

Research Statement

Paarijaat Aditya

January 2017

Motivation: Privacy threats in a world of ubiquitous computing

Sophisticated mobile computing, sensing and recording devices are now commonplace. Smart phones have already achieved significant penetration and novel devices like Snapchat Spectacles & Microsoft HoloLens are imminent. Moreover, these devices are carried by their users virtually round the clock, blurring the distinction between the online and offline world and enabling transformative new applications. For instance, mobile apps can provide location- and activity-sensitive information, which can be overlaid onto a user's field of view using smart glasses. Mobile devices can also maintain a detailed record of a user's life, recording everything a user does, sees, hears and who he meets, enabling creative new applications that have become an integral part of our lives.

However, these mobile applications and services have also introduced a range of new threats to a user's privacy. While a detailed personal record of a user's life is useful for his own reference, it is also highly sensitive and inherently private. This private data is accessed by mobile apps running on a user's device and is often uploaded to the app provider for further processing. Unfortunately, once a user's personal data leaves his device, he has little (or no) control over how this data is used by the app provider. This creates a serious privacy challenge and numerous real world instances of misuse of users' personal data for profit by organizations have only aggravated these concerns [1, 2].

Even in cases where a user does not use these mobile applications or devices himself, his privacy might still be at risk, e.g., when a user is captured in photos taken by nearby strangers that are later shared on online social networks, revealing the whereabouts of everyone photographed. Such unwanted image capture is perceived to be such a serious privacy threat that it led to Google Glass being banned at numerous venues and was likely one of the factors that led to its discontinuation [3]. This failure in adoption of latest technology strongly highlights that for future mobile technology to be broadly accepted, privacy concerns cannot be treated as an afterthought but rather must be a first-order concern while designing mobile apps and hardware.

My agenda: Privacy by design while preserving functionality

On the first look it may appear that loss of privacy is inherent in ubiquitous context-sensitive applications. However my research has shown that it possible to achieve privacy by design in these mobile apps without losing intended functionality.

My research enables the design of mobile apps in a way that puts the users back in control of what personal data is shared and how it is used. In particular, I investigated two specific contexts in which privacy problems arise: 1) risks to a user's personal information introduced by an increasingly popular class of apps called mobile social apps, 2) privacy risks due to ubiquitous digital capture, where bystanders may inadvertently be captured in photos and video recorded by other nearby users.

1. **EnCore: Private, Context-based Communication for Mobile Social Apps:** [4]

The first part of my dissertation focused on mitigating privacy risks introduced by mobile social apps which consider users' location, activity and nearby devices to provide context-aware services; e.g. sharing captured images with nearby users, detecting the presence of friends in close vicinity, sharing news and gossip with nearby people and helping people find missed connections. Most of the currently deployed mobile social apps rely on a trusted cloud service to match and relay information, requiring users to reveal their whereabouts (potentially including a continuous trace of their location), the perils of which have been extensively noted [5, 6, 7, 8, 9].

To mitigate this problem, I built and deployed EnCore, a mobile platform that builds on secure encounters between pairs of devices as a foundation for privacy preserving communication for mobile social apps. An encounter occurs whenever two devices discover each other within Bluetooth radio range and generate a unique encounter ID and an associated shared secret. EnCore detects nearby users, groups them in named communication abstracts called events, and enables encrypted communication and sharing among event participants, all while relying on existing network, storage, and online social network services. By relying on device-to-device communication to discover nearby users, EnCore puts users in control of the privacy and confidentiality of the information they share with the app provider and other users.

2. **I-Pic: A Platform for Privacy-Compliant Image Capture:** [10]

In the second part of my dissertation I focused on the privacy risk introduced by ubiquitous digital capture facilitated by smartphone cameras, smart glasses, and life-logging cameras. Bystanders may be photographed (intentionally and/or inadvertently) without their consent, which poses a significant risk to their privacy and security.

To mitigate this risk, I built and deployed I-Pic, a trusted software platform that integrates digital capture with user-defined privacy. In I-Pic, users choose a level of privacy (e.g., image capture allowed or not) based on social context (e.g. in public vs. with friends vs. at workplace). Privacy choices of nearby users are advertised via BLE (Bluetooth Low-Energy), and I-Pic-compliant capture platforms generate edited media to conform to the privacy choices of the captured subjects. I-Pic uses a state-of-the-art deep neural network for face recognition and combines it with secure multiparty computations to ensure that users' visual features and privacy choices aren't revealed publicly, regardless of whether they are the subject of an image capture. Just as importantly, I-Pic preserves the ease-of-use and spontaneous nature of image capture and sharing between trusted users.

Summary:

Both my dissertation projects were aligned with the overall aim of providing a platform for building mobile apps in a privacy-compliant manner that puts users back in control of what personal information is collected and shared. Specifically I demonstrated that one can preserve most of the functionality of spontaneous image capture and device-to-device communication without giving up privacy. I believe that even though a single platform alone is unable to provide a comprehensive end-to-end privacy-preserving infrastructure, technical innovations that mitigate specific risk vectors will provide a strong basis for facilitating a broad societal conversation about the value of user privacy, how basic services are deployed, and shape future research in privacy-compliant ubiquitous computing.

