
The Implications of Mobile Assessments of Wellbeing for Liberty

Reuben Kirkham

OpenLab

Newcastle University

r.kirkham@newcastle.ac.uk

Abstract

It has recently become possible to compute assessments of mobile wellbeing, using Human Activity Recognition combined with inertial and locational data. When used properly, this could be of great assistance, for instance by the medical profession in more accurately tracking the effect of an individual's impairment. However, these systems can also potentially be used in assessments of mental capacity, and impliedly, whether or not someone should be deprived of their liberty. Depending on how this issue is addressed, it could have a significant negative impact upon the legal process and also reduce end user engagement with wellbeing technologies which are otherwise intended to help them.

Author Keywords

Disability; Mental Capacity; Wellbeing

ACM Classification Keywords

H.5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous

Introduction and Background

For most people, Mobile Wellbeing might simply amount to an app or system that helps them track some aspect of their mood or general health (e.g. a Fitbit). However, for a minority who have significant mental health disabilities, or impairments that might impact upon their

mental capacity (the cognitive ability to weigh facts and to take decisions about their day to day lives [1]), pervasive systems that assess or measure 'wellbeing' could be beneficial for treatment and management of their condition. However, there is a pre-condition in that these systems need to be designed and developed to take into account the individuals individual impairment. This is because individual impairments generally impact upon people on markedly different ways, even within a given 'named' impairment group.

This point is especially true for sensing systems based upon Human Activity Recognition (HAR), including those which track mood and behavior. These systems are based upon machine learning, which uses 'training data' from people performing a given set of actions in order to learn the underlying rule. However, if someone performs the actions in question in a different way from other people within the training data, then it is inevitable that a system would mischaracterize their behavior. Normally with these systems, especially those based upon inertial sensing (which is normally the only practical approach), there is no video or other data to evaluate the accuracy of the system post-hoc on a given individual.

There are some inherent problems that arise out of this situation. HAR systems cannot be relied upon to have the same accuracy levels across individuals (presuming that an appropriate statistical measure is identified in the first place, which is a challenge in and of itself – see [2]). There is little to prevent these systems being deployed in court to assist in determining the liberty of someone who is said to have restricted 'mental capacity'

without any argument being made on either side as to whether or not these systems are appropriate. On some occasions, it is possible that the use of these systems will not even be noted, for instance if an expert relies upon them without identifying that there is a HAR system in play. And, even if the risk of a HAR system is identified, there is little chance that the legal system will get it right, especially in a closed tribunal (see [3] for a discussion of this general risk)¹. As such, these potentially empowering systems need to be explored with caution.

Needless to say, if the conclusions produced by a HAR system are wrong and are relied upon in respect of deciding personal liberty, then this is seriously problematic and ought to be prevented by the designer of these technical systems. At the same time, there are potentially great benefits, in that these systems could empower end users to challenge deprivation of liberty. With some disabilities, existing assessments by professionals can be unscientific and lack objectivity, for instance in the case of London Borough of Hillingdon v. Steven Neary [2011] EWHC 1377 (COP), where unscientific behavioral logs (the person in this case had severe autism and engaged in 'problem behavior', which was logged on a day to day basis) were used as a basis for depriving someone of their liberty. A HAR based behavioral logger would have likely prevented this, showing that the concern cuts both ways: whilst people need to be protected from the inaccuracies of HAR systems, they could also use the assessment capabilities of these systems as a form of empowerment.

¹ In some very recent caselaw (Pyrrho Investments Ltd v MWB Property Ltd & Ors [2016] EWHC 256 (Ch)), the court used a rather unusual phrase for machine learning, describing it as

'predictive coding' and going on to fundamentally misunderstand the way that these systems operate.

Implications for the Design of Mobile Wellbeing Systems

It follows for the above narrative that there is a need to consider the risks and opportunities implied by the use of HAR in systems designed to track wellbeing in some form. This is a question of design, including how the risks and opportunities of HAR are appropriately communicated within a system. I therefore make some general observations on this topic as a series of starting point's (or precautions) to be considered:

Declaration of Human Activity Recognition

This point might seem obvious, but it is fundamental. The way that HAR works is rather different to a range of systems that the general public have previously encountered, especially in terms of how it can provide inaccurate measures. It should be observed that most people would assume that these systems are accurate, or mostly accurate. An activity tracker is an example of this: most people take the numbers produced by the system at face value and even if they didn't, they would assume that the values would be correct on a 'broad brush' level (i.e. +/- a given tolerance).

