

The impact of smart wearables on the decisional autonomy of vulnerable persons

Niël H. Conradie, Sabine Theis, Jutta Croll, Clemens Gruber und Saskia K. Nagel

Abstract

Smart wearable technologies have seen an explosive growth over recent years, with some research indicating that the wearable technology industry is expected to grow from USD 24 billion today to over USD 70 billion in 2025. This proliferation has extended across disparate domains, ranging from medical applications and fitness and social technologies to military, industrial, and manufacturing applications. As with any emergent technology, these wearables present opportunities for our moral benefit as well as moral challenges to be addressed. A crucial dimension of this discussion is the ethical evaluation of the impact of smart wearables on the autonomy of human decision-making. Nowhere is this a more pertinent concern than when dealing with persons uniquely vulnerable to autonomy infringement. This contribution, undertaken from an explicitly normative and ethical perspective, investigates the potential impact of smart wearables on various dimensions of the decisional autonomy of vulnerable persons.

1. Introduction

Smart wearable technologies – by which we mean here, wearable technologies that have the capacity for the collection and algorithmic processing of data, often including a wireless connection to a user network or the internet, in order to produce corrective output via means of actuators integrated into the worn item¹ - have seen an explosive growth over recent years, with current predictions being that this trend is set to accelerate (Seneviratne et al 2017; Dian et al 2020). Indeed, some research indicates

¹ This definition is our own but is informed by several foregoing attempts – see Viseu 2003; Bower and Sturman 2015; Xue 2019 as examples, and Niknejad et al 2020 for a good overview of other existing definitions.

that the wearable technology industry is expected to grow from USD 24 billion today to over USD 70 billion in 2025. And, unlike other sectors, the COVID-19 pandemic proved to be an accelerator for many of the key trends driving wearable technology, such as the move towards remote patient monitoring and digital healthcare as well as the current embrace of fitness and wellbeing (IDTechEx Research 2021). This proliferation has extended across disparate domains, ranging from medical applications and fitness and social technologies to military, industrial, and manufacturing applications. As with any emergent technology that promises to leave such an extensive footprint on society, wearables present opportunities for our benefit as well as moral challenges that we should not be negligent in addressing. A crucial facet in this discussion concerns the impact of smart wearables on human decision-making. Through applications such as smart wearables, decisions are increasingly being made with the support of algorithms. This raises the question of the extent to which a person's ability to make autonomous decisions is impaired or promoted by this digital support. For example: If machine learning algorithms and artificial intelligence (AI) relieves the need to make decisions, is the resultantly freed up cognitive resources productive for other cognitive and decision-making processes, or does the relief lead to deskilling or unlearning decision-making? The former possibility is what we should undoubtedly hope for, but if the latter were the case it could result in significant harm for both society and individuals - thus elaborating on the answer to the question is of serious ethical importance. Nowhere is this a more pertinent concern than when dealing with persons who are uniquely vulnerable to infringements on their autonomy. *In this regard, the present contribution investigates, from an explicitly normative and ethical perspective, the potential impact of smart wearables on various dimensions of the autonomy of decision-making of persons uniquely vulnerable to autonomy infringement.* By "person uniquely vulnerable to autonomy infringement" here, we mean a person who, in generality, either displays a lack or deficit of the capacities necessary for, or an oversensitivity to unacceptable external influences on, their decision-making. Such vulnerability is overwhelmingly determined by the internal cognitive processes relevant to decision-making and their often (but certainly not exclusively) age-related changes. We restrict our consideration in this work to formative agents and persons with autonomy disabilities, knowing full well this is far from an exhaustive accounting of all uniquely vulnerable persons. Furthermore, we limit ourselves to a consideration of commercially available smart wearables for private use – particularly within a health or fitness context as users from vulnerable groups are particularly prevalent in this context. This is not to say that state-employed

or industrial smart wearables – such as those used in military contexts, in factory production settings or by the judicial system – are not morally pertinent, quite the contrary, but they introduce confounding complexities into the discussion that we hope to avoid here. They would undoubtedly be excellent targets for future research. Finally, this investigation is explicitly undertaken from a normative perspective, which, though informed by empirical and descriptive work, seeks to preemptively identify ethical considerations that we foresee as relevant to the interaction between smart wearables and vulnerable persons.

We proceed in five steps: we begin in Section 2 by presenting how we take the process of decision-making to be best understood, and then identifying and affirming the central moral importance of autonomy in such a process – particularly what we call here *decisional autonomy* – when addressing the potential impact of commercial smart wearables for commercial or private use.

