Unlocking the Black Box of Wearable Intelligence: Ethical Considerations and Social Impact

Lauri Tuovinen
Insight Centre for Data Analytics
Dublin City University
Dublin, Ireland

Biomimetics and Intelligent Systems Group
University of Oulu
Oulu, Finland
lauri.tuovinen@oulu.fi

Alan F. Smeaton
Insight Centre for Data Analytics
Dublin City University
Dublin, Ireland
alan.smeaton@dcu.ie

***Abstract*—Computational intelligence is making its way into a variety of popular consumer products, including wearable physiological monitors such as activity trackers and sleep trackers. Such products are very convenient for the user, but this convenience is the result of a trade-off that has ethical implications, since in almost all cases it denies the user access to their own raw data underlying the easy-to-understand analyses that the products generate for them. One problem with this is that the user is not made aware of the uncertainty of the conclusions or analyses drawn from the data; another is that it is difficult for the user to reuse his or her data in other contexts, such as to combine data from multiple sources. Even if the user did have full control of the data, this would only solve part of the problem, because most people do not have the special skills required to analyze such data. This overall problem could be solved through collaboration between the data owner and a data analysis expert, though this again introduces further problems, notably that of preserving the data owner's privacy. In this paper we analyze the aforementioned issues pertaining to the ethics of wearable intelligence, propose possible approaches to handling them, and discuss the potential social impact of the technology if the issues can be successfully overcome.**

I. Introduction

Some form of computational intelligence is nowadays found in many high-tech consumer products. One popular subset of these comprises wearable devices that record sensor data about the wearer's physiological state or movement and process these data to generate information that is meaningful to the wearer. This can include information on physical activity, health and wellness indicators, sleep or combinations of these. The sensors and the hardware and software required to process the data are packaged as a watch to be worn on the wrist or even a ring to be worn on a finger, making the device convenient to be worn at all times. The information generated from the raw data is presented in a form that makes it easy to track one's own performance with respect to quantities that are relevant to one's health and well-being, such as the number of steps taken during the day, the amount and quality of sleep had during the night, or one's recovery and readiness for the day based on heart rate metrics.

The convenience and simplicity of using these products, which we refer to as *wearable intelligence* (WI), has presumably played an important part in making them attractive to the general public, as we shall see later in Section II. However, despite their widespread popularity and adoption, using such WI is not without problems. The general rationale for purchasing and using an activity, health or sleep monitor is to be able to make better-informed decisions on one's lifestyle and behavior based on its outputs. For the device to fulfill this purpose, it is not enough that the outputs are understandable - they must also be reliable. The problem here is that there is necessarily a degree of uncertainty associated with the aggregated information as presented to the user, but making the user aware of this uncertainty would make it less straightforward for the user to interpret the information. Understandability and reliability as two well-known design goals are thus in conflict in a way that cannot be easily resolved.

The hiding of information on the reliability of WI device outputs from the user is, in fact, just one manifestation of a broader issue having to do with who controls the data generated by WI devices. Increasingly, it is held that the data should be controlled by the users who generate it, and there are various technical solutions being developed that would enable them to do so. Having control over their WI data would allow individuals to process the data in a more versatile fashion than is possible using only the tools provided by the product vendor, potentially increasing not only the reliability of the data but also their overall value by enabling the data to be analyzed to extract additional knowledge not provided by the device itself or by its associated software or cloud service. In the European Union, the right of individuals to access personal data collected from them is asserted by the Charter of Fundamental Rights [1] and the General Data Protection Regulation [2], which is a step

The work of Lauri Tuovinen is funded by the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement number 746837. The Insight Centre is funded by Science Foundation Ireland under grant number SFI/12/RC/2289, and is co-funded under the European Regional Development Fund.

in the right direction but not the same thing as the individuals themselves controlling the data.

Besides control of raw data, getting the maximum benefit out of WI data requires data analysis tools and expertise, and most people do not have either. This problem could be circumvented by facilitating collaboration between the owner of the data and someone who does have the required expertise, but this again raises further problems having to do with, for instance, ensuring a fair exchange among the collaborators and addressing any privacy concerns of the data owner. However, if these problems can be solved, the potential social impact of WI technology is considerable, because it represents a new way for people to take responsibility and to look after their own well-being, possibly resulting in measurable positive effects at the population level.

In this paper we examine some of the ethical and social implications of WI devices, focusing on health and wellness monitors. We posit that a long-term vision of individuals maximizing the benefit they get from such devices through data sharing and collaborative analysis is attainable, but as discussed above, there are several issues with connections to ethics that are blocking this and that need to be resolved before this can happen. The paper discusses each of these issues in turn, aiming to identify the fundamental problems underlying them and potential approaches to handling them. We conclude that in general, a successful approach will involve both value choices and technological advances. The principal contributions of the paper are listed below:

- A definition of WI as a subset of computational intelligence and an analysis of what makes it interesting from the perspective of ethics;
- An analysis of ethical issues associated with epistemic opacity and centralized control of WI data, with suggestions for how these could be alleviated;
- An empirical case study carried out using two different sleep trackers, demonstrating notable unexplained discrepancies between their outputs;
- A description of the concept of collaborative data analysis and an analysis of its challenges and potential impact when applied to WI data.

