Activity Recognition Research: The Good, the Bad, and the Future

A.J. Bernheim Brush, John Krumm, and James Scott,
Microsoft Research

I. INTRODUCTION

Successfully recognizing people's activities enables a wide range of pervasive computing applications. Recent research on activity recognition, particularly for elder care and health applications, has demonstrated that it is possible to recognize a variety of activities such as driving, walking, and using stairs and elevators [e.g. 2,5,6,9,10,13]. While we are inspired by existing activity recognition research, we believe as a community there are steps we can take together to enable future breakthroughs that are robust and reproducible.

Our general interest in this workshop stems from our research in ubiquitous and pervasive computing. A.J. Brush's main research interest is technology for homes and families, in particular supporting sharing, sustainability, and helping with everyday problems such as scheduling and coordination. John Krumm focuses on techniques for measuring a person's location and for using location data to benefit the user. He has worked on predicting driving routes and destinations and looked for routines in logs of activity data. James Scott conducts research on sensors and devices, mobile interaction, energy management, security and privacy.

Two current projects led to our particular interest in this workshop. First, as part of a study to understand arrival and departure prediction for households, we have been collecting GPS data from 12 households which we would like to provide to other researchers. This has required additional effort to collect and anonymize data to address privacy and legal concerns, which we would be interested in discussing with other members of the community. Second, in a project that involves recognizing activities using sensors on mobile devices, we are more interested in building experiences based on recognized activities than conducting foundational research on activity recognition ourselves. Ideally we would be able to build on previous activity recognition research, but we have not found this easy to do. In the rest of this paper, we describe what we think the community is doing well, places we believe improvement would be beneficial, our recommendations for reporting on activity recognition research, and next steps we feel could benefit the community.

Authors are listed in alphabetical order. A.J. Brush, and John Krumm are with Microsoft Research Redmond Lab, One Microsoft Way, Redmond, WA, 98052 USA (e-mail: {ajbrush, jckrumm @ microsoft.com).

James Scott is with Microsoft Research Cambridge, 7 JJ Thomson Ave, Cambridge CB3 0FB, UK (e-mail: jws@microsoft.com).

II. DOING WELL

Past research on activity recognition has had some notable successes. In particular, there are examples of end-to-end applications, particularly in health monitoring (e.g. UbiFit [3]) and navigation for people with mild cognitive disabilities (e.g. OpportunityKnocks [11]). Special purpose hardware such as the Multi-modal Sensor board (MSB) [12] built by researchers at Intel Research Seattle and University of Washington, has enabled exploration of the value of different sensors in inferring activities. Researchers have also demonstrated the possibility of recognizing many different physical activities from sitting and walking, to sit-ups and teeth-brushing. A variety of machine learning techniques including decision trees, Bayes classifiers, and nearest-neighbor algorithms have also been explored for activity recognition.

Collecting annotated sensor data can be a challenge for activity recognition and we are aware of at least two examples where researchers have provided datasets that others can use. The PlaceLab project at MIT [7] makes available several multi-modal sensor datasets from the PlaceLab live-in laboratory, and the MIT Reality Mining project [8] makes available data collected on mobile phones including location, communications, nearby devices and phone status.

III. NEEDS TO IMPROVE

Given activity recognition is a relatively young field of research, there are ways in which we believe the community could improve.

First, there is no standard taxonomy of activities used by researchers. The closest taxonomy is the Activities of Daily Living [1], but these are very general. For example "Moving Around" rather than the types of specific activities that researchers have been trying to recognize (e.g. walking, running, going up stairs). Having a shared set of activities and definitions would help ensure that researchers are collecting and labeling ground truth data in a standard way. In addition to a lack of agreed on set of activities, there is typically no common sense reasoning about activities. For example, I'm unlikely to be driving and riding a horse at the same time or brushing my teeth while walking. Lastly, there are no agreed on prior probabilities for activities. The probability that a person is watching television is likely much greater than the probability he or she is fixing a car.

Second, in our experience it has been difficult to build on past work and leverage classifiers others have trained. This is particularly important for researchers interested in building prototypes that could leverage activity recognition. They may not be interested or have the machine learning skills to advance basic research in activity recognition. We recognize there are many potential hurdles to sharing classifiers and algorithms, but at the same time would like to explore how fundamental building blocks of activity recognition could be made more widely available to the research community.

IV. Recommendations

When reviewing activity recognition papers we appreciate the following:

- Confusion matrices for the classifiers that make it clear how well activities were recognized and which activities where most often confused with each other.
- Precision and recall data for the classifiers.
- Details about the data collected including definitions of the activities, and information about how the data was collected and labeled.
- Information about which sensors, set of sensors, and numerical features were most valuable for recognizing activities and which were not helpful.

