

Improving Activity Classification for Health Applications on Mobile Devices using Active and Semi-Supervised Learning

Brent Longstaff, Sasank Reddy, Deborah Estrin
Center for Embedded Networked Sensing
University of California Los Angeles
blongstaff, sasank@ucla.edu, destrin@cs.ucla.edu

Abstract—Mobile phones' increasing ubiquity has created many opportunities for personal context sensing. Personal activity is an important part of a user's context, and automatically recognizing it is vital for health and fitness monitoring applications. Recording a stream of activity data enables monitoring patients with chronic conditions affecting ambulation and motion, as well as those undergoing rehabilitation treatments. Modern mobile phones are powerful enough to perform activity classification in real time, but they typically use a static classifier that is trained in advance or require the user to manually add training data after the application is on his/her device. This paper investigates ways of automatically augmenting activity classifiers after they are deployed in an application. It compares active learning and three different semi-supervised learning methods: self-learning,

Activity data are also useful as contextual data or meta-data. They provide context for other health measures such as physiological self-tests (blood pressure, blood glucose, weight) as well as reporting of symptoms and side effects. Finally, activity traces can improve the user interface mechanisms across a range of applications by increasing the relevance and adherence of its users. For example, they could be reminded or triggered for action or input at a convenient moment, or a moment of interest to the study. Because of the importance of this activity data stream, we investigate how to improve the performance of activity classification using smartphones.