Association for Computing Machinery

Advancing Computing as a Science & Profession

NEWS RELEASE

Contact: Jim Ormond
212-626-0505
ormond@hq.acm.org

LATEST IN CYBERSECURITY RESEARCH UNVEILED AT ACM CONFERENCE ON COMPUTER AND COMMUNICATIONS SECURITY

CCS 2020 Program Addresses Security in Variety of Environments, Including Cloud, Software, Web and Machine Learning Security, Privacy and Censorship

New York, NY, November 4, 2020 —The Association for Computing Machinery's Special Interest Group on Security, Audit and Control (ACM SIGSAC) will hold its flagship annual conference on Computer and Communications Security (CCS 2020) November 9-13. The conference will be conducted virtually for the first time ever due to the global pandemic. Now in its 27th year, CCS presents the leading scientific innovations in all practical and theoretical aspects of computer and communications security and privacy.

"Cybersecurity challenges continue to be a major concern for governments, organizations and individuals," said CCS 2020 General Co-chair Jay Ligatti, University of South Florida. "Our cybersecurity conference provides an opportunity for the world's leading security researchers and practitioners to share research and discuss how to best address the myriad challenges facing a number of industries and computing areas."

Added CCS 2020 General Co-chair Xinming Ou, University of South Florida, "With more than 100 research papers and a robust schedule of presentations, CCS 2020 addresses machine learning and neural networks, blockchain, network security, privacy and censorship, and more. Advances in cybersecurity often begin on the research level, and these papers represent foundational efforts to create a more secure computing environment."

CCS 2020 features research papers in nine categories, including blockchain and distributed systems; machine learning security; network security; privacy and censorship; and software and web security.

2020 CCS HIGHLIGHTS

Keynotes

"Machine Learning and Security: The Good, The Bad, and the Ugly"

Wenke Lee, Georgia Tech University

Wenke Lee explores the good, the bad and the ugly of interactions between machine learning and security and posits that attackers will keep exploiting holes in ML, and automate their attacks using ML. Security professionals need to prepare for ML failures, and ultimately, humans have to be involved. The question is how and when? The cybersecurity challenges threaten ML efforts at a point when there is more data, more powerful machines and algorithms, and better yet, no need to always manually engineer features. However, state-of-the-art models such as deep neural networks are not as intelligible as classical models, such as decision trees. How are deep learning-based model deployed for security when it is unknown for sure if the model is learning correctly, and if the data is poisoning intentionally in misinformation campaigns.

Realistic Threats and Realistic Users: Lessons from the Election

Alex Stamos, Stanford University, former chief security officer at Facebook

Alex Stamos will leverage his experience from inside Facebook during the 2016 and 2018 elections, and from running an election integrity war room in 2020, to discuss the ways that technology fails society. He will discuss the realistic assumptions made about threats, and the expectations of users, and try to chart a path forward for how cutting-edge security research might better inform the engineers and product designers who end up putting computing technologies in the hands of billions.

Research Papers (Partial List) *For a full list of papers, visit the CCS 2020 Program Page.

Gotta Catch'Em All: Using Honeypots to Catch Adversarial Attacks on Neural Networks

Shawn Shan, Emily Wenger, Bolun Wang, Haitao Zheng, Ben Y. Zhao, University of Chicago; Bo Li, University of Illinois at Urbana-Champaign

A new "honeypot" approach to protecting deep neural networks (DNN) is explored in this research paper, which describes the implementation of a trapdoor-enabled defense. The authors analytically prove that trapdoors shape the computation of adversarial attacks so that attack inputs will have feature representations very similar to those of trapdoors. Second, it is experimentally shown that trapdoor-protected models can detect, with high accuracy, adversarial examples generated by state-of-the-art attacks (PGD, optimization-based CW, Elastic Net, BPDA), with negligible impact on normal classification.

Phantom of the ADAS: Securing Advanced Driver-Assistance Systems from Split-Second Phantom Attacks

Ben Nassi, Ben-Gurion University of the Negev, Georgia Institute of Technology; Dudi Nassi, Raz Ben-Netanel, Yuval Elovici, Ben-Gurion University of the Negev; Oleg Drokin, Independent Researcher

In this paper, the authors investigate "split-second phantom attacks," a scientific gap that causes two commercial advanced driver-assistance systems (ADASs) to treat a depthless object that appears for a few milliseconds as a real obstacle/object. To counter this threat, a countermeasure is proposed which can determine whether a detected object is a phantom or real using just the camera sensor, and

demonstrate the countermeasure's effectiveness (it obtains a TPR of 0.994 with an FPR of zero) and test its robustness to adversarial machine learning attacks.

BDoS: Blockchain Denial-of-Service

Michael Mirkin, Ittay Eyal, Technion; Yan Ji, Ari Juels, Cornell Tech; Jonathan Pang, Ariah Klages-Mundt, Cornell University

The authors present Blockchain DoS (BDoS), the first incentive-based DoS attack that targets Proof of Work cryptocurrencies. Unlike classical DoS, BDoS targets the system's mechanism design: It exploits the reward mechanism to discourage miner participation and can cause a blockchain to grind to a halt with significantly less resources, e.g., 17% as of February 2019 in Bitcoin according to the researchers' empirical study. Beyond its direct implications for operational blockchains, BDoS introduces the novel idea that an adversary can manipulate miners' incentives by proving the existence of a secret longest chain without actually publishing blocks.