In Section 3, we examine the constitutive features of smart wearables that structure our definition and thereafter outline the qualities and capacity of smart wearables that give them the potential to be uniquely effective vectors for impacts on the autonomy of user decision-making, looking at three qualities: proximity, convenience, and ubiquity. In addition, the capacity of smart wearables for facilitating cognitive offloading for the user will be considered.

Equipped with the insights from the first two sections, in Section 4 we sketch four opportunities for autonomy-promotion (increased informational input, freeing cognitive resources, extending the range of agency, nudging) and four concerns about autonomy-reduction (privacy, overchoice, dependency and deskilling, sludging and overnudging) raised by smart wearables.

In Section 5, we motivate the need to consider impacts on vulnerable persons. We put forward three reasons for focusing on vulnerable persons: (a) They are especially vulnerable to harms and manipulations and have a reduced ability for recourse in the face of such, (b) they stand to benefit from support provided by wearables that compensate for autonomy impairments or fosters developing capacities necessary for autonomy, and (c) as a society, we often permit violations of the decisional autonomy of these persons in the name of other values where such violations would be intolerable when applied to others. We then briefly touch on some of the unique ways in which these vulnerable groups could be differentially impacted by smart wearables.

2. Decision-making and decisional autonomy

For our purposes we will understand the process of decision-making as follows: the cognitive process of choosing between two or more alternatives, ranging from the relatively clear to the complex. *Decisions* describe the choice between at least two options or alternatives based on personal preferences. Often the relation between an option and its consequences is probabilistic, so that the degree of uncertainty of possible consequences (i.e., the risk) is an important characteristic of a decision (Edwards 1954). However, in numerous decision-making situations, especially when dealing with complex, dynamic technical systems, either the consequences or the probabilities of their occurrence are unknown. Decisions are called “risky” above all when some of the possible but uncertain outcomes are particularly unpleasant or associated with high costs. While, for example, human factors psychology is interested in how people *should* make decisions according to an optimal framework, the decision-making research investigates to what extent errors or biases in the decision process can be attributed to limited human attention, working memory, or selection strategies or familiar decision routines. In order to ensure an empirically informed perspective on the concept of decision making, this will be briefly discussed below. The aim of this description, however, is not to introduce empirical research.

The information processing relevant to a decision begins following a human factors perspective with the user extracting cues of certain modalities from the environment and briefly storing them in the short-term memory (Wickens et al. 2021). Subsequently, the sensory stimuli are filtered. Here a selection process (clue filtering) transmits only those stimuli to conscious processing (perception), which are considered as relevant to a certain situation, based on the experience of the decision-maker. This “selective attention” is centrally controlled and binds attentional resources depending on the complexity of the problem. Selective Attention is considered the first step in the decision-making process. As humans are not passive information processors but actively engage in the process, the filtering can be initiated by the stimuli themselves (bottom-up) or from information from the long-term memory (top-down). Subsequent perception of selectively perceived stimuli serves their identification and interpretation. Based on the selectively perceived and processed information an understanding and assessment of the decision situation in terms of a diagnosis is created. Cognition and working memory are considered as central, supporting the planning and diagnostic process, and organizing a reciprocal exchange of information with the long-term memory. One main

goal of the diagnosis phase is to build hypotheses about the external world and the decision-space, based on which an adequate response selection is made. For the development of a diagnosis, the concept of situation awareness (Endsley 1995; Durso et al. 2007) is of major importance. Building situational awareness includes three stages: In the first stage, all relevant information is perceived from the environment. The perceived information is then integrated top-down or bottom-up to an appropriate understanding of the current situation so that the further dynamic development of the current situation can be correctly predicted, and anticipation of future information can be derived. Across all stages, a general understanding of the system is built, from which hypotheses about system behavior and diagnoses can be derived. Based on the diagnosis, the process of action selection is then initiated, evaluating the expected consequences and the associated values of a decision (cost-benefit consideration), which in turn triggers the execution of the action. A significant factor influencing the choice of action is also the awareness of one's knowledge. Good decision-makers are aware of information lack and therefore search particularly attentively or, if necessary, wait for essential information before deciding. Since situation awareness involves the evolving decision process, it also shows a clear connection to meta-cognition (Edwards 1954; Rousseau et al. 2010).