Future Research:

In the long term, I want to conduct multidisciplinary research and development that aims to bring mobile apps and technologies to developing nations. I am passionate about building research prototypes and testing them in real world deployments. In both my dissertation projects I conducted multiple real world deployments that I want to continue doing in my future projects as well.

In the short term, I want to work towards making I-Pic part of a mobile operating system (e.g. Android), so that mobile devices, in particular smart glasses, could advertise themselves as privacy aware devices. I believe this would help alleviate some the privacy concerns that plagued the first release of Google Glass and the concerns that surround more recent devices, such as Snapchat Spectacles [11]. Towards this larger goal, there are business and legal challenges that need to be tackled as well, but my initial focus will on the numerous technical challenges that are still left.

Deep neural networks on resource constrained mobile devices: One of the technical challenges is how to execute resource intensive deep learning algorithms on resource constrained (in terms of battery, memory and cpu) mobile devices. Deep neural networks have become extremely popular in the recent past and are increasingly being adopted for major machine learning tasks ranging from text, speech, and image analysis. Motivated by this opportunity, I want to systematically explore the space of possibilities for designing and executing deep learning algorithms given the limited resources on mobile devices. Energy-efficient solutions to this problem would also encourage app developers to process user data locally on the device instead of shipping it to the cloud, which in turn would help keep user data private by design.

Securely offloading user data and machine learning computations to the cloud: Even with efficient solutions for processing data locally on mobile devices, there are many apps that require shipping user data to the cloud, for e.g. an augmented reality app provider may not want to store their proprietary deep learning network on a user's device for privacy reasons or, e.g., a cab sharing app provider may want to combine data from multiple users before providing relevant information to individual users. To enable these apps while preserving privacy of user data one could use homomorphic encryption to securely offload data and computations to the cloud. Homomorphic encryption provides a way to perform basic computations, like addition and multiplication, over encrypted data. Unfortunately these basic operations are not sufficient to support all the computations required by deep neural networks, which additionally require operations like comparison and max-pooling. As future work I want to explore alternative approximations for these operations and design & test deep neural networks with these approximations in place.

Possibilities at the intersection of EnCore and I-Pic: To make an actual deployment of I-Pic more feature rich, I want to add the capabilities provided by EnCore to I-Pic. E.g., since encounters in EnCore enable communicating with people who were present in the same vicinity as the user, one could use these encounters for I-Pic to support delayed communication of privacy preferences by the captured subjects. Using past encounters photographed subjects could first receive a copy of the images they are part of and then later communicate their privacy preference to the photographer. I believe such a functionality would make I-Pic more practical and fun to use and also contribute towards the larger goal of making I-Pic more accessible.

References

- [1] Paarijaat Aditya, Bobby Bhattacharjee, Peter Druschel, Viktor Erdélyi, and Matthew Lentz. Brave new world: Privacy risks for mobile users. In *The Workshop on Security and Privacy Aspects of Mobile Environments (SPME 2014)*, 2014.
- [2] Natasha Lomas. WhatsApps privacy U-turn on sharing data with Facebook draws more heat in Europe. <https://techcrunch.com/2016/09/30/whatsapp-privacy-u-turn-on-sharing-data-with-facebook-draws-more-heat-in-europe/>.
- [3] Victor Luckerson. Google will stop selling Glass next week. <http://time.com/3669927/google-glass-explorer-program-ends/>.
- [4] Paarijaat Aditya, Viktor Erdelyi, Matthew Lentz, Elaine Shi, Bobby Bhattacharjee, and Peter Druschel. EnCore: Private, Context-based Communication for Mobile Social Apps. In *The 12th ACM International Conference on Mobile Systems, Applications, and Services (MobiSys'14)*, 2014.
- [5] Google fires engineer for violating privacy policies. <http://www.physorg.com/news203744839.html>.
- [6] Joseph A. Calandrino, Ann Kilzer, Arvind Narayanan, Edward W. Felten, and Vitaly Shmatikov. "you might also like: " privacy risks of collaborative filtering. In *The IEEE Symposium on Security and Privacy*, S&P '11, 2011.
- [7] Keir Thomas. Microsoft cloud data breach heralds things to come. http://www.pcworld.com/article/214775/microsoft_cloud_data_breach_sign_of_future.html, 2010.
- [8] Liana B. Baker and Jim Finkle. Sony PlayStation suffers massive data breach. <http://www.reuters.com/article/2011/04/26/us-sony-stoldendata-idUSTRE73P6WB20110426>, 2011.
- [9] Fahmida Y. Rashid. Epsilon data breach highlights cloud-computing security concerns. <http://www.eweek.com/c/a/Security/Epsilon-Data-Breach-Highlights-Cloud-Computing-Security-Concerns-637161/>, 2011.
- [10] Paarijaat Aditya, Rijurekha Sen, Seong Joon Oh, Rodrigo Benenson, Bobby Bhattacharjee, Peter Druschel, Tongtong Wu, Mario Fritz, and Bernt Schiele. I-Pic: A Platform for Privacy-Compliant Image Capture. In *The 14th ACM International Conference on Mobile Systems, Applications, and Services (MobiSys'16)*, 2016.
- [11] Sasha Lekach. Privacy Panic? Snapchat Spectacles raise eyebrows. <http://mashable.com/2016/11/16/snapchat-spectacles-privacy-safety/>.