

Activity Recognition Research is More Than Finding the Ultimate Algorithms or Parameters

Benjamin Poppinga
OFFIS – Institute for
Information Technology
Oldenburg, Germany
poppinga@offis.de

Susanne Boll
University of Oldenburg
Oldenburg, Germany
boll@informatik.uni-oldenburg.de

Abstract—In this paper we discuss the current situation of the activity recognition research field. In our opinion activity recognition approaches of today focus too much on algorithmic solutions for detecting whatever activities. The overall motivation or a usage scenario is often missing. Evaluations are often using an artificial data set and thus, many recognition approaches have not been tried in any real application context. Our suggestion is to unify publications and workflows to improve the overall quality and increase the reusability of knowledge even across other research domains. We propose a standard set of sections each publication has to include: motivation, design, evaluation, and lessons learned. Especially the sections motivation and evaluation should be particularly aware of the humans’ needs. We think there is also a need for more different conferences or workshops focussing on one single application domain. Finally, researchers could start to accept activity recognition as a technology, enabling and supporting established or emerging fields like e.g. wellness informatics.

I. INTRODUCTION

The field of activity recognition research is already present at several conferences. However, compared to other established research fields, activity recognition is still struggling with non-standardized methodologies for e.g. experiments or publication outlines. One of the reasons for this issue might be the diversity of the activity recognition research community itself. In this paper we reflect on recently conducted research in the field of activity recognition and try to identify the causes of the problem. Finally we give four ideas on which next steps might take the activity recognition research to the next level.

II. WHAT IS THE COMMUNITY DOING WELL?

Browsing through publications dealing with topics on activity recognition, there are several strengths most of the papers have to offer. Combining all the papers, authors, and contributions, the biggest strength of the community is its diversity. Only one indicator for this variety is the application field for what activity recognition has been used: health, wellness, elderly, sports or simply the daily routine. The following lines give some impressions how wide the range of topics in the field of activity research is.

Almost every paper describes the architecture that have been used to conduct the activity recognition. While some of the architectural designs are rarely described, some of them are designed universally enough to fit multiple use cases.

Froehlich et. al. [4] designed and developed a flexible and reusable system to capture in situ mobile data. They described their architecture at a very detailed level, already presenting the XML structures they used for defining sensors and triggers. Another approach is to describe the architecture by splitting up components into modules and link them with each other by drawing lines (boxology). The paper by Raento et. al. [9] uses this way to describe their architecture. A description of the architecture is important, as it enables other researchers to understand and trace how components co-operate.

The technical challenges during the development or application of the activity recognition application are of particular interest as they allow to assess the limitations of the proposed technique or system. Kanjo et. al. [6] describe that they had problems with power consumption, GPS and Bluetooth connectivity, and storage of the gathered data in mobile devices’ memory. Some papers are addressing a single non-functional requirement or limitation, trying to solve this particular issue. Wang et. al. [11] investigated the issue on power consumption and propose an energy efficient mobile sensing system. Maurer et. al. [7] are also dealing with the power consumption problem and describe how they trade of between energy consumption and prediction accuracy.

Almost every paper describing a built system has a section evaluating some important parameters under whatever conditions. While most papers prefer to evaluate in a more formal way, setting a ground truth and comparing particular aspects, e.g. [5], some papers have particularly lovely evaluations. Consolvo et. al. [3] spend a huge effort in evaluating and analysing their UbiFit Garden system in a more or less non-formal way, involving many qualitative user statements. Both ways of evaluating a system have advantages and disadvantages. Inspired by the field of activity recognition there even exist creative new analysis techniques [2].

III. WHERE DOES THE COMMUNITY NEED TO IMPROVE?

Only a few papers are greatly balanced on describing the motivation, the used technologies, the conducted user study, and the gained experiences while conducting the research. On the opposite there are many papers focusing on only one of those aspects, while the used technology is probably the most preferred aspect. However, these publications could be much

better if taking into account *every* of the mentioned aspects. Activity recognition research is more than finding the ultimate algorithms or parameters!

Only focussing on algorithms, parameters, and technologies does not contribute to the field of activity recognition in a more meaningful and general way. The knowledge on how to select technologies under whatever circumstances is quite important and has not yet gained the needed importance within publications. Of course the selection depends on a given use-case or motivation. This is something which also has to be described in a good activity research paper. Something which is missing very often is a suitable evaluation. As outlined above, most papers rely on a formal evaluation, comparing their algorithms to any artificial ground truth. However, most publications do not take into account the user and the user's specific needs given within the motivation or usage scenario. Finally, what is often missing is the experience the authors gained dealing with the topic during their research. However, these experiences might beware other researchers from doing the same fault twice. Additionally these experiences can be re-integrated into future research as a non-functional requirement prior to the development process. These additional aspects can increase the reusability not only of the algorithms but also of the experiences.