Though there is certainly a rich and ongoing study of human decision-making,² we take the above description to be sufficient, though far from exhaustive, in order to launch our central aim: the investigation, from a normative and ethical perspective, of the impact of smart wearables on the decisional autonomy of certain vulnerable persons. Viewed through such an ethical lens, we take there to be two ways in which technologies such as smart wearables can impact this process: (1) impacting the quality of the decision made and (2) impacting the autonomy of the decision-making process. The positive face of the former sort of impact is usually the selling point for the technology in question, the promise that using the device will lead to the user making better choices – i.e., choices that promote the user's wellbeing or personal utility. This can be achieved, for example, through the device bringing information to the user's situation awareness that they would otherwise lack, providing more accurate risk assessment than the user could hope to, compensating for limited attention and memory, or through the device counteracting some error or bias in the user's

2 As a small sample of the available work, see Resnik 1987; Ben-Haim 2001; Bermúdez 2009; Martin 2009; Buchak 2010.

decision-making that would have resulted in suboptimal outcomes. Of course, this sort of impact also has a negative face. Technologies can also serve, despite the promises of their purveyors, to result in users making worse choices. The adjudication of which of these impacts dominate for a given technology within a given context is a crucial part of ethical reflection on the justifiability of its use. Important as this consideration is, however, we will largely be setting it to one side for the remainder of this piece to focus on the possible impact of smart wearables on the autonomy of our decision-making.

In everyday life, we value not only making good decisions but also our ability to make our own decisions (Christman 2005). We may struggle in the face of difficult choices and wish for someone else to take them out of our hands, but it is exceedingly rare for us to warm up to the idea that someone or something else will or should make our decisions for us in any general sense. Even if I cede a choice to someone else, I will want to retain a veto. This resistance, and the value given to this self-government or autonomy that it reveals, is at least in part a *moral* value, and so any infringement upon it is open to a demand for ethical justification – and if this is lacking, should be prohibited (though not universally supported, this view is widely endorsed across a wide spectrum of normative theorists – for examples see Dworkin 1988; Korsgaard 1996; Veltman and Piper 2014). But what does it mean to infringe on the autonomy of a person's decision-making? As extreme examples, it seems clear that if I am hypnotised to vote for a certain candidate in an election, for example, then this decision was not autonomous. If I am physically addicted to a narcotic and out of addictive compulsion choose to sell my most prized possession for a fix, this is (at least) not fully autonomous. For our purposes, we will follow Niker et al. (2021) in holding that to decide autonomously is for that decision to be:

1. The result of your own (evidence- and reasons-responsive) decision-making processes.
2. Guided by your authentic aims and values.
3. Without undue external influence.

It is far beyond the scope of this work to provide a defence for the legitimacy of autonomy as a moral value, so from here we will follow the assumption that, barring reasons for exemption in exceptional cases, human beings have morally significant legitimate claim to autonomy over their decision-making. As a stronger claim, we will assume that within the limited space of what we will call *commercial smart wearables*, those designed, developed, purveyed, and supported by commercial entities such

as corporations for use by corporations or private citizens, a policy of *autonomy priority* should be followed. According to this position, autonomy should be the first moral value we worry about, and downstream benefits of potential consequences should be considered only after autonomy impacts have been accounted for.

Autonomy priority is not as controversial an assumption as it may *prima facie* seem, as it falls in line with well-established facts of both everyday and legal practice. Though states and their organs are sometimes empowered (rightly or wrongly) to violate the autonomy of its citizens – usually in the name of well-being or some other moral value (Feinberg 1983) – we do not, nor should we, permit commercial entities the same power over their customers. That their product would improve the quality of decision-making for users is not sufficient justification for a commercial entity to infringe on the autonomy of its customers. This is reflected in the much-discussed General Data Protection Regulation of the EU (GDPR), which is primarily aimed at improving the data sovereignty of EU citizens by regulating various facets of data gathering, processing, and use. The concerns the GDPR seeks to address, from a moral dimension, are the misuse of data and the violation of privacy. As the moral harm of privacy violation is plausibly best understood as the result of a violation of autonomy (Altman 1975; Debatin 2011), this urgent push for its protection is what we would expect if the underlying assumption were that corporate actors should not be permitted to violate autonomy.