The remainder of the paper is organized as follows: Section II presents some essential background information and an overview of related work. Section III discusses the concept of epistemic opacity and how it pertains to WI devices. Section IV examines issues related to data control and ownership. Section V looks into the idea of collaborative data analysis, its potential benefits and the factors that currently make these benefits difficult to achieve. Section VI presents a critical discussion of selected topics, and Section VII concludes the paper.

II. Background

No longer an exclusive domain of fitness and quantified-self enthusiasts, wearable health and activity monitors are breaking into the mainstream. For example, a recent survey conducted among the adult population of Alberta, Canada [3] found that one-fifth of the sample own and use a wearable physical activity tracker. The companies that design and manufacture these devices cater for a wide range of market segments, and an entry-level product can now be bought for well under 100 USD. Mobile applications that use sensors commonly included in smartphones as data sources can even be downloaded free of charge, so for an individual who already owns a phone that is compatible with one of these free apps, there is not necessarily any additional cost for them to be able to start tracking their activity, health and/or sleep.

While on the subject of "smart" devices, it is worth considering the question of what justifies our use of the term "wearable intelligence". For example, a so-called smart TV may be basically a TV set with Internet connectivity and no features that could meaningfully be referred to as intelligent, so why is WI not similarly a misnomer? To answer this question, we need to take a closer look at the composition of WI devices and the theoretical basis of their functionality.

The essential components of a WI product are sensors, such as accelerometers, gyroscopes, optical heart rate sensors and body temperature sensors, and software for processing the sensor readings, some of which may be running on the device itself while other functions are implemented by an external app or cloud service. The intelligence lies in the software, in models that aim to predict the values of variables such as energy expenditure or sleep quality based on input variables derived from the sensor readings. In the case of energy expenditure, for instance, reference techniques such as the doubly labelled water method [4] cannot be implemented as convenient wearable devices, making it a practical necessity to estimate expenditure using a computational model instead. The model applied may be, for example, a linear regression model, a linear mixed model or an artificial neural network taking accelerometry data as its inputs [5].

Some of the outputs of WI products are not even measurable quantities in the usual sense but inherently dependent on the application of computational intelligence techniques. For example, several products are capable of recognizing specific activities of the wearer, such as walking, running or cycling; the reference in this case is that the wearer obviously knows his/her own activities, but there is no equipment that could be used to "measure" them. Therefore the only way to collect this information without depending on manual input is to train a classification model to recognize activities based on variables that are measurable; a wide variety of classifiers may be applied to detect activities in data generated by sensors commonly found in WI devices [6]. Besides being useful in themselves for the purpose of presenting the wearer with an automatically generated exercise log, the classification results can be used to improve the accuracy of energy expenditure estimation by selecting the model to be applied at a given point in time based on the current activity [7].

Given the rising popularity of fitness and sleep tracker devices and apps, it would appear that a lot of people are finding them useful, but there is a certain doubt concerning how well justified this perception is. The obvious question to

ask is whether the metrics generated by the devices accurately represent the phenomena they purport to measure; this question has been investigated in numerous studies, but because of the rate at which new device models are being developed and launched, peer-reviewed research which investigates the quality of the output metrics necessarily lags behind the market. However, the most recently published surveys indicate that the consistent accuracy of consumer wearables is not yet comparable to that of reference equipment [8], [9], and this should be a cause for concern.

The uncertainties associated with WI devices have previously been examined in [10], where a distinction is made between *input uncertainty*, referring to uncertainties concerning the reliability of the input data used by the devices' computational models, and *output uncertainty*, referring to difficulties experienced by users when trying to assess the significance of the results generated by the models. A third type of uncertainty, called *functional uncertainty*, is also introduced, referring to uncertainties concerning who has access to the users' data, how they use the data and what for. As noted in [10], privacy concerns are concrete manifestations of this type of uncertainty, and this appears to be the one aspect of the ethics of WI on which there is already a considerable amount of published research, e.g. [11]–[15].

The focus on privacy of data in the context of WI is hardly surprising, given that it has been an active issue in computer ethics since well before the emergence of consumer wearables. This topic is, of course, important, but from another perspective it is just one aspect of the broader issue of who owns and controls the data generated by these devices. Typically the data are uploaded to the device manufacturer's cloud service, where the user can view them and possibly export them, but this is not equivalent in either principle or practice to the user him-/herself controlling the data. In the decentralized Web (DWeb) community there has been some interesting work aiming to enable individuals to become their own data controllers, such as the Solid platform [16], which was recently released to the public.

For an ordinary person to be able to truly control his/her data would be a major step forward, but there are also substantial problems to be solved concerning data analysis. Special software platforms have been developed to facilitate data sharing and collaborative development of data analytics solutions; some are already available as commercial Web services (e.g. [17], [18]), while others are more experimental (e.g. [19], [20]). However, the problem with all of these is that they have not been designed specifically to accommodate non-expert users, and therefore they do not adequately address the special requirements introduced by the presence of such non-expert users in the collaboration process. Dealing with this problem will require a multidisciplinary approach involving certain ethical considerations, as we shall observe in Section V.