Text Captcha Is Dead? A Large Scale Deployment and Empirical Study

Chenghui Shi, Shouling Ji, Qianjun Liu, Zhejiang University; Changchang Liu, IBM Research; Yuefeng Chen, Yuan He, Alibaba Group; Zhe Liu, Nanjing University of Aeronautics and Astronautics; Raheem Beyah, Georgia Institute of Technology; Ting Wang, Pennsylvania State University

The development of deep learning techniques has significantly increased the ability of computers to mitigate the security of existing captcha schemes. In this work, the authors introduce advCAPTCHA, a practical adversarial captcha generation system that can defend against deep learning-based captcha solvers--and deploy it on a large-scale online platform with near billion users.

Practical Recommendations for Stronger, More Usable Passwords Combining Minimum-strength, Minimum-length, and Blocklist Requirements

Joshua Tan, Lujo Bauer, Nicolas Christin, Lorrie Faith Cranor, Carnegie Mellon University

Multiple mechanisms exist to encourage users to create stronger passwords, but there is little definitive, scientific guidance on how these mechanisms should be combined and configured to best effect. Through two online experiments, the authors evaluated combinations of minimum-length and character-class requirements, blocklists, and a minimum-strength requirement that requires passwords to exceed a strength threshold according to neural-network-driven password-strength estimates. The results lead to concrete recommendations for policy configurations that produce a good balance of security and usability.

Impersonation-as-a-Service: Characterizing the Emerging Criminal Infrastructure for User Impersonation at Scale

Michele Campobasso, Luca Allodi, Eindhoven University of Technology

This paper provides evidence of an emerging criminal infrastructure enabling impersonation attacks at scale. Impersonation-as-a-Service (ImpaaS) allows attackers to systematically collect and enforce user profiles (consisting of user credentials, cookies, device and behavioral fingerprints, and other metadata) to circumvent risk-based authentication system and effectively bypass multi-factor authentication

mechanisms. Researchers analyzed the operation of a large, invite-only, Russian ImpaaS platform providing user profiles for more than 260,000 Internet users worldwide.

Pre- and Post-Conference Workshops (Partial List)
For a full list of workshops, visit the CCS 2020 Workshops Page

13th ACM Workshop on Artificial Intelligence and Security (AISec)
The AISec workshop is the leading venue for presenting and discussing new developments in the intersection of security and privacy with AI and machine learning. Recent years have seen a dramatic increase in applications of artificial intelligence, machine learning, and data mining to security and privacy problems. The use of AI and ML in security-sensitive domains, in which adversaries may attempt to mislead or evade intelligent machines, creates new frontiers for security research. The recent widespread adoption of deep learning techniques, whose security properties are difficult to reason about directly, has only added to the importance of this research.

2020 Cloud Computing Security Workshop (CCSW)
This workshop will explore cloud-based security concerns. Clouds and massive-scale computing infrastructures are starting to dominate computing and will likely continue to do so for the foreseeable future. Major cloud operators are now comprising millions of cores hosting substantial fractions of corporate and government IT infrastructure. CCSW is the world's premier forum, bringing together researchers and practitioners in all security aspects of cloud-centric and outsourced computing.

18th Workshop on Privacy in the Electronic Society (WPES)
The goal of this workshop is to discuss the privacy problems that result as well as their solutions. The Information Revolution has thoroughly transformed society. One of the major implications of this technological shift has been a massive increase in the collection, sharing, and analysis of personal data. The workshop features research from academia, government, and industry on all theoretical and practical aspects of electronic privacy, as well as experimental studies of fielded systems, and systematization of knowledge (SoK) submissions.

7th ACM Workshop on Moving Target Defense (MTD)
The workshop brings together researchers from academia, government, and industry to report on the latest research efforts on moving-target defense, and to have productive discussion and constructive debate on this topic. The static nature of current computing systems has made them easy to attack and hard to defend. Adversaries have an asymmetric advantage in that they have the time to study a system, identify its vulnerabilities, and choose the time and place of attack to gain the maximum benefit. The idea of moving-target defense (MTD) is to impose the same asymmetric disadvantage on attackers by making systems random, diverse, and dynamic and therefore harder to explore and predict.

2nd Workshop on Cyber-Security Arms Race (CYSARM)
Cybersecurity is a complex ecosystem that is based on several contradicting requirements. For this reason, it is often defined as an arms race between attackers and defenders: The goal of the CYSARM

workshop is to foster collaboration and discussion among cyber-security researchers and practitioners to better understand the various facets and trade-offs of cybersecurity and how new security technologies and algorithms might impact the security of existing or future security models.

About SIGSAC

The mission of the ACM Special Interest Group on Security, Audit and Control (SIGSAC) is to develop the information security profession by sponsoring high quality research conferences and workshops. SIGSAC conferences address all aspects of information and system security, encompassing security technologies, secure systems, security applications, and security policies.

About ACM

ACM, the Association for Computing Machinery, is the world's largest educational and scientific computing society, uniting educators, researchers and professionals to inspire dialogue, share resources and address the field's challenges. ACM strengthens the computing profession's collective voice through strong leadership, promotion of the highest standards, and recognition of technical excellence. ACM supports the professional growth of its members by providing opportunities for life-long learning, career development, and professional networking.

###