This concern can also be seen within the existing draft of the EU's Artificial Intelligence Act (AIA), which adopts a risk-based approach to categorizing AI technologies (CNECT 2021). This takes the form of a proverbial pyramid of risk (Kop 2021), ascending to the top. At the summit of the pyramid – labelled “Unacceptable risk” – we find *prohibited AI practices*. Four examples of such practices are provided:

- AI systems that deploy harmful manipulative ‘subliminal techniques’
- AI systems that exploit specific vulnerable groups
- AI systems used for social scoring purposes
- “Real-time” remote biometric identification systems in publicly accessible spaces for law enforcement purposes, except in a limited number of cases

Of these, the use of harmful manipulation, the exploitation of vulnerable groups, and social scoring are all definite examples of autonomy impairment. Even the prohibition on real-time remote biometrics, though more obviously aimed at preserving privacy (as mentioned already, unquestionably an important moral consideration and one intimately tied to auto-

nomy itself), contains risks of autonomy impairment as it might coerce citizens out of public spaces. This is unsurprising, given that the protection of human rights is the chief moral principle at work in the draft law, and such rights have a long history of association with considerations of autonomy (Pateman 2002; Thrasher 2019; Niker et al. 2021). It should be noted, however, that the EU draft law explicitly takes *harm* as the foremost consideration for evaluating the risk posed by a particular AI technology. Depending on definitions of harm, this may not constitute autonomy priority. That said, the AIA is not aimed exclusively at regulating commercial technologies, but also those in industrial, military, or other domains. As we have already mentioned, these domains add other moral considerations to the table, and so we should not expect complete overlap between its approach and ours. With this in mind, we will take *decisional autonomy* as the moral value of central importance when discussing *impacts on decision-making* and when limited to the case of commercial wearables for private use. By doing this we do not discount or devalue other morally relevant concerns in the vicinity.

3. Qualities and capacity of smart wearables relevant to decisional autonomy

In this section, we unpack the constitutive features of smart wearables that inform our definition, and thereafter explore the qualities and capacity of smart wearables that are most salient for decisional autonomy. To reiterate, we take smart wearables to refer to wearable technologies that have the capacity for the collection and algorithmic processing of data, often including a wireless connection to a user network or the internet, and often in order to produce corrective output via means of actuators integrated into the worn item. Even with our restriction to only commercial wearables for private use, this still casts an intentionally wide net, including as it does activity trackers, augmented reality devices, E-textiles, EEG and ECG belts, smart watches, and an ever-increasing catalogue of new developments. There is a plethora of ways to classify the contents of this net depending on the perspective adopted. Ometov et al. (2021: 6-9) lays out five possibilities for the classification of smart wearables, each organised around a different factor: classification by application/functionality, classification by device type, classification by worn location, classification by energy-consumption, and classification by battery type. Though classification by energy-consumption or battery type might be of great ethical import in light of questions of sustainability, given our interest lies in how these devices impact the decisional autonomy of their users, we take classification around functionality (illustrated in Table 1) to be the most salient for our purposes.

Table 1: Classification based on the wearable application/functionality types (ordered alphabetically)

Type	Brief description
Communication functionality (C)	Provides the potential not to process the data locally but to exchange it with surrounding nodes and/or remote cloud.
Control/input functionality (CI)	A broad area of input devices ranging from smart buttons to sophisticated gesture recognition devices. This group's main task is to extend conventional Human-Computer Interaction (HCI) input focusing on the usability of the devices keeping a small form-factor as a rule.
Education and professional sports (ES)	Aim at improving the education and training by monitoring assistants.
Entertainment, gaming and leisure functionality (E)	The improvement of the perception experience including e.g., audio systems, personal entertainment displays, etc.
Heads-up, Hands-free Information (H)	Extend the conventional ways of the data delivery to the user utilizing personal assistants, AR, XR, Remote Expert Devices, wearable cameras, etc.
Healthcare/medical functionality (HM)	Separated from conventional sensing and monitoring ones due to the need to obtain medical device status that requires significant effort in the device development and testing as well as providing a high level of the obtained data trustability and the need for additional certification, however, covering similar devices, e.g., Electrocardiogram (ECG), Electroencephalogram (EEG) monitors, relaxation devices, neural interfaces, exoskeletons, etc.
Location tracking functionality (LT)	Requires having either some Global Navigation Satellite System (GNSS) on board or, at least, a wireless communication technology. On the one hand, the concept here corresponds to location awareness from the node's perspective and, on the other hand, to remote localization of the device if needed.
Notification functionality (N)	Ranges from simple vibration notification to complex AR extensions. Similarly to sensing functionality, almost any personal device connected to the cloud directly or via the gateway can carry this functionality.
Output functionality (O)	Various visual, audio, or haptic-enabled devices to provide the user and/or people around with prompt information from the personal ecosystem.
Safety and Security functionality (S)	Personal safety devices, emergency assistants, etc.
Monitoring functionality (M)	Extremely straightforward and cheap to implement this functionality. Generally, any device that has an accelerometer on board can already provide some level of sensing. (Fitness and preventive healthcare – Activity Trackers, ECG, EEG monitors, etc.)
Wearable devices for pets and animals (PF)	Mainly covers smart collars, bark collars, smart clothes, etc.

