JARON MINK

Thomas M. Siebel Center for Computer Science
201 North Goodwin Avenue
Urbana, IL 61801-2302

Phone: +1 (909) 362-0734
Email: jaronmm2@illinois.edu
Web: https://jaronm.ink

RESEARCH INTERESTS

Human Factors in Security and Privacy; Trustworthy Machine Learning; System Security

EDUCATION

University of Illinois at Urbana-Champaign — Aug. 2019 - Present
Ph.D. in Computer Science
Advisor: Gang Wang

University of California, Los Angeles — Sep. 2015 - Mar. 2019
B.S. in Computer Science

AWARDS AND HONORS

NSF Graduate Research Fellowship (GFRP) - *Awarded: $138,000*	Aug. 2021 - Aug. 2024
Magna Cum Laude	2019
UCLA Deans Honor List, UCLA	2019

PUBLICATIONS

Refereed Publications

- [CHI 2024] **Jaron Mink**, Miranda Wei, Collins W. Munyendo, Kurt Hugenberg, Tadayoshi Kohno, Elissa M. Redmiles, Gang Wang. “It's Trying Too Hard To Look Real: Deepfake Moderation Mistakes and Identity-Based Bias”.
  ACM CHI Conference on Human Factors in Computing Systems. Honolulu, HI, USA. May 2024.
  (acceptance rate: 26%)

- [USENIX Sec. 2023] **Jaron Mink***, Harjot Kaur*, Juliane Schmüser*, Sascha Fahl, Yasemin Acar. ““Security is not my field, I’m a stats guy”: A Qualitative Root Cause Analysis of Barriers to Adversarial Machine Learning Defenses in Industry”.
  32nd USENIX Security Symposium. Anaheim, CA, USA. August 2023.
  (acceptance rate: 29%; * equal contribution) [pdf]

- [IEEE SP 2023a] **Jaron Mink**, Hadjer Benkraouda, Limin Yang, Arridhana Ciptadi, Ali Ahmadzadeh, Daniel Votipka, Gang Wang. “Everybody’s Got ML, Tell Me What Else You Have: Practitioners’ Perception of ML-Based Security Tools and Explanations”.
  44th IEEE Symposium on Security and Privacy. San Francisco, CA, USA. May 2023.
  (acceptance rate: 17%) [pdf]

- [IEEE SP 2023b] Muhammad Adil Inam, Yinfang Chen, Noor Michael, Jason Liu, **Jaron Mink**, Sneha Gaur, Adam Bates, Wajih Ul Hassan. “SoK: History is a Vast Early Warning System: Auditing the Provenance of System Intrusions”.
  44th IEEE Symposium on Security and Privacy. San Francisco, CA, USA. May 2023.
  (acceptance rate: 17%) [pdf]

- [ACSAC 2022] Muhammad Adil Inam, Akul Goyal, Jason Liu, **Jaron Mink**, Noor Michael, Sneha Gaur, Adam Bates, Wajih Ul Hassan. "FAuST: Striking a Bargain between Forensic Auditing's Security and Throughput".
  38th Annual Computer Security Applications Conference. Austin, TX, USA. December 2022.
  (acceptance rate: 24.5%) [pdf]

- [USENIX Sec. 2022] **Jaron Mink**, Licheng Luo, Natã M. Barbosa, Olivia Figueira, Yang Wang, Gang Wang. "DeepPhish: Understanding User Trust Towards Artificially Generated Profiles in Online Social Networks".
  31st USENIX Security Symposium. Boston, MA, USA. August 2022.
  (acceptance rate: 16%) [pdf] [New Scientist] [Futurum]

- [WWW 2022] Ziyi Zhang, Shuofei Zhu, **Jaron Mink**, Aiping Xiong, Linhai Song, Gang Wang. "Beyond Bot Detection: Combating Fraudulent Online Survey Takers".
  The ACM Web Conference. Lyon, France. April 2022.
  (acceptance rate: 18%) [pdf] [The Transmitter]

- [CHI 2022] **Jaron Mink**, Amanda Rose Yuile, Uma Pal, Adam J Aviv, Adam Bates. "Users Can Deduce Sensitive Locations Protected by Privacy Zones on Fitness Tracking Apps".
  ACM CHI Conference on Human Factors in Computing Systems. New Orleans, LA, USA. May 2022.
  (acceptance rate: 24%) [pdf] [The 21st Show]

- [ACSAC 2020] Noor Michael, **Jaron Mink**, Jason Liu, Sneha Gaur, Wajih Ul Hassan, Adam Bates. "On the Forensic Validity of Approximated Audit Logs".
  36th Annual Computer Security Applications Conference. Austin, TX, USA. December 2020.
  (acceptance rate: 23%) [pdf]

Pre-Prints

- [Revision 2024b] Miranda Wei, **Jaron Mink**, Tadayoshi Kohno, Elissa M. Redmiles, Franziska Roesner. "SoK or So(L)K? On the Quantitative Study of Sociodemographic Factors and Computer Security Behaviors".
  Under Revision at USENIX Security, 2024.