Besides the problems of collaboration, Section V also looks at the opportunities. Previously, this theme has been studied in [21], which surveys ethical challenges and opportunities pertaining to lifelogging technologies. WI devices are a significant subset of the range of data sources that can be used for lifelogging, and consequently, both the challenges (e.g. infringements of privacy, shortcomings of the technology) and the opportunities (e.g. personalized services, health benefits) overlap with those associated with WI. As demonstrated by the survey in [22], using such technologies for ethically commendable purposes – in the case of the survey, helping people with dementia – can give rise to a wide range of further ethical issues, but exploring these is outside the scope of this paper.

III. Epistemic Opacity of Wearable Devices

The concept of epistemic opacity was defined in [23] as the quality of a computational process that makes it impossible for an observer to know all of its epistemically relevant elements. For the purposes of this paper, the epistemically relevant elements can be informally defined as those elements of the process that the observer needs to know in order to understand why the process arrives at a given conclusion when supplied with a given set of inputs. The context in which the definition is given is a discussion of the use of computer simulation in science, the implication being that relying on simulations as a source of scientific knowledge raises philosophical issues because it renders a significant portion of the discovery process impenetrable to inspection and evaluation by human cognition.

A similar argument can be made about the use of computational intelligence techniques, especially so-called black-box models that are by their nature particularly difficult to explain in terms of why they yield the results that they do. The opaque and probabilistic nature of such models means that there is necessarily some uncertainty concerning the reliability of metrics generated by WI devices using them: the numbers presented to the user are not direct measurements but approximations computed from proxy variables, and the models used to compute them are never perfect. However, the user of the device is not made aware of this uncertainty. Instead, what is essentially the best (computational) guess of the algorithms used to process the sensor readings is presented as if it were an objective measurement.

From the user's point of view, the hiding of these forms of uncertainty adds another level of epistemic opacity on top of the inherent opacity of the algorithms. The user is not only unable to understand or explain the numbers displayed by the device or the associated software, but also unable to tell when it would be useful to seek an explanation, except in cases where the outputs of the device are in clear conflict with the user's own knowledge and experience. How problematic this is in practice depends on circumstances; for example, as pointed out in [10], some users may be mainly interested in tracking their progress over time, in which case it is not necessarily a problem if the absolute figures are not always accurate, as long as the long-term trends are.

Vendors of WI products are careful to point out that the devices are not medical instruments and should not be treated as such, but this disclaimer tends to be overshadowed by

the more prominent and emphatic marketing message that depicts the products as providing useful information to support a healthy lifestyle. In fact, in many cases they provide not just information, but also recommendations on actions to be taken based on the information. The risk of any serious harm being caused by these recommendations may be low, but it is nevertheless somewhat problematic that the quality of WI device outputs is not regulated; for instance, under guidance issued by the United States Food and Drug Administration (FDA) [24], most fitness trackers and sleep trackers can be classified as low-risk general wellness products, which the FDA does not intend to subject to requirements that medical products must fulfill. This leaves it entirely up to the user to decide whether the recommendations are to be trusted, yet the user is not given all the information that would be relevant in making such a decision.

There is no obvious solution to this problem, because as we have already observed, there are conflicting interests at work here. It is far from evident how information about the uncertainty of the metrics computed by a WI device should be presented such that it is meaningful to the user and not confusing, and thus the overall effect of including this information in the outputs of the device might well be to make the main outputs more difficult to understand. This would arguably improve the reliability of the device, but would also damage its usability as a product intended to provide a simple way to track variables relevant to one's health and well-being. Even simple things like a graph with error bars reflecting a confidence level in the values can be difficult for the non-mathematically minded person to fully comprehend [25].

One way to reduce the opacity of the information generated by WI devices would be to expose the underlying algorithms, but this is also a problematic proposal for a number of reasons. Understanding such algorithms requires highly specialized knowledge, so for the average WI user it would be not helpful at all to know how the input data are being processed to generate the results displayed by the device. Exposing the algorithms would, of course, enable independent evaluators to scrutinize them and inform the public of their findings; a parallel may be drawn here with computer security, where the publicity of designs, protocols and even source code is considered good practice, as opposed to so-called security through obscurity, which relies on the secrecy of such information and is largely discredited [26].

The analogy here is based on the idea that with both WI algorithms and security algorithms, openness contributes to trustworthiness. It should be kept in mind, however, that the trustworthiness of algorithms used to, say, make online banking transactions secure is considerably more critical than that of algorithms used in non-medical devices to estimate energy expenditure or sleep quality. Therefore, given that these algorithms are sometimes important business secrets for the producers of WI devices, enforcing their disclosure would be a drastic measure. Still, it is worth thinking about if the producers could somehow be incentivized to partially expose their designs and/or implementations to facilitate some form of external scrutiny that would boost their trustworthiness significantly.

There is a stronger argument to be made in favor of disclosing the underlying data, not to the public but to each individual WI user, because while the algorithms are unquestionably intellectual property of the company that created them, the ownership status of the data is a much more ambiguous issue. This would still leave the problem that the data by themselves would not be any more meaningful to the average user than the algorithms, but there are some research directions suggested in [10] that would help change this. One is that the user should be provided with access to confirmatory independent evidence supporting the conclusions presented by the device; another is that the communication of the uncertainties associated with the conclusions should be tailored to the specific requirements of the context in which they are used.

