Introducing New Sensors for Activity Recognition

Dawud Gordon, Hedda Schmittke and Michael Beigl
Distributed and Ubiquitous Systems
Technische Universität Braunschweig,
Braunschweig, Germany
Email: {gordon, schmidtke, beigl}@ibr.cs.tu-bs.de

***Abstract*—When introducing a novel sensor to the context recognition community, one of the major challenges is to support reproducability under similar conditions. In order to get a grasp on this process, we divide the context recognition into three subsections: the physical environment and how it is affected by the context, the sensor and how it can represent the attributes of the physical world, and the classifying method and how it deciphers the sensory representation. We then outlined our recommendation for a methodology to formally describe the context and sensor subdivisions in order to isolate and quantify error within the system. The result would develop a basis of standard models of activities and contexts within the community which would serve to improve evaluation of novel sensors and classification algorithms.**

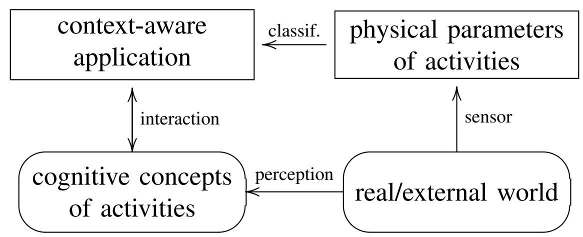

Fig. 1. Three sub-fields of activity recognition research: evaluation of context-aware computing (top left) relies on HCI research methods; activities as human cognitive concepts can be studied using methods from cognitive and social sciences (bottom left); and evaluation of the physical parameters measured by sensors requires methods from the natural sciences (top right).

I. Introduction

As human beings, we mentally structure the physical and social reality surrounding us in a way that is both efficient and practical for survival in conditions for which the evolution process optimized us for. Highly optimized perception and cognitive processes help us to, for instance, to separate moving objects as a foreground from a stable background environment. But we not only carve out solid objects from the perceptual stream, such as a table or the moon, but also less tangible objects of cognition, such as a headache or the orbit of the moon [5]. Activity recognition research sets out to chart a particularly difficult terrain of objects of cognition. Activities we aim to recognize include comparatively simple to describe mechanical processes [2], [6], such as 'walking', which has received scientific study and clear definitions in the area of ergonomics, for instance, but also complex socio-spatial processes, such as meetings [1].

In order to develop a system that can successfully recognize and distinguish a range of activities, we have to implement a measurement and classification mechanism in a way so that it carves up reality correctly, that is, so that it is compatible with the categories of activities that we apply. Figure 1 illustrates our perspective on the three main sub-fields of activity recognition research and how they are linked by the components of activity recognition systems. We argue in this paper that activity recognition research can be separated into the three parts of sensing, classification and interaction, if the three interfacing research efforts can be carried out in a coordinated way. In particular, we support the claim that activity recognition can be made more independent from analysis of classification and interaction parts by benchmarking it with a set of standard classification methods []. In addition however, we emphasize that the physical parameters of the activities to be recognized with the sensor must be understood from related works on cognitive concepts of activities. We hypothesize that introducing a new sensor for activity recognition with a careful analysis of whether the sensor can measure the relevant physical properties of certain activities with the required amount of reliability ensures the usefulness of the sensor for consecutive research. We argue that with this method reproducibility of evaluation experiments would be facilitated and improved, while retaining completeness of exposition in papers.

II. Novel Sensors in Activity and Context Recognition

Due to the highly innovative nature of activity and context recognition, new sensors are continuously being introduced as inputs for activity classification. The sensors themselves vary greatly in terms of physical phenomena measured, data output format, size, accuracy, reliability, and resource consumption. This introduces a dilemma for researchers looking to select the appropriate sensor to recognize a specific activity or context, as the multitude of sensors available in the literature and their suitability for a given scenario can be daunting.

In this section we introduce a recommendation for an improvement to the standard approach for introducing a novel sensor to the activity and context recognition community. This approach is designed to provide the community with valuable information as to the uses of the sensor in context recognition, how to best construct a recognition system based on the novel sensor, and how to integrate the sensor into existing systems.

A. What is the community doing well?

The approach for evaluating novel sensors for activity or context recognition has coalesced into a fairly standardized process. The common approach is to use well-known, standard classification mechanisms, e.g. those of the WEKA toolkit[1], which are widely-used in the community and have been proven effective at classifying similar activities using standard sensors. Using three or more mechanisms gives the reader an impression of the distribution of recognition rates, which can aid in the sensor and classification mechanism selection process. Using only a single classification method represents a disadvantage for readers in the community who would like to use the sensor in an existing system with a different method, as it does not allow for an estimation of how the system would perform under the untested conditions. Good examples of this approach can be found in [7], [3], [4].

