

THE CENTER FOR ADVANCED COMPUTER STUDIES

at

the University of Louisiana at Lafayette

Lafayette, Louisiana

Proudly announces a presentation

Dr. Philippe Burlina

*Associate Research Professor in Computer Science
Johns Hopkins University*

and

Dr. Yinzhi Cao

*Assistant Professor in Computer Science
Johns Hopkins University*

will give a presentation entitled

Fairness and Privacy In AI Applied to Healthcare

Abstract: Current AI frameworks using deep learning (DL) have met or exceeded human capabilities for tasks such as classifying images, recognizing objects (ImageNet) or facial expressions. Medical AI is also approaching clinicians' performance for diagnostics tasks. This success masks real issues in AI assurance that will impede the deployment of such algorithms in real life production systems. Two of the most critical concerns affecting AI assurance are privacy and bias. Privacy leakages can invalidate HIPAA or HITECH compliance, interdicting the use of DL diagnostic models in healthcare. Additionally, DL systems' performance strongly depends on the availability of large, diverse, and representative annotated training datasets, which are often unbalanced with regard to factors such as gender, ethnicity and/or disease type, resulting in diagnostic bias. In this talk, we will introduce our recent progress in developing algorithms to address bias as well as approaches to assess possible risks in existing algorithms for privacy/membership attacks and propose ways to effectively defend against such privacy attacks.

DATE: FRIDAY, APRIL 30, 2021

TIME: 11:00 A.M. – 12:00 NOON

LOCATION: Via Zoom

Biography: Philippe Burlina is an associate research professor in Computer Science at Johns Hopkins University. He earned his M.S. and Ph.D. in Electrical Engineering at The University of Maryland, College Park and Diplome d'Ingenieur at the Universite de Technologie de Compiègne in France. His research work is focused on computer vision and machine learning challenges that impact autonomy and healthcare, with emphasis on problems including fairness in AI, robustness and domain generalization, low shot and zero shot learning, anomaly detection, and semantic approaches to generative modeling.

Yinzhi Cao is an assistant professor in Computer Science at Johns Hopkins University. He earned his Ph.D. in Computer Science at Northwestern University and worked at Columbia University as a postdoc. Before that, he obtained his B.E. degree in Electronics Engineering at Tsinghua University in China. His research mainly focuses on the security and privacy of the Web, smartphones, and machine learning. His past work was widely featured by over 30 media outlets, such as NSF Science Now (Episode 38), CCTV News, IEEE Spectrum, Yahoo! News and ScienceDaily. He received two best paper awards at SOSP'17 and IEEE CNS'15 respectively. He is a recipient of the Amazon ARA award.

Zoom Meeting:

<https://ullafayette.zoom.us/j/95570241062?pwd=QTlRV3FYVjdWScTnVTFRN2c2QmErUT09>

Passcode: Fri11Semi