EXPERIENCE

University of Washington – Visiting Researcher Jun. 2023 - Aug. 2023
Mentor: Tadayoshi Kohno

- [**CHI 2024**] Investigated identity-based biases that occur during deepfake content moderation.
- [**Revision 2024b**] Investigated how security literature quantitatively analyzes and interprets demographic factors.

Human Computing Associates – Consultant Mar. 2022 - Mar. 2023
Mentor: Elissa Redmiles

- Consulted with "Partnership on AI" to identify areas of shared concern and high priority in AI safety.

Max Planck Institute for Software Systems – Visiting Scholar May. 2022 - Aug. 2022
Mentor: Elissa Redmiles

- [**CHI 2024**] Investigated identity-based biases that occur during deepfake content moderation.

Max Planck Institute for Security and Privacy – Research Fellow May. 2021 - Aug. 2021
Mentor: Yasemin Acar

- [**USENIX Sec. 2023**] Investigated ML developers' perceptions of adversarial machine learning and barriers to defense deployment.

Viasat – Network Engineer Intern Mar. 2019 - Aug. 2019
Mentor: Andrew J. Acalinovich

- Researched and developed a network alert algorithm and dashboard for Security Operations Center.

Novacoast – Software Developer Intern Mar. 2017 - Mar. 2019
Mentor: Rouel Soberano, Renato Untalan

- Designed and developed application for experimental remote autism therapy with clinical therapists.
- Developed mobile application to control "FLIR" UAS thermal cameras.

POSTERS

Conference Posters

- **Jaron Mink***, Harjot Kaur*, Juliane Schmüser*, Sascha Fahl, Yasemin Acar. "Security is not my field, I'm a stats guy": A Qualitative Root Cause Analysis of Barriers to Adversarial Machine Learning Defenses in Industry. *The 32nd USENIX Security Symposium*. Anaheim, CA, USA. August 2023. (* equal contribution)

Workshop Posters

- **Jaron Mink**, Miranda Wei, Collins W. Munyendo, Kurt Hugenberg, Tadayoshi Kohno, Elissa M. Redmiles, Gang Wang. "It's Trying Too Hard To Look Real: Deepfake Moderation Mistakes and Identity-Based Bias". *Center for Privacy and Security for Marginalized and Vulnerable Populations Workshop.* Anaheim, CA, USA. August 2023. (Co-located at USENIX)

CONFERENCE TALKS

- "Security is not my field, I'm a stats guy": A Qualitative Root Cause Analysis of Barriers to Adversarial Machine Learning Defenses in Industry. *The 32nd USENIX Security Symposium*. Anaheim, CA, USA. August 2023.
- Everybody's Got ML, Tell Me What Else You Have: Practitioners' Perception of ML-Based Security Tools and Explanations. *44th IEEE Symposium on Security and Privacy*. San Francisco, CA, USA. May 23, 2023.
- DeepPhish: Understanding User Trust Towards Artificially Generated Profiles in Online Social Network. *31st USENIX Security Symposium.* Boston, Massachusetts USA. August 11, 2022.
- Users Can Deduce Sensitive Locations Protected by Privacy Zones on Fitness Tracking Apps. *ACM CHI Conference on Human Factors in Computing Systems*. New Orleans, LA, USA. May 3, 2022.

- On the Forensic Validity of Approximated Audit Logs. *6th Annual Computer Security Applications Conference.* Virtual. December 9, 2020.

INVITED TALKS

- Human Factors in Secure and Non-Abusive AI. *J.P. Morgan: Artificial Intelligence Research Group.* Virtual. January 24, 2024.
- Human Perceptions and Roles Under Emerging Machine Learning Threats. *Georgia Tech: School of Cybersecurity - Student Security Seminar.* Atlanta, GA, USA. November 16, 2022.
- Human Perceptions and Roles Under Emerging Machine Learning Threats. *CISPA.* Hannover, Germany. July 28, 2022.
- Human Perceptions and Roles During Emerging Machine Learning Threats. *Cyber Security in the Age of Large-Scale Adversaries (CASA) Colloquium.* Bochum, Germany. July 6, 2022.
- DeepPhish: Understanding User Trust Towards Artificially Generated Profiles in Online Social Network. *Capital Area Colloquium on Trustworthy and Usable Security/Privacy*. Washington D.C., USA. May 31, 2022.