During the evaluation of the classification systems the community uses multiple subjects to gather data. This counteracts the influence of interpersonal variation on the evaluation, as well as the influence of spread characteristics of different instances of the same sensor (sensitivity, calibration, etc). Another approach which is helpful to the community is the analysis of the effects of the novel sensor when it is integrated into an existing classification system. This informs the community as to the effects of the novel sensor on conventional systems, in particular, it indicates what types of information the sensor can sample which other sensors cannot.

B. Where does the community need to improve?

We see a need for improvement in reproducibility of evaluation experiments as well as completeness of exposition in context and activity recognition publications on novel sensors, in terms of their usefulness to and impact on the context recognition community as a whole.

a) No sensor is perfect: Each sensor has the ability to convert specific physical phenomena in its environment into signals which relay this information. Every type of sensor is specialized to measure specific parameters while ignoring other properties of the environment; for example, an accelerometer does not relay information about light intensities in its environment: it is specialized to convert acceleration levels into electronic signals. There is no "perfect" sensor which can relay a complete representation of its environment in all of its facets, but rather each sensor is specialized in representing a single facet of the environment at a certain level of detail.

This implies that each sensor can only create a representation for a minute subset of the physical phenomena in its environment. In context recognition, machine learning algorithms search these representations for dependencies and correlations in various situations. Using this information, recognition systems attempt to generate knowledge about how these situations affect the attributes in the physical environment by examining the representations relayed by the sensor.

[1]http://www.cs.waikato.ac.nz/ml/weka/

b) No classifier is perfect: Recognizing an activity in a given situation correctly would therefore mean that a) the activity which has been recognized actually is present in the physical environment at the time, b) that the sensor is able to create and relay a useful and reliable representation of physical parameters of the environment that are affected by the activity, and c) that the classification algorithm is able to decipher these intricacies, yielding a correct activity classification. Not recognizing an activity on the other hand, could be a result of many different influences: a) the activity or context which is meant for classification may not be affecting the physical environment which is being monitored, b) the sensor which is being used cannot relay sufficient information about the physical parameters which are affected by the activity, or c) the classifier is not able to classify the data although the the information is theoretically available.

In addition, activities and contexts are rarely 100% recognizable, but rather each system generates varying recognition rates. This can also have many sources. A low classification percentage for a specific context could be due to the inability on the part of the classifier to consistently recognize the patterns created by the activity. It could also be due to using a sensor which cannot measure the parameters of the environment which are affected by the activity, or, the problem could be that the context does not affect the environment in a fashion which is consistently correlated with the presence of the activity. Any error is passed up the chain: the sensor can only deliver a measurement relevant for an activity if the activity is mainly determined by parameters that can reliably be measured by the sensor; the classifier can only classify on the basis of what is presented to it by the sensor. In other words the quality of context is compounded at every level, beginning with the quality with which a context affects the environment, compounded by the ability of the sensor to represent and relay those influences, and for the classifier to recognize and label the context in the representation.

c) A need-to-know basis: When publishing a context recognition paper which introduces a novel sensor as a basis for context, the goal should be to provide the community with all information necessary to use the technology presented. This implies that results in the introductory paper should document *experiments that can be reproduced in a cost-effective way*, like for instance, using multiple classifiers is a cost-effective standard approach to addressing the problem of possible error in the classifier. The multiple classification method approach prevents false values based on the inadequacies of a single classification method, and also provides colleagues with a basis for assessment using easy to obtain systems.

Where the community falls short, is the evaluation of the sensors themselves and their representations of attributes of the physical environment: papers should clearly specify physical limitations of the sensors, so as to provide a *complete* account of what the sensor can be used for and how its specification relates to that of other sensors with respect to relevant physical parameters of a range of intended types of activities. The commonly used method is to select a range of activities for classification and to test these across multiple classification algorithms. Detected variations in recognition rates then give the reader, and therefore the community, an impression of what is recognizable using the novel sensor; but what they do not provide is information about the classes that could not be recognized well. The question then remains where in the chain the information about the activity is being lost: was it never there to begin with? Was the sensor not able to sample and relay the information reliably? Are the classifiers insufficient? It is here that we see the crux of where activity research can improve in order to increase impact and reproducibility of context recognition applications using novel sensors.

C. Our best recommendation

In order to provide the community with as much pertinent information as possible as to the usage and application of the novel sensor, two vital analyses should be carried out during the evaluation. First, a study should be done on the contexts or activities which are to be recognized, in order to find out exactly how they affect the physical environment. The results of the study should be physical specifications of the context or activity in terms of a definition of what is and what is not the activity. Second, a study of the novel sensor must be conducted in order to create a physical/mathematical model of how exactly the sensor creates a representation of its environment and exactly which parameters of the environment can be relayed.