A particularly interesting suggestion found in [10] is that there should be a way to preserve the *provenance of uncertainty* when the outputs of a WI device are exported and used in another context. In other words, it should be possible to export not just the outputs but also some form of metadata representing the uncertainties associated with them. This would be important when, for instance, WI devices are used for data collection in scientific research, because knowledge of the uncertainties could affect the evaluation of the research results considerably, yet at the moment there is no choice but to accept the device outputs as the ground truth.

We can easily expand this idea and consider what other additional information the user of a WI device should be able to export besides the provenance of uncertainty. From a certain point of view, everything the device records about the user, up to and including raw sensor readings, is the user's personal data. In fact, there is a clear rationale for why the user should have access to his/her own raw data; to illustrate, we conducted a small experiment where one individual used two different sleep tracking products simultaneously over a period of approximately three months in order to see how closely their outputs correlate. One device was the S+ contactless sleep tracker from ResMed[1] and the other was the Ōura ring, billed as the "most accurate sleep and activity tracker"[2]. Each device processes the raw data that it senses (motion, light intensity in the room, room temperature and noise level in the case of the S+ and movement, body temperature and heart rate in the case of the Ōura ring) into an aggregate score for the previous night's sleep, both in the range 0..100.

During the experiment period there were some interruptions, resulting in a dataset covering 75 days altogether, with some points of non-contiguity. The results, shown in Figure 1, indicate that the two devices do broadly agree with each other in that their curves do tend to rise and fall with each other, but there are some days where the devices contradict, highlighted by the light gray filled region between the curves. Prominent discrepancies can be seen, for example, between days 30 and

[1]https://www.resmed.com/us/en/consumer/s-plus.html

[2]https://ouraring.com/

35 as well as between days 71 and 75. Furthermore, the scores from the two devices seem to be calibrated differently, suggesting that they use different algorithms to calculate their sleep quality/efficiency values, which of course they do since they have data from different sensors. How the respective companies compute their sleep scores is proprietary information, but in the spirit of providing open and transparent access to our own data, should companies not provide us with our own raw data so we can use an independent third party to calculate sleep score values and not be bound to how a single company does this?

WI companies may argue that retaining the raw data would be impractical, but if we accept the argument that the user, not the company, owns the data, then this should be a decision made by the user, not the company. Practical issues aside, the principle that individuals should have control over what happens to their personal data stands in any case and is now affirmed by legislation such as the General Data Protection Regulation (GDPR) of the European Union, which we mentioned briefly earlier in Section I. If WI companies are not willing to make provisions for the storage and exportation of all the data that may be of interest to the users, then it should be possible for the users to obtain the data and source the required storage from another provider. We will examine the implications of this in the next section.

IV. Whose Data Is It Anyway?

The usual way in which the user of a WI device gains access to the data generated by the device is that the device uploads them to a cloud service provided by the company that created the device. By logging in to the service, the user can view the data, along with a variety of higher-level analyses computed from them. However, should the user wish to export and reuse the data, it is largely up to the benevolence of the company how much of the data can be exported and how convenient the exportation procedure is for the user. Normally it is the processed or analyzed data that are exported, and almost never it is the raw, unprocessed data from the sensors.

In theory, the GDPR [2], having become effective in May 2018, has changed the situation in favor of WI users within the European Economic Area. According to Article 4 of the regulation, which is now the law, any data that can be connected with a specific identifiable natural person are considered personal data, and Articles 12–22 specify certain rights that are guaranteed for the individuals concerned, referred to in the regulation as *data subjects*. Among these data subject rights are the right to get a copy of the data from the *data controller* (in the case of WI data, usually the provider of the cloud service) and the right to have the data transmitted to another controller.

Laws such as the GDPR are important, as they affirm and enforce the principle that the data subjects are the owners of their personal data and have, as a rule, the right to determine how their data are used by others. However, the principle alone is not enough – for the data subjects to be able to actualize the full potential of their data, it must be practical for them to access the data at any time and to process them without any limitations imposed by the data controller. Submitting GDPR requests to the data controller is far too unwieldy a mechanism for this purpose, so a different approach is needed to end the dependence of WI users on what WI device producers are willing to allow them to do with their data.

The fundamental problem with the status quo is that even under the GDPR, data controllers are largely free to determine the practicalities of how they share the data being processed with the data subjects. WI users who are interested in reusing their own data can choose devices and services that make it convenient to do so, but this assumes that there are such devices and services available, and that if there are, they are not inadequate in some other respect such as battery life or on-device storage capacity. To eliminate this problem, we argue that it is necessary to change the status quo to reflect the principle that individuals are the owners of their personal data. In other words, instead of companies controlling the data and sharing it with the individuals concerned, it should be the individuals controlling the data and sharing it with companies of their choice.

To achieve this, the storage of WI data would have to be decoupled from the generation of said data, whereas currently these two functions generally come as a package that an individual user can take or leave but not alter the terms of. Technical solutions such as the Solid platform [16] are necessary enablers, but they are not going to make any real difference unless WI companies can be persuaded to let go of their role as data controllers. Therefore what is needed is a change of values whereby WI users being in control of their data becomes the new norm, and there is likely to be considerable resistance to such a change, since controlling large quantities of personal data is a major business asset in today's economy.