We would prefer not to see the following:

- Only positive results. We feel systems should be tried to the point of failure. If the paper has only positive results it suggests that the system was not adequately validated across a wide enough range of scenarios
- Unrealistic groups of activities: Some papers include many similar activities and then one that seems quite different, for example walking, using the elevator, and brushing teeth. This choice of many different activities should be justified, as it is difficult as a reader to understand why such different activities have been chosen.

V. Next Steps

For a community interested in furthering activity recognition research, we believe there are several possible next steps that would be valuable.

First, we could define a set of activities to recognize, which could be augmented and expanded as necessary. This would help us exploit existing common sense reasoning work like the Cyc project [4] to make taxonomies of activities. For example, TV watching might be a subset of relaxing. Similarly, a valuable research contribution would be to discover and catalog probabilities surrounding activities. These include simple priors (e.g. watching television is generally more probable than fixing a car), conditionals (e.g. given a local time of 3 a.m., sleeping is more probable than being awake), and Markov sequences (e.g. after riding a bus, walking is more probable than flying).

Pragmatically, we could also create a location for hosting shared datasets, or even just an index listing their locations. This would make it easier for other researchers to leverage shared data. It would also be nice to explore methods for recognizing and rewarding the extra effort necessary to provide shared data to others in the community. Another possibility is creating an activity recognition contest, similar to those held in natural language processing or speech recognition, where labeled data is provided and people compete to do the best recognition. This type of contest could also include building applications that use activity or context recognition for end to end applications. A "grand challenge" contest for applications in a particular domain (e.g. healthcare, eldercare) might also encourage innovation.

We would also be interested in exploring ways of sharing classifiers and algorithms for activity recognition. While this likely has many hurdles including concerns about intellectual property, asking each group that wishes to use activity recognition to reinvent the algorithms or recreate classifiers that have been mentioned in research papers strikes us as counter-productive.

Finally, pervasive computing also has a long history of research related to location. While activity recognition and location are often explored separately (with of course some exceptions) we would like to explore ways to encourage applications that combine activity and location recognition.

VI. Conclusion

Research on activity recognition is critically important to enable a wide range of pervasive applications. Working together as a community will help increase the impact of future research in activity recognition and we would be excited to participate in this workshop.

References

[1] Activities of Daily Living, http://en.wikipedia.org/wiki/Activities_of_daily_living

[2] Bao, L. and Intille, S., Activity Recognition from User-Annotated Acceleration Data. In Proc. Pervasive 2004, Springer (2004) 1-17.

[3] Consolvo, S., McDonald, D., Toscos, T., Chen, M.Y., Froehlich, J., Harrison, B., Klasnja, P., LaMarca, A., LeGrand, L., Libby, R., Smith, I., & Landay, J.A. Activity Sensing in the Wild: A Field Trial of UbiFit Garden. In Proc. CHI 2008, ACM (2008) 1797-1806.

[4] Cyc, http://en.wikipedia.org/wiki/Cyc

[5] Huynh, T., Blanke, U. and Schiele, B. Scalable Recognition of Daily Activities from Wearable Sensors. In Proc. LOCA 2007, Springer (2007).

[6] Lester, J., Choudhury, T., Borriello, G., A Practical Approach to Recognizing Physical Activities. In Proc. Pervasive 2006, Springer (2006) 1-16.

[7] MIT PlaceLab Datasets, http://architecture.mit.edu/house_n/data/PlaceLab/PlaceLab.htm

[8] MIT Reality Mining Data set, http://reality.media.mit.edu/

[9] Ofstad, A., Nicholas, E., Szcodronski, R., and Choudhury, R. R. AAMPL: accelerometer augmented mobile phone localization. In Proc. MELT Workshop. ACM (2008).

[10] Patterson, D. Liao, L, Fox, D., and Kautz, H. Inferring High-Level Behavior from Low-Level Sensors, In Proc. Ubicomp 2003, Springer-Verlag (2003) LNCS 2864, pp. 73-89.

[11] Patterson, D., Liao, L, Gajos, K., Collier, M., Livic, N, Olson, K. Wang, S., Fox, D., Kautz, H., Opportunity Knocks: A System to Provide Cognitive Assistance with Transportation Services. In Proc. Ubicomp 2004. Springer. pp. 433-450.
[12] Sensor Boards, http://ubi.cs.washington.edu/wiki/index.php/Sensor_Boards
[13] Ravi, N., Dandekar, N, Mysore, P., and Littman, M. Activity Recognition from Accelerometer Data, In Proc IAAI 2005, AAAI (2005) 1541—1546.