Using these two studies, a relational mapping from the presence of the context to the output of the sensor can be inferred, specifically, which environmental parameters affected by the context are represented in the information relayed by the sensor, and which information is lost due to sensor-activity incompatibility. Such a model would greatly benefit the community by providing scientists with a method of evaluating if the context which they wish to classify (assuming they can model its affects on its environment) can be represented using this specific sensor. If models are accurately portrayed, they can serve to greatly reduce the amount of effort required for error diagnostics in context recognition. Using these models it is possible to localize errors due to inadaquate classification methods, as the formal activity and sensor descriptions isolate any error in the classifier algorithm.

rapidly expand our knowledge about activity recognition methods: results on errors could then be used to demonstrate and remedy inadequacies of a given classification method for a given sensor-activity combination.

D. Next steps

Since models for how the most commonly applied sensors (accelerometer, light, temperature, GPS, etc.) represent the environment already exist, the first big step is to create physical and mathematical models for widely employed contexts and activities and how they affect their physical environment. These can be verified based on past classification results using algorithms and sensors which are well understood to infer the effects of the sensors. The result is a basis of standard models of activities and contexts within the community which serve to improve evaluation of novel sensors and classification algorithms. Using this basis, each new sensor can then be evaluated using known context and classifier models, creating a faster as well as more complete process for introducing a novel sensor to the context recognition community.

E. Our work

We are currently working on an activity classification system based on a single ball switch which yields frequency-based vibration levels. The system evaluation is being conducted using the guidelines presented in this paper, namely extensive evaluations of the mechanical properties of the sensor itself, as well as studies into formal descriptions for the activities which we attempt to classify. An initial analysis of the properties of the ball switch indicate that at certain frequencies it is well equipped to convert mechanical motion into a digital event time-line.

III. Conclusion

Human cognitive abilities represent an extremely powerful context classification structure, compared to which machine-based systems are still dwarfed in terms of quality and complexity. We introduced a division of context recognition into three subsections: the physical environment and how it is affected by the context, the sensor and how it can represent the attributes of the physical world, and the classifying method and how it deciphers the sensory representation. We outlined our recommendation for a methodology to formally describe the context and sensor subdivisions in order to isolate and quantify error within the system. We argue that using these methods the reproducibility of evaluation experiments under slightly different conditions would be facilitated and improved, while retaining completeness of exposition in papers.

The results of this process would be a growing basis of knowledge within the community. First, a collection of formally defined activities would accrue among context recognition researchers. Second, with each publication of context recognition with novel sensors, a new specification of a sensor and its capabilities is added to the knowledge base. Finally the only open point left to be addressed are benchmark datasets for evaluation of the context classifying methods. This separation of concerns would allow for faster, iterative progress of context recognition research.

References

[1] Michael Beigl, Hans-Werner Gellersen, and Albrecht Schmidt. Mediacups: experience with design and use of computer-augmented everyday artefacts. *Computer Networks*, 35(4):401–409, 2001.

[2] Seon-Woo Lee and Kenji Mase. Activity and location recognition using wearable sensors. *IEEE Pervasive Computing*, 1:24–32, 2002.

[3] Georg Ogris, Matthias Kreil, and Paul Lukowicz. Using fsr based muscule activity monitoring to recognize manipulative arm gestures. In *ISWC '07: Proceedings of the 2007 11th IEEE International Symposium on Wearable Computers*, pages 1–4, Washington, DC, USA, 2007. IEEE Computer Society.

[4] G. Pirkl, K. Stockinger, K. Kunze, and P. Lukowicz. Adapting magnetic resonant coupling based relative positioning technology for wearable activitiy recogniton. In *Wearable Computers, 2008. ISWC 2008. 12th IEEE International Symposium on*, pages 47–54, 28 2008-Oct. 1 2008.

[5] Barry Smith. Fiat objects. *Topoi*, 20:131–148, 2001.

[6] Emmanuel Munguia Tapia, Stephen S. Intille, William Haskell, Kent Larson, Julie Wright, Abby King, and Robert Friedman. Real-time recognition of physical activities and their intensities using wireless accelerometers and a heart rate monitor. *Wearable Computers, IEEE International Symposium*, 0:1–4, 2007.

[7] Kristof Van Laerhoven and Hans-Werner Gellersen. Spine versus porcupine: A study in distributed wearable activity recognition. In *ISWC '04: Proceedings of the Eighth International Symposium on Wearable Computers*, pages 142–149, Washington, DC, USA, 2004. IEEE Computer Society.