The MyData declaration [27], particularly the principles posited in its third section, provides a good overview of the kind of change we are proposing. Bringing about a change like this will not be a simple matter, but there are several forces that can help it happen:

- Grassroots pressure: as awareness of the problems with the status quo and the benefits of the MyData model increases, demand for products and services that bring these benefits to consumers will also increase;
- Regulation: as enabling technology becomes more widely available and less costly to deploy, data controllers can be required to make it more convenient for data subjects to exercise their rights;
- New business opportunities: if individuals can truly control their personal data and have them all stored in one place, it will be possible to develop new kinds of services based on integrating data from multiple sources.

Transferring control of personal data to the data subjects would go a long way toward dispelling the privacy concerns currently associated with the processing of personal data by companies, as it would enable the subjects to fully control the boundaries of their privacy according to their own individual

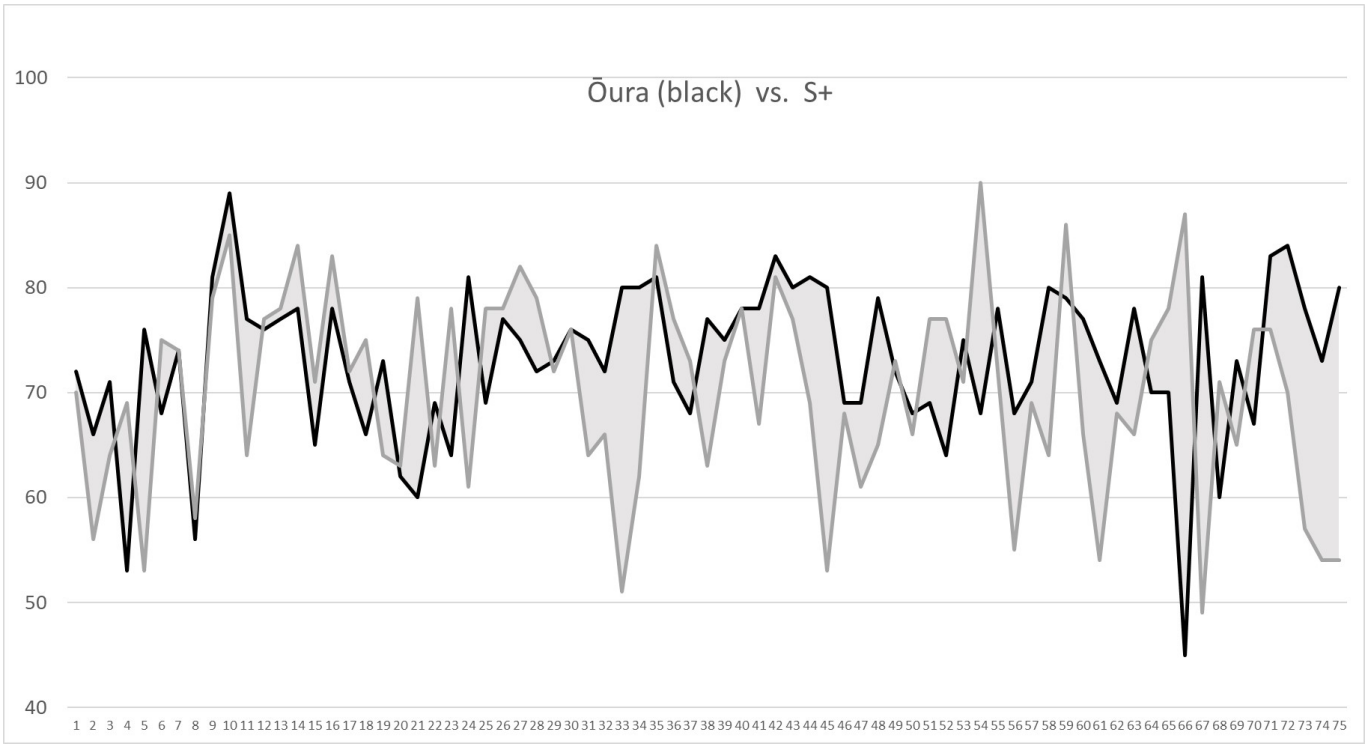

Fig. 1. A day-by-day comparison of the sleep scores output over a period of 75 (non-contiguous) days by two sleep trackers, the ResMed S+ and the Ōura ring. The light gray filled region between the two curves highlights instances of substantial discrepancy.

preferences. Furthermore, it would enable people to maximize the benefit they get from their personal data, including WI data, through sharing and collaboration on their own terms. The ethical challenges and potential social impact of this are discussed in the next section.

V. Challenges and Potential of Data Sharing and Collaborative Analysis

Shifting control of personal data from companies to individuals would be a major positive step, but for most people, refining their data into useful knowledge about themselves would require some form of collaboration with others, because the range of data analysis tools that can be applied effectively without special expertise is limited. The mobile apps and online dashboards available to owners of WI devices fit this description, but the relationship between the device user and the application provider cannot be properly described as collaborative, because there is little or no interaction between the two parties involved. In true collaboration, the owner of the data would have an active role in determining the objectives of the collaboration and the terms and conditions under which they are pursued.

To give a simplified description of how collaborative analysis would work: Initially the data owner carries out a search on the collaboration platform, looking for users with expertise relevant to what he/she is hoping to achieve. Experts can be identified based on information from multiple sources, including self-assessment, referrals and known past activities. Having found an expert, the data owner invites him/her to collaborate and the two proceed to negotiate on what they will do and what the contributions and expectations of each collaborator are. They will then work together using the collaboration platform to produce what the data owner is looking for, which may be some feedback on a given set of data or a piece of program code that the data owner can use to carry out a specific analysis task without having to involve the expert again. As a concrete example of what kind of analysis might be carried out collaboratively, discovery of periodicities from longitudinal data [28] is something that off-the-shelf WI products do not typically provide.