Source: (Ometov et al: 7)

It can be immediately seen that the types in this classification are of two distinct kinds, those that track features possessed by the wearable in question, and those that track the purpose of its application. C, CI, H, LT, N, O, and M are all features that a wearable can possess, whereas ES, E, HM, and PF are domains of application. Differing domains of application will undoubtedly give rise to unique use cases and so in turn unique

moral opportunities and challenges, and a fine-grained examination of these is, and will remain, an important task for as long as we employ these technologies. However, as the moral opportunities and challenges we seek to identify in this work are those that we think of as *applicable to the entire class of smart wearables*, our target here must be those features that are *constitutive* of a smart wearable.

Our definition of smart wearables already presents constitutive features of smart wearables: the wearable must possess the sensors capable of collecting data (LT, M), they are able to algorithmically process this data or share it so that it can be processed elsewhere (C), they can employ this data in order to produce aim-guided output (H, N, O). Using Ometov et al.'s framework, we take LT to be a special example of the more general M and see this as a constitutive feature of smart wearables. C is not strictly constitutive but rather contingent – a wearable need not have a networked uplink to an external processing system to qualify as a smart wearable – however, given the limitations in processing power of the worn devices themselves, the outsourcing of this work is likely to be an all but universal feature (Vijayan et al. 2021). In terms of output, we take O to capture this notion in its widest sense, while H and N are all more specific instantiations thereof. Devices that provide hands-free data delivery to a user is providing a certain sort of output, and a notification is a form of visual, audio, or haptic output providing “prompt” information to a user. CI is a more complicated case. Our contention is that while the sort of usable extension to HCI input that Omertov et al. envisage might increasingly be a normative expectation we have about smart wearables, it is not a constitutive feature itself. A wearable that possesses this functionality, will possess one of the other functionalities as well. Furthermore, there are two differences between the sorts of outputs delivered by smart wearables and those delivered by other (non-smart) wearables: firstly, the smart wearable can be imparted with an aim, and will provide output in an attempt to meet the imparted aims (within some, sometimes severe, limits), which directly leads into the second difference, that the output in some cases can serve as a corrective to the device's previous outputs if these have missed the mark – an ability that is only possible thanks to both the features of data collection and algorithmic processing. A fitness watch, for example, can learn to give a user recommendation based on certain health aims, and then alter the output it delivers to achieve this end in the face of the data it continually collects. There are several other smart technologies that have similar features, but as their name gives away, smart wearables represent the application of these features to *worn items*, items that will be

physically in contact with the user's person. Indeed, any technology with these features that is physically worn by a person is a smart wearable.

With this definition and constitutive features in mind, our observation is that smart wearables possess three qualities that, individually and together, work to amplify their potential impact on decision-making. We contend that it is their presence directly on a user's person – *proximity* – combined with their ease of accessibility and use – *convenience* – and omnipresence – *ubiquity* – that make them prime vectors for interventions that impact a user's decisional autonomy. There are *normative expectations* of users towards wearables illustrated by these qualities. As such, *developers and purveyors aim to deliver these qualities, and users expect to be able to make use of them*. Studies involving smartphones – the technology most similar to smart wearables in terms of the three identified qualities – have shown that if a device is in close physical proximity and readily accessible to a user it can quickly become an almost unquestioned part of a person's day to day activity and decision-making (Hamilton and Yao 2018; Reiner and Nagel 2017; Kutscher 2015). The behavioural changes (and dependencies) that this possible unreflective adoption introduces are not always consciously clear to the person being impacted. And since a key benefit to wearables is precisely their convenience as ready to hand tools about which we don't have to give too much thought in day-to-day use, we do not see these as aspects of wearables that can simply be designed away. Nor – presumably – would this be desirable, as many of the benefits of these technologies are the result of these very qualities. A further symptom of the proximity and ubiquity of wearables is that they will also collect data of a highly personal and intimate nature, including health, movement, and location data. This enables interventions for improving a user's life or even promoting their autonomy but raises concerns about possible infringements on privacy and the prospect of this data being misused to undermine the user's autonomy (Ashworth and Free 2006; Belanger and Crossler 2011; Fuller 2019).