IV. OUR BEST RECOMMENDATIONS

To understand our argumentation, we would like to quote the text given when asked to think about the recommendations:

Imagine you are reviewing an activity recognition paper. The authors have identified a set of activities they wish to automatically detect using a specific sensor system. They have collected training and test data.

From our point of view this description of a possible activity recognition paper is already way too focused on the technical parts. Activity research is more than having any activities, someone wants to detect by using whatever sensor systems. Why do the authors want to use activity recognition and what is the situation they want to observe? To conduct successful research in the field of activity recognition an in-detail motivation should be given as only this would allow reviewers to decide if the proposed system does make sense at all.

What we would therefore like to see in a paper is first of all a suitable motivation on why activity recognition has to be applied to the domain the author wants to analyze. This should include a reasonable set of activities their system should have to recognize to fulfill their motivation and requirements. Given this, the authors should describe how they identified where to place the sensors, which algorithms to choose, and how to adapt these algorithms. Then the authors should both validate the algorithms technical accuracy under common aspects, e.g. repeatability, but they should also check if the system performed well objectively. Therefore they could ask the participants of the study how they judge the systems decisions or how they felt during the study. Only these results can give valuable

feedback on if the activity recognition system is suitable for the motivated scenario or use case.

What we would not like to see in a paper is therefore easy to deduce. A missing motivation or usage scenario could already be a reason to reject a paper, if any technical considerations are not justified at all. What we would also dislike is if the validation of the system is done in an unfounded and non-comprehensible way – this is probably the regular case without any given use-case. Note that some artificial comparisons on completely different sensor platforms regarding power consumption or size do indeed fill two pages with text, but mostly do not enrich the papers contribution.

V. NEXT STEPS

There are four steps necessary to improve the overall quality of activity recognition publications and increase the global awareness for this field of research. The first step is to find a reasonable and suitable structure each paper investigating the field of activity recognition has to fulfill. An increased involvement of the humans' needs and wishes in the sections motivation and evaluation is the second step. The third step is to find sub communities, focusing on the special characteristics each application field has. The fourth step might be the most difficult one: establish activity recognition as technique and enabler for emerging research fields, e.g. wellness informatics.

There is a strong need to unify the global structure of a publication to increase the overall reusability, finally leading to a global acceptance of activity recognition as a technology enabler. This would also unify the expectations a reviewer has, when reviewing a activity recognition paper. Unifying a structure is in general nothing new. If any author wants to write about a conducted user study or evaluation, a typical structure consists of methodology, results, discussion. For activity recognition we would like to propose the following structure: motivation, design, evaluation, and lessons learned. However, this is only an idea which necessarily should be discussed within the community.

At the moment the involvement of the humans' needs and wishes is way too few. As there is no activity without any human being involved, researchers should take care and involve the humans sufficiently. Thus, a part of the motivation should be *why* exactly it is necessary to detect *which* activities of a human. In the evaluation part of the research it is actually unavoidable to reflect on the initial motivation and state how the motivation is fulfilled in which ways. As the only persons who can state this are those for who the activity recognition has been designed for, again those humans' impressions have to be involved into the evaluation.

The field of activity recognition already has a lot of different use cases and motivations. From our point of view, most of the fields have something very unique, making them a bit incomparable with other application fields. There should be more specialized conferences or workshops dealing with the specific requirements and problems each field has.

Activity recognition isn't a bare research field any more. Nowadays activity recognition can easily serve as enabler for

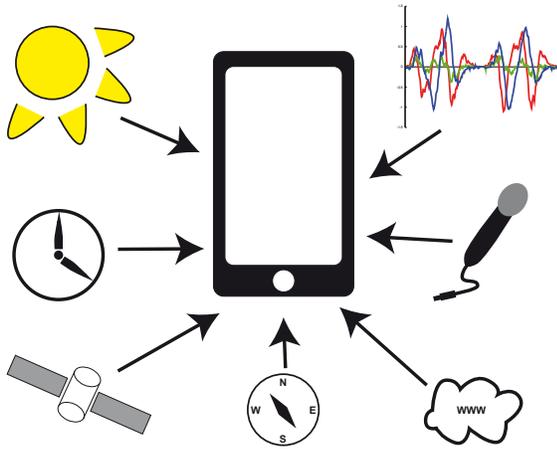


Fig. 1. One of our research ideas is to use a common smart phone for evaluation purposes. The inclusion of multiple sensor inputs can lead to some completely new evaluation parameters.

emerging fields, like e.g. Wellness Informatics. The community has to take this chance and refocus on the application of activity recognition as an established technique.

VI. OUR WORK

What leads us to activity recognition is two fold. First we have developed a gesture recognition library *wiigee*¹, which is able to recognize almost arbitrary movements recorded by an accelerometer [10]. This library gives us some basic understanding what is possible with even a single 3D acceleration sensor by applying machine learning technologies. On the other hand we are interested in the field of Human Computer Interaction (HCI), where technology often disappears to the background and user studies are very important.