The ethical issues that arise from collaborative analysis of personal data have to do mainly with how the collaborators can establish a mutually satisfactory relationship. In a collaboration between an expert and a non-expert, there is an obvious imbalance of power where, even though the non-expert is technically empowered to negotiate an acceptable set of terms, he or she does not necessarily have sufficient knowledge to understand all the relevant privacy implications, for example. Collaboration platforms should therefore be designed to support non-expert users in achieving and maintaining an awareness of such implications, so that they can be sure they are not sharing more data with their collaborators than they are actually comfortable with or intend to. In [29], we proposed that this could be accomplished by using a collaborative data analysis ontology and a reasoner to detect privacy issues and to identify potential ways to resolve them.

Another issue here that has an ethical dimension is trust. Assuming that the non-expert collaborator has the competence, either independently or with the support of the collaboration platform, to negotiate with an expert collaborator on equal terms, that still leaves the question of how the non-expert – or the expert for that matter – can be sure that the outcome of the negotiation will be honored. To frame this in terms of privacy again, the non-expert may be able to specify which data are to be shared with the expert, but this alone does not guarantee that the data will not be used by the expert in ways that the non-expert did not intend. On the other hand, the expert could have trust issues to be addressed as well, if, for instance, the collaboration involves sharing some data analysis code with the non-expert. Ideally, the collaboration platform should provide mechanisms for encoding the result of the negotiation in a machine-readable format and enforcing it automatically.

Indeed, the perspective of the expert collaborator is not to be neglected, even though the concerns of the non-expert are more evident. It is worth considering, for example, what the motivation of the expert is for participating in the collaboration; it could be that the non-expert is willing to pay for the service, but alternatively, it could be that the expert is interested in gaining access to the non-expert's data, which could be the case if the expert is, for instance, looking for research data. In the latter case it may prove considerably more complicated to negotiate and guarantee an exchange that satisfies all of the legitimate expectations of both parties, because the expert is not simply providing a service for the data owner but has interests of his/her own concerning the extraction of knowledge from the data.

If the obstacles in the way of effective collaborative analysis of WI data can be overcome, the most immediate benefits will come to the individual data owners, but there is also potential for a significant positive impact at the population level. If the popularity of self-measurement using WI devices continues to increase, and if the users of WI products can be provided with everything they need to make the best possible use of their data, the potential net effect on public health would be beneficial to society as a whole. In fact, sharing of WI data with medical practitioners is one of the most obvious collaboration scenarios, and also one of the least problematic in terms of privacy and trust issues, thanks to the special nature of the doctor-patient relationship.

Another way in which widespread collection and sharing of WI data could contribute to the common good is that it would enable large and rich datasets to be built for purposes such as scientific research. In a certain sense this is already being done by the WI companies currently in business, and the centralized model of data control is undeniably an efficient way to accumulate data in large quantities, but the availability of these datasets for research is entirely up to whether the companies controlling them are willing to release them. In the decentralized model it would be up to each individual person to select the purposes for which his or her data may be used, eliminating the business interests of WI companies as an obstacle to data sharing. Datasets built in this way could be considerably more diverse than those currently held by WI companies, because they could integrate data from multiple sources for each data subject. However, here again there are substantial privacy and trust issues to be addressed, such as how to achieve irreversible anonymization of the data.

Finally, we should also consider the possible ill effects that may come about if collection and sharing of WI data becomes the norm. It is conceivable that there will be increasing pressure for these data to be shared with, for example, insurance companies or employers, to be used for decision-making that may have negative effects on the individuals concerned. In theory, there is no problem as long as the sharing is genuinely voluntary, but there is a problem if the person in question does not fully understand how the data will affect the decisions, and likewise if refusing to share the data automatically results in some kind of penalty. There are already many examples of health insurers offering self-tracking bonus programs [30], and while in these the customers are never penalized, they do raise some questions, as the authors point out, concerning the availability of these programs to customers who are disabled or otherwise disadvantaged. Furthermore, the conceptual leap from here to using self-tracking data, or the refusal to share such data, in determining the availability and cost of insurance for a given customer is not very long, so this is a prospect that needs to be taken seriously.

VI. Discussion

Much of the current discourse on the ethical and social implications of artificial intelligence (AI) is concerned with how the application of AI in various sectors of society affects people and how its harmful effects can be averted or alleviated. The topics covered in this paper overlap with this discourse to some extent, but there is a key difference in the way the people concerned are viewed, not simply as passive objects affected by AI but as active subjects who can use AI to produce effects of their own desiring. WI is, at its heart, a technology that enables people to do this, even though the way it is currently implemented falls short of its full potential by a considerable margin.

In the previous section we implicitly assumed, for the sake of simplicity, collaboration involving two individuals, one data owner and one data analysis expert. Doing so enabled us to identify the types of ethical issues that are likely to arise in a fairly straightforward fashion, but it is worth noting that this is not necessarily a realistic or comprehensive model of what collaboration on WI data entails. In the real world, collaborations may turn out to be considerably more complex, involving potentially any number of data owners, domain experts and technology experts acting either as individuals or as members of an organization. This, in turn, may considerably complicate the known ethical issues, as well as introduce some entirely new ones.

