THE UNIVERSITY OF CHICAGO

MACHINE-AUGMENTED HUMANS AS A PRIVACY ARMOR

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE DIVISION OF THE PHYSICAL SCIENCES
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

DEPARTMENT OF COMPUTER SCIENCE

BY
YUXIN CHEN

CHICAGO, ILLINOIS
AUGUST 2022

ABSTRACT

As a variety of smart and connected sensors are being deployed everywhere, significant privacy issues arise since these devices can constantly capture our (private) behaviors in forms of image, video and sound. These sensor data can be used by adversaries to attack personal privacy. For example, leaked audio data can be processed using machine learning models to extract private conversation, track user activity, identify human speakers or even generate any speech in the voice of the speaker. These privacy attacks are fully automated and can be launched at scale. As a result, they pose a real security and privacy threat to everyone. Yet protecting users against such intrusive sensing is challenging. Privacy laws and policies can help regulate the use of sensors and machine learning models, but they are known to be difficult and slow to deploy.

In this dissertation, we explore a user-centric approach for protecting personal privacy against intrusive sensing. We propose to develop low-cost wearables that users can carry and turn on/off to prevent their private information from being extracted by unauthorized parties. Along this line, we design and engineer novel wearables that emit safe and carefully designed signals to protect both content and identity privacy. Our wearables also leverage the inherent properties of the human body to improve protection strength and coverage. Together, these wearables and the human body form a powerful privacy armor, providing users with full agency in privacy control.

This dissertation makes three key contributions.

First, to protect our speech privacy, we engineer a wearable microphone jammer as a bracelet, which disables surrounding microphones, including hidden ones. Our design leverages a hardware property that, when exposed to ultrasonic noise, commodity microphones will leak the noise into the audible range, which disrupts the speech recording. Our jamming bracelet also leverages natural body movements to increase protection coverage and effectiveness.

Second, we study privacy protection for our typing content, where we examine the idea of using VR headsets to disable keystroke inference attacks. We show that typing in VR environments can naturally defeat existing attacks because the keyboard and its layout are invisible in the physical world. We then develop a new, more sophisticated attack that can successfully infer VR keystroke content using just a RGB camera. This presents a new threat against VR typing privacy and the need for additional protection methods.

Finally, we also study identity privacy and its impact on user authentication. Since our normal biometrics data, such as face, voice, and fingerprint, can be easily captured by sensors and leaked to attackers, we develop an alternative, wearable-based authentication method based on muscle stimulation. Our proposed system authenticates a user by stimulating the user's forearm muscles with a sequence of electrical impulses (a challenge) and measuring the user's involuntary finger movements (response to the challenge). Our system produces 68 million challenges per user, using just one second of muscle stimulation. Attackers replaying used responses will be rejected, making our system highly robust against data breach and leakage.

In summary, this dissertation develops user-centric solutions to protect personal privacy against intrusive sensing, by augmenting the human body with wearables to form a ubiquitous privacy armor. We hope our work sheds light on the development of personal privacy protection in the physical world.

CHAPTER 1

INTRODUCTION

A variety of smart and connected sensors are populating into our physical world and revolutionizing the way we live, work and play. Today, our devices, homes, offices, private and public spaces are full of sensors, including cameras, microphones, biometric sensors, and many others.

However, as these sensors continue to flourish, significant privacy issues arise since they can constantly capture and save our (private) behaviors in forms of image, video and sound, either maliciously or by misconfiguration [82, 198, 176]. Using powerful machine learning (ML) models, adversaries can process these sensor data to extract our private information, and thus launch significant and unacceptable attacks against user privacy. Taking audio as an example. Leaked data can be processed by ML models to extract confidential conversation [198, 41, 40], track user activity [18], infer typed text and handwriting content [15, 219, 205], identify speakers [87] or even generate any speech in the voice of the speaker [86]. As these attacks are becoming fully automated, they pose a significant threat to all of us.

Despite significant concerns against intrusive sensing and privacy attacks, there are few tools available to protect users against them. Privacy laws and policies could help regulate the use of sensors and ML models, but they are known to be difficult and slow to deploy. A notable example is facial recognition. Despite the significant media backlash against intrusive facial recognition services like *Clearview.AI* [43], the legislative efforts to address these services remain elusive in the US [100, 8, 27].

Overview of My Work In this dissertation, we explore the idea of protecting personal privacy against intrusive sensing by low-cost wearables that users can carry and turn on/off to prevent their private information from being extracted by unauthorized parties. Our wearables also leverage the inherent properties of the human body to improve protection

strength and coverage. Together, these wearables and the human body form a powerful privacy armor, providing users with full agency in privacy control.

Along this line, we design and engineer novel wearables that emit safe and carefully designed signals to protect both content and identity privacy. We start with engineering a wearable jammer to protect speech privacy by preventing unauthorized microphones from capturing and extracting our speech. We then focus on protecting typing privacy, where we examine the idea of using VR headsets to disable keystroke inference attacks. Finally, we propose a novel user authentication system using on-arm electrical muscle stimulation, which relies on disposable, privacy-insensitive muscle stimulation-response records rather than face/voice/fingerprint.

In the following, we briefly introduce the work included in this dissertation.

1.1 Wearable microphone jamming for speech privacy protection

Voice-based smart devices are everywhere and becoming an integral part of our life. Their implementation requires them to equip microphones that can always monitor and record our speech. Recent studies have shown that these devices can be exploited by adversaries to extract the content of our private speech. In the first part of my thesis, we seek to build a wearable that protects users' speech privacy.

