanced Threat Protection

On the Robustness and Vulnerability of Graphs
April 2020 Georgia Institute of Technology

D^{2}M: Dynamic Defense and Modeling of Adversarial Movement in Networks
Aug. 2019 Microsoft Advanced Threat Protection Research Expo

Mining Marked Nodes in Large Graphs

Dec. 2018 Microsoft Advanced Threat Protection Group

May 2018 Arizona State University

Local Partition in Rich Graphs

Dec. 2018 IEEE International Conference on Big Data

Rapid Analysis of Network Connectivity

Nov. 2017 ACM International Conference on Information and Knowledge Management (CIKM)

Network Connectivity Analysis and Visualization in Large Graphs

April 2017 Keynote Speaker: ASU Fulton Undergraduate Research Initiative (FURI)

Nov. 2016 ASU FURI Research Symposium

Press

June 2021 "New NVIDIA Partnership Bridges Education Gap for Data Science and Machine Learning",

April 2021 "ML Student Earns Prestigious IBM Ph.D. Fellowship Award",

April 2021 "IBM PhD Fellowship Awardees Announced",

April 2021 "Accelerated Data Science in the Classroom: Teaching Analytics and Machine Learning with RAPIDS",

April 2020 "Georgia Tech and Intel Awarded Multimillion-Dollar Program to Defend Against Attacks on AI",

April 2020 "DARPA Snags Intel to Lead its Machine Learning Security Tech",

April 2020 "Machine Learning Technique Helps Wearable Devices Get Better at Diagnosing Sleep Disorders and Quality",

Feb. 2019 "Raytheon Awards Two ML@GT Students Graduate Research Assistantships",

July 2018 "NSF Graduate Research Fellow wants to use computer science to solve society's toughest problems",

Grants and Funding

2021 **IBM PhD Fellowship**

IBM PhD Fellowship Awardee

Funded: $95,000

2020 **Google Cloud Research Grant**

Large Scale Malware Analysis

Funded: $5,000 Google cloud credits

2018 — 2022 **Guaranteeing AI Robustness against Deception (GARD)**

DARPA Research Grant

Co-PIs: Jason Martin, Duen Horng (Polo) Chau

Funded: multi-million

Helped formulate adversarial defense techniques

2018 **Amazon AWS Research Grant**

Adversarial Re-Training and Model Vaccination for Robust Deep Learning

Funded: $5,000 AWS cloud credits

2018 **Nvidia GPU Grant**

Defending Adversarial Attacks by Robust, Inference-time Local Linear Approximation

Funded: Nvidia Titan V GPU worth $3,000

2019 **Raytheon Research Fellowship**

Extracting Knowledge For Adversarial Detection and Defense
Funded: $25,000

2018 — 2023 **NSF Graduate Research Fellowship Program (GRFP)**
Multi-level Interdiction and Assistance Modeling for Natural Disasters
Funded: Full tuition + $102,000

2016 — 2017 **FURI Grant**
Network Connectivity Analysis and Visualization in Large Graphs
Funded: $3,000

Teaching

Spring 2021 **Graduate Teaching Assistant**
Georgia Institute of Technology, Atlanta, GA
Data and Visual Analytics, Instructor: Duen Horng (Polo) Chau

Fall 2020 **Graduate Teaching Assistant**
Georgia Institute of Technology, Atlanta, GA
Data and Visual Analytics, Instructor: Duen Horng (Polo) Chau

Fall 2013 **Undergraduate Teaching Assistant**
Arizona State University, Tempe, AZ
Fulton Undergraduate Research Experience (FSE 294), Instructor: Joshua Lyon
Designed and taught introductory lesson plans to new engineering students

Mentoring

Present — Spring 2020 **Omar Shaikh**
B.S. in Computer Science, Georgia Institute of Technology

Present — Spring 2020 **Jon Saad-Falcon**
B.S. in Computer Science, Georgia Institute of Technology

Present — Spring 2020 **Frank Zhou**
B.S. in Computer Science, Georgia Institute of Technology

Present — Summer 2020 **Kevin Li**
B.S. in Computer Science, Georgia Institute of Technology

Service

Program Committee
Association for the Advancement of Artificial Intelligence (**AAAI**) at AAAI 2021
ACM International Conference on Information and Knowledge Management (**CIKM**) at ACM CIKM 2020

Reviewer
Practice of Knowledge Discovery in Databases (**ECML-PKDD**) 2021
International Conference on Computer Vision (**ICCV**) 2021
Conference on Computer Vision and Pattern Recognition (**CVPR**) 2021
ACM SIGKDD Conference on Knowledge Discovery and Data Mining (**KDD**) 2019
International Conference on Machine Learning (**ICML**) 2019