Another implicit assumption we have made that it is worth taking a critical look at is that people who use WI products are going to be willing to engage in this type of collaboration.

With so many problems that have yet to be solved before the vision can become reality, the most we can say for now is that we do not know that they are not going to be willing. If asked about it now, many current WI users might well be skeptical, so a definitive answer to this question may not be forthcoming until all the problems have been solved and the actual nature and scale of the benefits to be had from collaborative analysis become evident.

One thing that does appear to be clear enough from currently available evidence is that as a general rule, people are not averse to sharing their personal data provided that they are getting something back in return, as a reward. The success of social media platforms such as Facebook depends entirely on this willingness to share, free access to the platform being what the users receive in return for their data. However, the long-term impact of high-profile data scandals such as the recent one involving Facebook and Cambridge Analytica remains to be seen; if the companies involved fail to respond to these in a convincing manner, this may significantly undermine the trust of the public in such companies as responsible controllers of personal data.

If trust in centralized control of personal data deteriorates, decentralization of control may begin to seem more attractive, paving the way for one of the main prerequisites of collaborative analysis. On the other hand, it may also be that if the current privacy concerns continue to grow in magnitude, sharing of personal data will become less attractive regardless of who is controlling the data and how great the potential benefits are. One of the major challenges of collaboration will therefore be to provide assurances that with decentralized control, the privacy of personal data is genuinely stronger than in the currently prevalent centralized model.

VII. Conclusion

In this paper we explored the ethical and social implications of using wearable intelligence (WI) devices and software to measure and analyze phenomena that affect the user's health and well-being, such as his or her physical activity and sleep. WI products are popular, but they are not always reliable because there is necessarily some uncertainty associated with their outputs. However, the user is not made aware of this uncertainty, even though having this knowledge would be important in enabling the user to judge the usefulness of the metrics and recommendations generated by the product. Another problem is that the user's ability to control the data recorded is often limited, which not only raises privacy issues but also restricts the user's options in utilizing the data to their full potential.

If individuals were to gain control of their own WI data, a promising approach by which they could benefit from the data is collaborative analysis. Here the data would be shared with experts chosen by the owner of the data and processed under terms negotiated by the collaborators, with the data owner being in full control of his or her privacy preferences. Collaboration like this raises further issues having to do with how the privacy of the data owner can be guaranteed and how the collaborators can trust one another, but if these can be overcome, there is potential here for a significant positive social impact in areas such as public health and scientific research. On the other hand, negative effects are also a possibility if WI users begin to be pressured to share their data for decision-making processes that may result in harmful consequences to them. The challenges associated with WI cannot be dealt with by engineering alone but will require contributions from other disciplines, including ethics.

References

[1] European Union, "Charter of fundamental rights of the European Union," *Official Journal of the European Union*, no. C 326, Oct. 26, 2012. [Online]. Available: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:12012P/TXT (visited on 01/17/2019).

[2] ——, *Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation)*, Apr. 27, 2016. [Online]. Available: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:02016R0679-20160504 (visited on 01/17/2019).

[3] S. Macridis, N. Johnston, S. Johnson, and J. K. Vallance, "Consumer physical activity tracking device ownership and use among a population-based sample of adults," *PLOS ONE*, vol. 13, no. 1, e0189298, 2018.

[4] K. R. Westerterp, "Physical activity and physical activity induced energy expenditure in humans: Measurement, determinants, and effects," *Frontiers in Physiology*, vol. 4, 2013. DOI: 10.3389/fphys.2013.00090.

[5] A. H. K. Montoye, M. Begum, Z. Henning, and K. A. Pfeiffer, "Comparison of linear and non-linear models for predicting energy expenditure from raw accelerometer data," *Physiological Measurement*, vol. 38, no. 2, pp. 343–357, 2017.

[6] Y. Saez, A. Baldominos, and P. Isasi, "A comparison study of classifier algorithms for cross-person physical activity recognition," *Sensors*, vol. 17, 2017. DOI: 10.3390/s17010066.

[7] B. Cvetković, R. Milić, and M. Luštrek, "Estimating energy expenditure with multiple models using different wearable sensors," *IEEE Journal of Biomedical and Health Informatics*, vol. 20, no. 4, pp. 1081–1087, 2016.

[8] L. M. Feehan, J. Geldman, E. C. Sayre, C. Park, A. M. Ezzat, J. Y. Yoo, C. B. Hamilton, and L. C. Li, "Accuracy of Fitbit devices: Systematic review and narrative syntheses of quantitative data," *JMIR mHealth and uHealth*, vol. 6, no. 8, e10527, 2018.

[9] L. Scalise and G. Cosoli, "Wearables for health and fitness: Measurement characteristics and accuracy," in *2018 IEEE International Instrumentation and Measurement Technology Conference (I2MTC)*, 2018, pp. 1–6.

[10] B. Knowles, A. Smith-Renner, F. Poursabzi-Sangdeh, D. Lu, and H. Alabi, "Uncertainty in current and future health wearables," *Communications of the ACM*, vol. 16, no. 12, pp. 62–67, 2018.

