Submission for NSF Workshop on College Student Health

Towards Privacy-aware Sensemaking On-the-go for College Student Health & Safety

Mani Srivastava (mbs@ucla.edu)

Modern computing systems can now fuse information from diverse sensors embedded in our environment and devices to draw and share rich inferences about our state, inform decisions and trigger interventions. This capability is behind many innovative applications and services in different domains and, as the Workshop Call also points out, is likely well-suited for helping college students manage their health and well-being.

Apps making use of sensory information could guide students towards healthier food choices at dining halls, help them find friends to safely walk back to the dorm with after a campus drinking party, perhaps even detect if drinks at a frat party have been spiked with drugs, detect signs of depression based on in-person social interactions and usage of social networks, and ensure healthy life with adequate sleep and fitness activities. While the higher technological literacy of College Students is a facilitator for adoption of sensing technologies, the unique attributes of college life and also present challenges.

Consider, for example, an app that might assist students going to campus parties with drinking but who are fearful of sexual assault. Often the hardest thing is to find a friend to walk back with, particularly as students hop across parties all over the campus. A fun partying style app could provide tracking for the duration of a party so that when someone, upon being done with the party, can locate others in their group. However, such an app would also present privacy and perception concerns, which if not addressed would prevent its adoption. For example, students would not want their use of such an assistive app to be known and therefore would prefer such tracking capability to be integrated into a well-accepted app that they already use, such as Facebook or Instagram. For the duration of the party the app would need to track fine-grained location. However, with current mobile OS the permission to share location is an install-time all-or-nothing decision with no context- or need-dependent control over the sharing permission, accuracy and precision. So giving the aforementioned app permission to track and share fine-grained location would be a permanent decision, leading to a continual violation of one's privacy.

More generally, the sharing of sensor information in current mobile and wearable devices at the OS-app interface is a poorly thought out one from a privacy perspective. Firstly, sensory information is shared in the form of high-dimensional raw measurements, allowing apps to make not only the necessary inferences but many additional and potentially sensitive ones. Secondly, user's control over the sharing of sensor information with apps is limited - a static one with no control over the fidelity. The current state of affairs results from smartphone's origin as not a sensory devices but as a networked computing device. Ergo, while mobile OS provide

sophisticated communication stacks that virtualize and police apps accesses to networks, the sensing stacks are barebones with apps given direct access to low-level high-dimensional sensor data with the attendant risks. Furthermore, the analytics on the sensor data is often done on the cloud-side, partly because of computing resource limitations as currently apps do not get direct access to specialized sensor processing cores that are present on mobile SoCs, but mostly because of monetization-driven data grab by app developers.

For smartphones to truly achieve their potential as devices with apps that continually assist based on sensory information (and intervene as necessary), the mobile OSs must handle sensory information in a responsible fashion by providing a proper privacy-sensitive "sensing stack" for sensor information to flow through. Such a stack would need to accomplish two things. Firstly, instead of exposing raw sensor measurements whose privacy implications are hard to fathom let alone exercise control over, the currency of sensor information exchange at the OS-app interface must be raised to one of semantically understandable "inferences" derived from the measurements. Realizing such a sensing stack would require that sensor analytics is pushed to the other side of the app-OS boundary without access to cloud computing resources, and thus require development of lightweight embedded machine learning algorithms and processing resources optimized for them.

Secondly, the sensing stack would need to control per-app context-dependent control over the fidelity with which the OS shares sensory information with the apps. This could, for example, be done via selective obfuscation of sensory information. However, given the large number of sensors and user contexts, a conventional user-configured permission system with obfuscation rules would be excessively burdensome. Instead, perhaps, a inference management system that learns user's privacy needs and configures the fidelity of information sharing automatically would be needed, for example, by adapting the fidelity with which location is shared with Facebook depending on user's context: more precisely at the party but coarsely at other times.

Brief Biography: Mani Srivastava is on the faculty at UCLA where he is associated with the EE Department with a joint appointment in the CS Department. His research is in the area of networked human-cyber-physical systems, and spans problems across the entire spectrum of applications, architectures, algorithms, and technologies. His current interests include issues of energy efficiency, privacy and security, data quality, and variability in the context of systems and applications for mHealth and sustainable buildings. He is a Fellow of the IEEE.

Acknowledgements: This submission is based in part on ideas from author's daughter Megha Srivastava, who is a sophomore at Stanford University and generously shared her perspective as a college student.