To meet this goal, we engineer a wearable jammer that is worn as a bracelet. When turned on, it emits inaudible ultrasonic noise that disables microphones in the wearer's surroundings, including hidden microphones. Our design leverages a hardware property that, when exposed to ultrasonic noise, commodity microphones will leak the noise into the audible range, which disrupts the speech recordings. By building the ultrasonic jammer as a bracelet and arranging ultrasonic transducers in a ring layout, our wearable emits jamming signals in multiple directions. This eliminates the need for the wearer to manually point the jammer to a microphone. Our wearable also leverages natural body movements to increase

protection coverage and effectiveness.

We confirm experimentally that our jammer is superior to existing stationary jammers by conducting a series of technical evaluations and a user study. The results demonstrate that (1) our wearable jammer largely outperforms existing static jammers in coverage; (2) it remains effective even if the microphones are hidden and covered by various materials, such as cloths or paper sheets; and, (3) our study participants feel that our wearable protects the privacy of their voice.

1.2 Examining the use of VR headsets for typing privacy protection

Besides speech, keyboard typing is another modality we regularly use to produce content. Prior work on keystroke inference attacks shows that attackers can infer our typing content when they know the keyboard and its layout. On the other hand, when using VR-based systems [131], the keyboard and its layout are visible to the user wearing the VR headset but remain hidden to any physical observers. This suggests that VR systems could naturally protect our typing privacy.

The second part of my thesis seeks to understand whether VR systems truly protect our keystroke privacy, where the attacker has no knowledge of the user's keyboard location, size, or layout, but can only observe their finger/hand movements at a distance. My research develops a new, more sophisticated attack that can successfully infer VR keystroke content using just a RGB camera. This presents a new threat against VR typing privacy and the need for additional protection methods.

Our proposed attack uses a two-layer self-supervised system. In layer 1, noisy results of hand tracking on the keystroke video are used to detect keystrokes, followed by a language model to recognize keystrokes. These initial labels are filtered using multiple consistency checks to produce high confidence labels on video frames. Next, in layer 2, these labels

and their corresponding video frames are used to train two 3D-CNN models that detect and recognize keystrokes from the raw video frames.

We evaluate this attack using IRB-approved user studies under a variety of conditions, varying the target (user/typing behavior, content typed, physical environment) and attacker behaviors (hand tracking tool, attack distance). The attack is highly effective in nearly all settings, and performs well across our user study participants, despite significant different typing styles and abilities. This presents a new threat against VR typing privacy and the need for additional protection methods.

1.3 On-arm electrical muscle stimulation for user authentication

Intrusive sensing also poses significant implications in biometric authentication. Today, when our biometric data (face/voice/fingerprint) are leaked to attackers, attackers can use them to bypass authentication systems used by banks and other critical services. There is nothing users can do to securely re-use their own data, as these biometrics are static.

To tackle this challenge, we explore a novel modality for active biometric authentication: electrical muscle stimulation (EMS). Our system, which we call ElectricAuth, stimulates the wearer's forearm muscles with an EMS challenge, i.e., a 1.2s sequence of electrical impulses and then measures the user's involuntary finger movements as a result of this challenge.

ElectricAuth authenticates users by leveraging what is typically seen as the biggest disadvantage of EMS: intersubject variability, i.e., the same electrical stimulation results in different movements in different users because everybody's physiology is different [96, 52, 58, 42, 132]. These differences arise from multiple compound factors in the field of muscle biomechanics and physiology [71, 107, 3]. All these differences add up to create individual responses to the same stimulus, which our system uses as the key feature to identify a user.

ElectricAuth also generates a very large pool of challenges by exploring an underutilized property of EMS: muscles respond differently depending on their current state of contraction,

which can be altered by varying the timing between two impulses. Using four muscles, six impulses and seven time gaps, ElectricAuth encodes one of 68M possible challenges in 1.2s. As such, ElectricAuth is robust against data breaches and replay attacks because it never reuses the same challenge twice in authentications – ElectricAuth rejects replay of recorded responses to any previously used challenges and can quickly recover from leak/breach of either authentication model or stored challenge-response pairs by asking the user to register responses to a new set of challenges (like registering new one-time passwords).

We evaluate our prototype of ElectricAuth by conducting a series of technical evaluations and user studies. The results demonstrate that: (1) ElectricAuth offers accurate user verification and resists three common biometric attacks: impersonation, replay and synthesis attacks; (2) ElectricAuth performs stably over 21 days against various muscle conditions (fatigue, humidity, etc.); (3) ElectricAuth can verify the user in 3ms on laptop's CPU and 35ms on a small embedded device after receiving a response, and can use either IMUs or a depth camera to track finger movements.

1.4 Structure of this dissertation

This dissertation is organized into five chapters. After this introduction (Chapter 1), we present a wearable microphone jammer that protects speech content privacy by disabling microphones in the wearer's surroundings (Chapter 2). We then focus on protecting keyboard typing, where we examine the idea of using VR headsets to disable keystroke inference attacks (Chapter 3). In the following chapter, we also study identity privacy and its impact on user authentication, where we propose a novel active biometric authentication system that authenticates a user by stimulating the user's forearm muscles with a sequence of electrical impulses and measuring the user's involuntary finger movements (Chapter 4). Finally, we conclude this dissertation by summarizing our contributions, and discussing the insights learned as well as potential future directions for this emerging field (Chapter 5).