Combining both fields leads us to the idea why not to apply activity recognition to the field of user studies. The approach to rely on a device to conduct a user study is in general nothing new; in common this is what is understood under the term *mobile living lab* [1]. However, compared to a technique called *Experience-Sampling Method* (ESM), we find it particularly interesting to rely on sensory input only (see Figure 1). This would allow to conduct user studies without any active experimenter involvement. We are sure that this approach can even lead the field of HCI to some new evaluation parameters far away from e.g. task completion time. Our most recent not yet published research showed that evaluation participants rate to have a logging device with them over the traditional observation by a human being.

Another interesting field where activity recognition is applicable and necessary is *wellness informatics*. The almost real time sharing of a human's personal health experience would allow revolutionary epidemiological analysis as well as an overall increase of the personal well-being [8]. The hidden activity recognition technology would probably not hamper any humans' feeling of well-being. We apply for this workshop

because we want to discuss what can take activity recognition to the next level and incorporate new and interesting ideas into our further research.

ACKNOWLEDGEMENTS

The authors are grateful to HaptiMap (FP7-ICT-224675). We like to thank our colleagues for sharing their ideas with us.

REFERENCES

- [1] Tom Broens, Henri ter Hofte, and Peter Ebben. A reference architecture for living lab measurement systems. In *Proceedings of the Mobile Living Labs Workshop in conjunction with Mobile HCI*, 2009.
- [2] Adrian K. Clear, Ross Shannon, Thomas Holland, Aaron J. Quigley, Simon A. Dobson, and Paddy Nixon. *Situvis: A visual tool for modeling a user's behaviour patterns in a pervasive environment*. In *Pervasive*, pages 327–341, 2009.
- [3] Sunny Consolvo, David W. McDonald, Tammy Toscos, Mike Y. Chen, Jon Froehlich, Beverly Harrison, Predrag Klasnja, Anthony LaMarca, Louis LeGrand, Ryan Libby, Ian Smith, and James A. Landay. Activity sensing in the wild: a field trial of ubifit garden. In *CHI '08: Proceeding of the twenty-sixth annual SIGCHI conference on Human factors in computing systems*, pages 1797–1806, New York, NY, USA, 2008. ACM.
- [4] Jon Froehlich, Mike Y. Chen, Sunny Consolvo, Beverly Harrison, and James A. Landay. Myexperience: a system for in situ tracing and capturing of user feedback on mobile phones. In *MobiSys '07: Proceedings of the 5th international conference on Mobile systems, applications and services*, pages 57–70, New York, NY, USA, 2007. ACM.
- [5] Tãm Huynh, Mario Fritz, and Bernt Schiele. Discovery of activity patterns using topic models. In *UbiComp '08: Proceedings of the 10th international conference on Ubiquitous computing*, pages 10–19, New York, NY, USA, 2008. ACM.
- [6] Eiman Kanjo, Steve Benford, Mark Paxton, Alan Chamberlain, Danae Stanton Fraser, Dawn Woodgate, David Crellin, and Adrain Woolard. Mobgeosen: facilitating personal geosensor data collection and visualization using mobile phones. *Personal Ubiquitous Comput.*, 12(8):599–607, 2008.
- [7] Uwe Maurer, Anthony Rowe, Asim Smailagic, and Daniel P. Siewiorek. ewatch: A wearable sensor and notification platform. In *BSN '06: Proceedings of the International Workshop on Wearable and Implantable Body Sensor Networks*, pages 142–145, Washington, DC, USA, 2006. IEEE Computer Society.
- [8] Jochen Meyer, Benjamin Poppinga, and Susanne Boll. Wellness 2.0 - sharing personal health experiences. In *Proceedings of the Wellness Informatics Workshop in conjunction with CHI*, 2010.
- [9] Mika Raento, Antti Oulasvirta, Renaud Petit, and Hannu Toivonen. Contextphone: A prototyping platform for context-aware mobile applications. *IEEE Pervasive Computing*, 4(2):51–59, 2005.
- [10] Thomas Schlömer, Benjamin Poppinga, Niels Henze, and Susanne Boll. Gesture recognition with a wii controller. In *TEI '08: Proceedings of the 2nd international conference on Tangible and embedded interaction*, pages 11–14, New York, NY, USA, 2008. ACM.
- [11] Yi Wang, Jialiu Lin, Murali Annavaram, Quinn A. Jacobson, Jason Hong, Bhaskar Krishnamachari, and Norman Sadeh. A framework of energy efficient mobile sensing for automatic user state recognition. In *MobiSys '09: Proceedings of the 7th international conference on Mobile systems, applications, and services*, pages 179–192, New York, NY, USA, 2009. ACM.

¹<http://www.wiigee.org/>