Toward Societal Scale Sensing using Mobile Phones

Andrew T. Campbell and Tanzeem Choudhury - Dartmouth College

Participants: Andrew T. Campbell leads the MetroSense project on people-centric sensing. His group developed a number of early sensing systems and applications for sensor-enabled mobile phones (e.g., CenceMe and SoundSense). Tanzeem Choudhury leads People-Aware Computing Group and has pioneered machine learning techniques for mobile systems (e.g., SoundSense, Mobile Sensing Platform) that can reason about human activities, interactions, and social networks in everyday environments.

Vision: The recent trend of incorporating an increasing number of sensors into mobile phones opens up a new research frontier that has the potential to significantly impact many aspects of our everyday life from healthcare, to safety, entertainment, and business. However, the broad impact of this vision is jeopardized without advances in the computational models we use to process the sensor data generated by mobile phones. These models turn raw data into inferences (ranging from inferring a physical activity, to tracking community wide patterns) that make these applications possible. Existing models are unable to deliver societal scale sensing because they are domain specific, do not adapt to new types of people and environment, and are trained at the level of a single person. When data is available for multiple people, each person is considered independent and predictions are made about each separately. Furthermore, existing approaches are too computationally expensive to be implemented on existing mobile phones without severely limiting the phones performance and classification accuracy. Therefore, new statistical models must be developed to scale and duty cycle on mobile phones without limiting the main function or lifetime of the phone.

To address these problems we envision community-guided learning that leverages contributions of sensor data from a diverse population of users to intelligently scale computational modeling and classification on the phone. Traditional sensor networks take advantage of many sensors in static environments in time and space. We propose utilizing the similarities in the behavioral patterns people exhibit as well as the proximity of people in time and space and propose to develop: 1) new sensing and learning techniques that exploit the structure present in community sourced data to enable scale; 2) learning algorithms that selectively leverages the human-in-the-loop to adapt and personalize applications; 3) techniques that allow groups of collocated mobile phones to exchange or pool classifier models to achieve better inference performance and robustness; and 4) open sensing and inference software and APIs for the implementation of duty cycled and computational light sensing and inference models across multiple vendors phone (e.g., Nokia, iPhone, Android, Window Mobile).

Impact: By developing new sensing and inference techniques based on the principles discussed above we will enable societal scale sensing using everyday mobile phones. Top end senor-enabled mobile phones will be cheap and globally ubiquitous within the next 5 years. Researchers and developers will be capable of potentially distributing their new idea/applications or sensing systems to 100s of millions of mobile phones across the globe with the potential for societal scale sensing. To get to this point there will need to be major advances in computational models for inferencing, ultimately pushing more intelligence to the phone without limiting its ability to function as a phone.