CODE AND DATASET RELEASE

- ““Security is not my field, I'm a stats guy”: A Qualitative Root Cause Analysis of Barriers to Adversarial Machine Learning Defenses in Industry”.
  https://osf.io/3q54p/
- “FAuST: Striking a Bargain between Forensic Auditing’s Security and Throughput”.
  https://bitbucket.org/sts-lab/faust
- “DeepPhish: Understanding User Trust Towards Artificially Generated Profiles in Online Social Network”.
  https://github.com/JaronMink/DeepPhish
- “Users Can Deduce Sensitive Locations Protected by Privacy Zones on Fitness Tracking Apps”.
  https://bitbucket.org/sts-lab/epz-game-chi22

PROFESSIONAL SERVICE

Program Committee

- [SaTML] IEEE Conference on Secure and Trustworthy Machine Learning 2023
- [EAAMO] ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization 2023

Poster Session Program Committee

- [SOUPS] Symposium on Usable Privacy and Security 2022, 2023

External Reviewer

- [CHI] ACM Conference on Human Factors in Computing Systems 2021, 2022, 2023, 2024
- [NDSS] Network and Distributed System Security Symposium 2021
- [SOUPS] Symposium on Usable Privacy and Security 2021

- [RAID] International Symposium on Research in Attacks, Intrusions, and Defenses 2022
- [TDSC] IEEE Transactions on Dependable and Secure Computing 2023

COMMUNITY SERVICE

SAIL: Cybersecurity Ninja Training Course, UIUC — Apr. 2021, 2022, 2023, 2024

- Designed and ran an interactive presentation to teach high school students about real-world security issues and how to protect themselves.

Safer Illinois: Security Verification Team, UIUC — Aug. 2020 - Dec. 2020

- Investigated security of UIUC's COVID contact tracing app.
- Disclosed findings and provided recommendations to developers.

TEACHING

Guest Lecturer: *CS 463. Computer Security II: Deepfakes - Threats and Mitigations.* UIUC — Fall 2021

Teaching Assistant: *CS 463. Computer Security II.* UIUC — Spring 2021

Learning Assistant: *CS 33. Intro to Computer Architecture.* UCLA — Fall 2018

ADVISING

Ezra Goodwin. B.S.: Univ. of Illinois at Urbana-Champaign — Sept. 2023 – Jan. 2024

- Investigating emerging deepfake use and misuse cases.
- UIUC: CS Student Ambassadors/Research Scholars (CS STARS)

Uma Pal. B.S.: Univ. of Illinois at Urbana-Champaign → PhD: UMass — Aug. 2020 – May. 2021

- [**CHI 2022**] Analyzed qualitative responses and elicited results.
- UIUC: Undergraduate research collaboration

Olivia Figueira. B.S.: Santa Clara Univ. → PhD: UC Irvine — Jun. 2020 – Sept. 2020

- [**USENIX Sec. 2022**] Developed profile stimuli and survey infrastructure.
- Computing Research Association-Widening Participation: DREU

Tooba Hashmi. B.S.: Univ. of Houston-Downtown → MS: Univ. of Houston — May. 2020 - Aug. 2020

- Investigated availability exploits in the routing protocol for Low Power and Lossy Networks (RPL).
- DHS: Summer Research Team for Minority Serving Institutions Program

MEDIA COVERAGE

- Formal Response to NIST's AI Executive Order RFI, Massive Data Institute at Georgetown University. *To be published shortly.* — Feb. 2024
- Scammers threaten quality of research survey data, The Transmitter https://www.thetransmitter.org/ethics/scammers-threaten-quality-of-research-survey-data/ — Aug. 2023

- Can you trust what you see online?, Futurum May. 2023
  https://futurumcareers.com/can-you-trust-what-you-see-online

- How Private is Your Health Tracking App Data?, The 21st Show Aug. 2022
  https://will.illinois.edu/21stshow/story/how-private-is-your-health-tracking-app-data

- People are bad at spotting fake LinkedIn profiles generated by AI, New Scientist Feb. 2022
  https://www.newscientist.com/article/2308977-people-are-bad-at-spotting-fake-linkedin-profiles-generated-by-ai/

- Four Illinois CS Students Earn Highly Coveted NSF Graduate Research Fellowship Jul. 2021
  https://cs.illinois.edu/news/2021-NSF-Grad-Research-Fellowships

- Autism Therapy in the Palm of Your Hand Aug. 2020
  https://news.ucsb.edu/2020/019985/autism-therapy-palm-your-hand