Artificial Intelligence and Cybersecurity Workshop

The Hotel at the University of Maryland

7777 Baltimore Ave. College Park, MD 20740

June 4-6, 2019

AI Cybersecurity Workshop Keynotes

PLENARY I: AI FOR SECURITY

The Dangers of the Subconscious Mind (of Cyber Reasoning Systems)
Yan Shoshitaishvili (Arizona State University)

ABSTRACT: Humans have goals, hopes, dreams, and fears. Humans are brilliant. They make incredible intuitive inferences. They conceptualize amazing algorithms to augment cybersecurity. But they can be misled; tricked; foole d into carrying out actions counter to their own best-interests.

The Cyber Grand Challenge propelled program analysis algorithms from mere tools to autonomous Cyber Reasoning Systems. These systems can operate independently to find, exploit, and mitigate vulnerabilities in sof tware, and under various programs and initiatives in the years since the CGC, they have continually improved on their humble beginnings. But similar to otherwise-intelligent and otherwise-autonomous humans, they can also be misled.

What weaknesses exist in the subconscious minds of Cyber Reasoning Systems? Can their dreams turn into nightmares? Can their hopes and goals be hijacked? Of course, the answer is yes. This talk will explore some of the concrete, technical routes to this sort of hijacking, both in terms of what existed in the Cyber Grand Challenge, what has been developed since, and what emerging disruptions and coersions might look like in the future.

BIO: Yan Shoshitaishvili is an assistant professor at Arizona State University, where he pursues research in automated program analysis and vulnerability identification techniques. As part of this, Yan led Shellphish's participation in the DARPA Cyber Grand Challenge, applying his research to the creation of a fully autonomous hacking system that won third place in the competition. Underpinning this system is angr, an open-source binary analysis project created by Yan (and others!) over the years. When he is not doing research, Yan is pushing the area of cybersecurity competitions into the future from his position on the Order of the Overflow, the organizers of DEF CON CTF.

PLENARY II: SECURITY OF AI

Michael Kearns (University of Pennsylvania)

BIO: Since 2002 Michael Kearns has been a professor in the Computer and Information Science Department at the University of Pennsylvania, where he held the National Center Chair. He has secondary appointments in the department of Economics, and in the departments of Statistics and Operations, Information and Decisions (OID) in the Wharton School. He is the Founding Director of the Warren Center for Network and Data Sciences, where his Co-Director is Rakesh Vohra. He is the faculty founder and former director of Penn Engineering's Networked and Social Systems Engineering (NETS) Program, whose current directors are Andreas Haeberlen and Aaron Roth. He is a faculty affiliate in Penn's Applied Math and Computational Science graduate program. Until July 2006 I was the co-director of Penn's interdisciplinary Institute for Research in Cognitive Science.

As of August 2018, He is affiliated with the Santa Fe Institute as an external faculty member.

Along with Yuriy Nevmyvaka (with whom he have also collaborated on a number of papers on algorithmic trading), he lead applied research in the AI Center of Excellence at Morgan Stanley,

He has worked extensively in quantitative and algorithmic trading on Wall Street (including at Lehman Brothers, Bank of America, and SAC Capital; see further details below). He often serves as an advisor to technology companies and venture capital firms. He is also involved in the seed-stage fund Founder Collective and occasionally invest in early-stage technology startups. He is a member of the Scientific Advisory Board of the Alan Turing Institute, and of the Market Surveillance Advisory Group of FINRA. He occasionally serves as an expert witness or consultant on technology-related legal and regulatory cases.

He is an elected Fellow of the American Academy of Arts and Sciences, the Association for Computing Machinery, the Association for the Advancement of Artificial Intelligence, and the Society for the Advancement of Economic Theory.

PLENARY III: INDUSTRY RESEARCH

Ulfar Erlingsson (Google)

BIO: Ulfar Erlingsson is a Senior Staff Research Scientist in the Google Brain team, currently working primarily on privacy and security of deep learning systems. Previously, Ulfar has led computer security research at Google, and been a researcher at Microsoft Research, Silicon Valley and Associate Professor at Reykjavik University. Ulfar was co-founder and CTO of the internet security startup Green Border Technologies and Director of Privacy Protection at deCODE Genetics. Ulfar holds a PhD in computer science from Cornell University.

PLENARY IV: ACADEMIC RESEARCH

Security Against Adversarial Examples
David Wagner (University of California at Berkeley)

ABSTRACT: Recent research suggests that modern machine learning methods are fragile and easily attacked, which raises concerns about their use in security-critical settings. I will survey several attacks on machine learning and directions for making machine learning more robust against attack. I will also briefly mention my own research in this area.

BIO: David Wagner is Professor of Computer Science at the University of California at Berkeley, with expertise in the areas of computer security and electronic voting. He has published over 100 peer-reviewed papers in the scientific literature and has co-authored two books on encryption and computer security. His research has analyzed and contributed to the security of cellular networks, 802.11 wireless networks, electronic voting systems, and other widely deployed systems.

PLENARY V: GOVERNMENT MISSION CONTEXT

Mona Lisa Talks
Dean Souleles (Office of the Director of National Intelligence)

Mr. Dean Souleles rejoined the Office of the Director of National Intelligence (ODNI) as the Chief Technology Advisor to the Principal Deputy Director of National Intelligence (PDDNI) in October of 2017. He provides subject matter expertise, assessment, and advice on all relevant technology matters to the PDDNI and senior leadership in the ODNI and the Intelligence Community (IC). Mr. Souleles leads an effort, with ODNI and IC partners, to assess, evaluate and develop IC wide strategies to accelerate the adoption of key technology priorities for the IC including Artificial Intelligence (AI), machine learning and automation to increase the IC's ability to assess, understand and add context to the rapidly accelerating volume, velocity and variety of data relevant to Intelligence.

Prior to rejoining ODNI, Mr. Souleles served as the founding Chief Technology Officer (CTO) of the National Counterintelligence and Security Center (NCSC) where he established the Office of the CTO, and was responsible for developing and communicating the National Counterintelligence and Security Technology Strategy and providing direct oversight of all of NCSC's technology activities and systems.

Prior to NCSC, Mr. Souleles served the ODNI as the Special Assistant for the Intelligence Community Information Technology Enterprise (IC ITE). He advised the PDDNI and the IC Chief Information Officer (CIO) on management and technical issues related to the development and implementation of IC ITE. Mr. Souleles founded and served as the co-chairman of the IC Architecture and System Engineering Team (ASET) and oversaw the operations of the IC ITE Mission User Group (the "MUG") which serves as the "Voice of Mission" to IC ITE.

Mr. Souleles has over 35 years of experience in information technology in a variety of industries prior to joining ODNI in 2012. Mr. Souleles served as Chief Operating Officer of Resolution Health, Inc. (RHI) a venture-funded healthcare, big data analytics and software company that was acquired by Anthem/WellPoint in 2008. Before that, he served as Executive Vice President and Chief Technology Officer of QuadraMed Corporation, a $140 million international public healthcare information technology firm (formerly AMEX: QD) Earlier in his career, Mr. Souleles spent 10 years at NASA's Jet Propulsion Laboratory (JPL) where he served as Principal Engineer and System Architect on various JPL space, civil and defense programs. And, after leaving JPL in the mid-1990s he founded SureNet, one of the earliest dial-up internet services providers.

Mr. Souleles was educated in Computer Science at the California State University at Northridge.