NSF Panel workshop – Pervasive Computing at Scale

Scalable Activity Recognition

Narayanan C Krishnan

Background and Experience:

I am a postdoctoral researcher working in the area of activity recognition in smart environments. In particular, I am exploring machine learning and pattern recognition frameworks for recognizing activities in large scale datasets(data collected over months) gathered from real-world environments. My doctoral dissertation focused on a computational framework for wearable accelerometer based activity and gesture recognition.

Vision:

Activity recognition in smart environments is a challenging problem critical to the development of a wide variety of applications. This research area has seen tremendous progress due to applications such as elder care (health care) related assistive technologies. Understanding human activities and intentions is vital to determine the assistance for promoting safe and independent living. Advances in pervasive computing has resulted in development of unobtrusive, wireless and inexpensive sensors for gathering activity information. This has in turn contributed to the development of a variety of techniques for recognizing activities of daily living (ADL). These techniques have demonstrated good performance in controlled environments with a single resident following a well-defined routine. However, there is still room for improvement for making these techniques scalable to new environments, users and sensing technologies. This white paper focuses on two research tracks that will help advancing current activity recognition systems to the next level for wide spread adoption.

Research Track 1: A typical activity recognition system focuses on recognizing a small set of activity labels. One has to deal with the problem of irrelevant or out-of-vocabulary samples when scaling these systems to work in real-world environments. Typical real-world data often consists of large amounts of irrelevant(out-of-vocabulary) data with sparse and sporadic occurrence of relevant activity data. Design and development of models of out-of-vocabulary samples are crucial for scaling current activity recognition systems as the focus is shifted towards correctly classifying only the relevant samples, by filtering out the irrelevant samples. This process has the potential to improve the recognition accuracies of relevant samples. The irrelevant samples can include noise, abnormalities or even relevant samples of activities that have not yet been considered. Thus in future, the out-of-vocabulary samples can be data mined to discover additional activities relevant to a particular user.

Current activity recognition approaches rely on models learned from relevant activity samples to derive a fixed probabilistic threshold to distinguish between relevant and out-of-vocabulary samples. Generative models such as Gaussian mixture models or hidden Markov models directly model the probability or likelihood of a sample belonging to a particular class. Probability values are approximated from classification margins for discriminative classifiers such as adaptive boosting and support vector machines(SVM). These models are insufficient when the threshold for classification changes according to the data sample, which is the case with real-world samples. Developing machine learning models that model out-of-vocabulary samples either directly as a separate class or indirectly through learning relevant samples will be of immense merit as these models can determine a classification threshold that depends on the sample. Our current experiments with a variant of one class SVM [1] on large real-world activity datasets collected as part of the CASAS project at WSU[2] has demonstrated promising results. The proposed approach performs discriminative density estimation similar to one class SVM on each class independently and then fuses the probability outputs from models belonging to each class. It was observed that filtering out irrelevant samples, before performing relevant sample classification resulted in an increase in precision by 3% and in recall by around 6% on data collected from three different pervasive environments.

Research Track 2:
Current computational frameworks for activity recognition are limited in terms of their ability to “quickly” adapt to new environments, users or sensing technologies. Overcoming problems such as changes in the layout of the environments, addition of new sensing technologies and variation in the activity traits across individuals make an activity recognition system robust and easily deployable. To overcome these problems, traditional systems rely on collecting labeled training data in the new context. However, collecting sufficient amounts of new training data is an expensive process and at times, infeasible. Thus alternate techniques have to be designed and developed to overcome this problem. Current advances in the machine learning with regards to transfer learning, multiple source learning and other related areas might hold a key to the above described problem. These approaches are primarily concerned with the problem of utilizing existing knowledge to train models that work on new domain or context.
Initial experiments conducted using novel transfer learning methodologies has shown promise in recognizing activity gestures through wearable accelerometer data across different contexts. The recognition accuracy improved by 5% when a variant of an instance based transfer learning technique was used for recognizing five different activity gestures such as 'pour', 'scoop' etc on data collected from unseen subjects in real-world environments[3]. There is also evidence in the literature [4] that illustrates superior performance of these techniques on activity data collected across different environments. While this is just a tip of the iceberg, further explorations are required to adapt and design new techniques to address the changes occurring in the activity data due to environments, users and sensing technologies. Advancement of these techniques have the potential to transform the activity recognition domain by improving the generalization of the algorithms across different contexts.

References:
[1] N. C. Krishnan, “A computational framework for wearable accelerometer based activity and gesture recognition”, PhD Thesis, Arizona State University, 2010.
[2] P. Rashidi and D. Cook, “Keeping the resident in the loop: Adapting the smart home to the user”,
IEEE Transactions on Systems, Man, and Cybernetics, Part A: Systems and Humans, 39(5):949-959, 2009.
[3] A. Venkatesan, N. C .Krishnan, and S. Panchanathan, “Cost sensitive boosting for concept drift” European conference on Machine Learning workshop on handling concept drift in Adaptive information systems, 2010.
[4] P. Rashidi and D. Cook, “Transferring learned activities in smart environments”, Proceedings of the International Conference on Intelligent Environments, 2009.