A clear communication of how a system can be inaccurate, including the fact that it relies upon HAR, is therefore essential. It means that people who wish to rely upon such a system are aware of its limitations. Whilst this cannot eliminate risk, it is a sensible step towards addressing the counterintuitive nature. At the least, it allows people the chance of knowing what they are dealing with.

Consider what Information is Stored

In some cases, it is possible to make it difficult or impossible for a HAR system to be used inappropriately in respect of mental capacity. This can be done by a

sensible data protection strategy of not storing information that is not relevant. For instance, a system need not retain the raw activity data which it used to compute the mental wellbeing of someone at a given point in time.

Taking such a step would help guard against perceived or actual misuse: to put this simply: if the original data is not available, or is only stored temporarily, rather than providing a full history, then it cannot be so easily misused. This consideration is line with the general expectation of modern data protection regimes, including Europe's GDPR. Whilst this point is in many respects 'common sense', it remains a point which is often ignored and therefore worth restating.

Appropriate Metrics

There are, needless to say, cases where HAR being used as part of an assessment would be of benefit, including the aforementioned Neary case. The challenge in these cases is to design appropriate metrics that minimize the risk of these systems being misconstrued. Ideally this would involve the collection of some test data (i.e. an annotated video) which would allow the effect of the system to be tested on that individual to ensure its accuracy.

There is another risk that should be considered. The more detailed a system purports to be (e.g. identifying events at a given point of time), the greater the risk of error by an evaluator. As such, metrics that work on the basis of counting events that fit into groups across different settings (e.g. someone who is in a care facility compared to at home with problem behavior) would be more appropriate than a system that simply states its predications at a given point in time.

DIY Systems

DIY systems can be particularly problematic, as I observed in my previous discussion of this issue in another pervasive computing context [4]. The issue is that these systems can be customized for the specific concerns of the user and thus enhance their autonomy, but at the same time, these systems are not necessarily implemented following modern data security and design practices. Indeed, there are a plethora of mobile apps that already assert that they track wellbeing, yet these are written by a range of developers following opaque protocols. Some of these may well be no better than Ouija boards. Although it is not possible to stop people developing these apps, there is a need to mitigate the risks arising from them, perhaps by providing a properly developed framework for adaption by DIY assistive technology developers.

Demonstrating safety to an end user

As soon as there is a case of a HAR system being misused in a legal context to assess someone's liberty, then this concern will become alive in the minds of people with mental health difficulties. Some would likely become concerned about the use of mood tracking technology generally and simply abandon the assistive technology in question. Plainly, mood tracking can be of great assistance, along with a broader range of digital health technologies. But it will only be of assistance if people with mental health disabilities are actually prepared to use such systems.

Accordingly, there is a need to distinguish between systems which can track an end user, and those which cannot. Achieving this would act as some kind of 'kite mark', enabling those with mental health difficulties to choose between systems which track them and those which do not. Such an approach would also accord with

existing human rights law, because it allows a choice by the disabled user, thus empowering them in line with Article 12 of the UN CRPD.

Conclusion

I have provided an overview of the potential problems that can arise if a system designed to track or promote mental wellbeing does not provide a sensible consideration of how HAR is used and communicated in a system. These arise when such a system is used in a formal assessment of wellbeing. Whilst the issues I raise are yet to be problematic on a wide scale, without action, they soon could be. As such, the design of mobile wellbeing systems which focus upon mental health should be improved in order to address the risks arising from their potential use to measure capacity and thus determine individual liberties for people with mental health related disabilities.

References

- [1] P. Gooding, A. Arstein-Kerslake, and E. Flynn, "Assistive technology as support for the exercise of legal capacity," *International Review of Law, Computers & Technology*, vol. 29, no. 2-3, pp. 245-265, 2015.
- [2] Ward JA, Lukowicz P, Gellersen HW. Performance metrics for activity recognition. *ACM Transactions on Intelligent Systems and Technology (TIST)*. 2011 Jan 1;2(1):6.
- [3] Meyerson MI, Meyerson W. Significant Statistics: The Unwitting Policy Making of Mathematically Ignorant Judges. *Pepp. L. Rev.*. 2009;37:771.
- [4] Kirkham R, Greenhalgh C. Social access vs. privacy in wearable computing. *IEEE Pervasive Computing*. 2015 Jan;14(1):26-33.