

Scaling Human Attention: End-User Interaction with Machine Learning In Pervasive Systems

James Fogarty

Computer Science & Engineering
University of Washington

jfogarty@cs.washington.edu

<http://www.cs.washington.edu/homes/jfogarty/>

Vision

Our ability to collect and store information has far outpaced our ability to actually understand and interact with it. In this context, the rise of large-scale pervasive computing creates an apparent paradox: people continue to demand that computing do more with less. Massive heterogeneous data is inherent to large-scale pervasive systems, so we expect more powerful methods for accessing, interacting with, and acting upon this data. But we also expect devices to be smaller and more mobile, expect to spend little or no time learning about or managing an interaction, and are unwilling to tolerate an application that seeks our attention at an inappropriate time or presents information that seems irrelevant to our current needs. In other words, people impose ever greater expectations for interacting with ever larger and more heterogeneous data, but remain unwilling to spend more space, time, or attention on computing.

The past twenty years have seen the emergence of search as a tool for managing an information explosion on the Web. An equally fundamental transformation is required to address the explosion of data created by large-scale sensor-based pervasive systems. I therefore believe the next twenty years will see machine learning embedded in every aspect of human-computer interaction. The familiar desktop metaphor, designed thirty years ago to support interaction with dozens of objects, will not scale to this future. This whitepaper briefly motivates a need for fundamentally new approaches to end-user interaction with machine learning in sensor-based pervasive systems, including: (1) new methods for creating and composing machine learning components, (2) new methods for end-user interpretation and feedback to machine learning, and (3) new methods for collaboratively defining machine learning components of pervasive systems.

Advancing the Field

Effective approaches to large-scale pervasive computing promise broad social contributions. One important application is helping elders and their caregivers manage chronic health conditions. Over 45% of non-institutionalized Americans have one or more chronic conditions, and their direct costs of care account for 75% of U.S. health care expenditures (Hoffman, Rice, and Sung, *Journal of the American Medical Association*). Existing research generally cannot address this need because it is limited by assumptions of sensor platforms too expensive or intrusive for widespread deployment, overly optimistic assumptions regarding the consistency and homogeneity of sensor data, or a lack of effective methods for non-expert configuration, querying, and analysis of large sensor streams. The rest of this section uses examples from this motivating application, but the applicability of our proposed research agenda is much broader.

We need new methods for end-user creation and composition of machine learning components across multiple scales of heterogeneous sensor data. In an application like home elder care, every home is different and every elder's needs are different. It is therefore critical that a

caregiver can quickly deploy sensors appropriate for a given situation and then effectively train those sensors to detect necessary information. Some activities should be trained by demonstration (e.g., teaching a fingerprint-based location tracking system to recognize when an elder is sitting in a favorite chair), while others need to be specified from a higher-level representation (e.g., setting an alarm when an elder has not moved from that chair for more than eight hours should not require one or more demonstrations of that scenario). The field needs new methods that unify these two approaches by allowing end-users to quickly demonstrate intermediate concepts that can then be used to learn or specify higher-level constructs.

We need new methods for end-users to interpret and give meaningful feedback to machine learning systems. Classic active learning approaches treat people as simple oracles capable of giving positive or negative labels to selected examples upon request, but people can potentially offer much richer information to machine learning systems. For example, a nurse might directly advise a system regarding the relevance of available sensors or features derived from those sensors (potentially reducing the complexity of a machine learning problem and the number of training examples needed for a system to reliably recognize a desired activity). This requires new methods for conveying how sensor-based machine learning systems work, what they have learned, and how a person can give meaningful feedback. People can also make mistakes or have incorrect assumptions regarding what information is predictive, so systems need to be able to re-evaluate end-user feedback in the presence of additional data collected as a system is used.

We need new methods for collaboratively defining machine learning components of pervasive systems. The diversity of real-world needs and situations makes it impossible to pre-package one-size-fits-all machine learning components for many sensor-based applications (e.g., the activities of “walking the dog” or “going to the grocery store” are different for everybody). But there are also many commonalities in how people conduct the same activity (e.g., the time of day an activity typically occurs, the typical duration of the activity, the objects a person brings with them for the activity). Instead of starting from scratch when training a system to recognize a particular person’s activities, we need new methods for re-using information about how other people conduct that same activity. This might include small-scale sharing (e.g., a community nurse partially re-using activity specifications they create when working with multiple elders) or massive analyses of anonymous aggregate data (e.g., learning which sensors are most reliable for detecting different activities, learning general models to bootstrap personalization processes).

Background and Experience

James Fogarty is an Assistant Professor of Computer Science & Engineering at the University of Washington and a key member of the DUB cross-campus effort in human-computer interaction. He is a regular member of top HCI program committees (e.g., UIST 2010, CHI 2010, CHI 2009, UIST 2008, UbiComp 2008, CHI 2008, UIST 2006) and has received NSF support through the CAREER, Human-Centered Computing, and Cyber-Enabled Discovery and Innovation programs.

His research has included extensive work on sensor-based statistical models of human interruptibility (CHI 2003, CHI 2004, Best Paper CHI 2005, CHI 2006, CHI 2007) and new approaches to home activity sensing (UIST 2006, Best Paper Nomination UbiComp 2009, Pervasive 2011). He has also extensively examined end-user interaction with machine learning on the Web (UIST 2007, CHI 2008, UIST 2008, Best Paper Nomination CHI 2009, UIST 2009, CHI 2010). He has recently begun to examine end-user interaction with machine learning in sensor-based systems (Pervasive 2010), as he believes this is critical to the future of computing.