[11] G. Paul and J. Irvine, "Privacy implications of wearable health devices," in *Proceedings of the 7th International Conference on Security of Information and Networks*, 2014, pp. 117–121.

[12] A. Aktypi, J. R. C. Nurse, and M. Goldsmith, "Unwinding Ariadne's identity thread: Privacy risks with fitness trackers and online social networks," in *Proceedings of the 1st International Workshop on Multimedia Privacy and Security*, 2017, pp. 1–11.

[13] C. Lidynia, P. Brauner, and M. Ziefle, "A step in the right direction – understanding privacy concerns and perceived sensitivity of fitness trackers," in *Advances in Human Factors in Wearable Technologies and Game Design*, T. Ahram and C. Falcão, Eds., ser. Advances in Intelligent Systems and Computing. Springer International Publishing, 2018, pp. 42–53.

[14] I. Torre, O. R. Sanchez, F. Koceva, and G. Adorni, "Supporting users to take informed decisions on privacy settings of personal devices," *Personal and Ubiquitous Computing*, vol. 22, no. 2, pp. 345–364, 2018.

[15] M. Zimmer, P. Kumar, J. Vitak, Y. Liao, and K. C. Kritikos, "'There's nothing really they can do with this information': Unpacking how users manage privacy boundaries for personal fitness information," *Information, Communication & Society*, 2018. DOI: 10.1080/1369118X.2018.1543442.

[16] E. Mansour, A. V. Sambra, S. Hawke, M. Zereba, S. Capadisli, A. Ghanem, A. Aboulnaga, and T. Berners-Lee, "A demonstration of the Solid platform for social Web applications," in *Proceedings of the 25th International Conference Companion on World Wide Web*, 2016, pp. 223–226.

[17] Mode Analytics, Inc. (2019). Mode website, [Online]. Available: https://modeanalytics.com/ (visited on 01/07/2019).

[18] Dataiku. (2019). Dataiku website, [Online]. Available: https://www.dataiku.com/ (visited on 01/07/2019).

[19] M. Deshpande, D. Ray, S. Dixit, and A. Agasti, "ShareInsights: An unified approach to full-stack data processing," in *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data*, 2015, pp. 1925–1940.

[20] E. Kandogan, M. Roth, P. Schwarz, J. Hui, I. Terrizzano, C. Christodoulakis, and R. J. Miller, "LabBook: Metadata-driven social collaborative data analysis," in *2015 IEEE International Conference on Big Data (Big Data)*, 2015, pp. 431–440.

[21] T. Jacquemard, P. Novitzky, F. O'Brolcháin, A. F. Smeaton, and B. Gordijn, "Challenges and opportunities of lifelog technologies: A literature review and critical analysis," *Science and Engineering Ethics*, vol. 20, no. 2, pp. 379–409, 2014.

[22] P. Novitzky, A. F. Smeaton, C. Chen, K. Irving, T. Jacquemard, F. O'Brolcháin, D. O'Mathúna, and B. Gordijn, "A review of contemporary work on the ethics of ambient assisted living technologies for people with dementia," *Science and Engineering Ethics*, vol. 21, no. 3, pp. 707–765, 2015.

[23] P. Humphreys, "The philosophical novelty of computer simulation methods," *Synthese*, vol. 169, no. 3, pp. 615–626, 2009.

[24] U.S. Department of Health and Human Services, Food and Drug Administration, *General wellness: Policy for low risk devices, Guidance for industry and Food and Drug Administration staff*, Jul. 29, 2016. [Online]. Available: https://www.fda.gov/downloads/medicaldevices/deviceregulationandguidance/guidancedocuments/ucm429674.pdf (visited on 01/17/2019).

[25] P. Shah and J. Hoeffner, "Review of graph comprehension research: Implications for instruction," *Educational Psychology Review*, vol. 14, no. 1, pp. 47–69, 2002.

[26] J.-H. Hoepman and B. Jacobs, "Increased security through open source," *Communications of the ACM*, vol. 50, no. 1, pp. 79–83, 2007.

[27] A. Poikola, D. Kaplan, and T. Mällo. (2017). Declaration of MyData principles, [Online]. Available: https://mydata.org/declaration/ (visited on 01/04/2019).

[28] M. P. Buman, F. Hu, E. Newman, A. F. Smeaton, and D. R. Epstein, "Behavioral periodicity detection from 24 h wrist accelerometry and associations with cardiometabolic risk and health-related quality of life," *BioMed Research International*, 2016. DOI: 10.1155/2016/4856506.

[29] L. Tuovinen and A. F. Smeaton, "Ontology-based negotiation and enforcement of privacy constraints in collaborative knowledge discovery," Presentation at the 2nd International Workshop on Personal Analytics and Privacy (PAP 2018), Dublin, Ireland, Sep. 10, 2018, [Online]. Available: http://kdd.di.unito.it/pap2018/papers/PAP_2018_paper_2.pdf (visited on 01/17/2019).

[30] M. Henkel, T. Heck, and J. Göretz, "Rewarding fitness tracking – the communication and promotion of health insurers' bonus programs and the use of self-tracking data," in *Social Computing and Social Media. Technologies and Analytics*, G. Meiselwitz, Ed., ser. Lecture Notes in Computer Science, Springer International Publishing, 2018, pp. 28–49.