In addition to these qualities is the capacity possessed by wearables to facilitate cognitive offloading by the user. It is this capacity that is perhaps most characteristically associated with smart wearables, and what distinguishes the role that they can serve from that of non-smart wearables. By cognitive offloading is meant the delegation of control over the performance of a cognitive task or over the making of a decision to some device

or system.³ This capacity is enabled through the wearable being able to collect and algorithmically process data to produce aim-guided and (at times) corrective output. This ceding of control can come by degrees (Dunn and Risko 2016; Risko and Gilbert 2016; Heersmink and Carter 2017). It may seem initially counter-intuitive, but this relinquishment of control can be part of a global promotion of autonomy (Kohler et al 2014; Carter 2018). Similarly, we sometimes cede control to others in such a way that they can support our autonomy, and with the result that we are better able to respond to the reasons that enable us to best achieve our aims: think of a fitness instructor who modifies our behaviour toward our desired end frequently by usurping control over aspects of our workout regimen. Where wearables are concerned, this allows us to delegate calculating an optimum workout plan to a fitness wearable or the need to search for the nearest Indian restaurant to a smartwatch. Where such delegation in narrow control results in a greater ability for the user to self-govern toward their more overarching ends, this would then be a case of autonomy promotion. However, as we shall see in the next section, this delegation of aspects or the entirety of our decision-making to these technologies can have undesirable outcomes.

4. The opportunities and perils for decisional autonomy

As is likely already clear, the combination of the three qualities and the capacity for offloading means that wearables are a double-edged sword in terms of its impact on the autonomy of our decision-making. This confluence gives rise to ample opportunities for our benefit, not least of which is the potential to promote and scaffold our autonomy. On the other hand, this selfsame combination allows wearables to act as particularly effective vectors for interventions that can reduce our autonomy. We will expand on each of the proverbial swords' edges in turn.

3 To be clear, we do not endorse the view that these technologies as they presently exist have the capacity to make decisions in the way that we have discussed for humans in Section 1. We assume that smart wearables do not possess the cognitive nor autonomous capacities that are distinctive of human decision-making, and which form the conditions of decisional autonomy. Even if we grant that they may possess the functional equivalents of some or all the cognitive capacities necessary, this is not so – at least not yet – for the capacities necessary for autonomous action. Where we talk of these devices “making decisions” or undertaking delegated “cognitive work”, this is merely a colloquialism to ease the discussion.

Smart wearables as a class can promote decisional autonomy along four general dimensions: (i) the freeing up of cognitive capacity, (ii) the provision of informational input, (iii) extending the range of agency, and (iv) nudging us toward an authentic aim of ours. The freeing up of cognitive capacity is self-explanatory: by facilitating cognitive offloading the wearable allows the user to focus on pursuits that they deem to be more valuable, increasing the possibility that the user is able to recognise reasons that would have otherwise been overlooked. Examples of this are easy to come by: offloading the cognitive work of tracking my fitness schedule to a fitness tracker can free up my resources to rather be spent on my more overarching aims. (ii) has already been mentioned in passing, but the provision of otherwise unavailable information can also better allow a user to recognise salient reasons – consider the benefits of an EEG monitor that provides a user with otherwise difficult, or impossible, to obtain information, which in turn permits them to self-govern more effectively. By “extending the range of agency” what we mean here is that the wearable makes directly possible options that were previously unavailable. Most of the obvious examples of this are where wearables are used to assist those dealing with reduced autonomy. A good example of this is the case of Simon Wheatcroft, a long-distance jogger who happens to be blind. Using a wearable device collecting movement and proximity data guides Simon through haptic cues. In the words of Wheatcroft, “As a blind person, you always strive for independence. But it’s a bit of a contradiction, because oftentimes, you’re using somebody with sight to become independent. What we’re trying to do is use this technology to really achieve true independence” (Sisson 2017).

Whereas (i)-(iii) are all easy enough to grasp colloquially, understanding (iv) requires clarifying some jargon, most pertinently: what is meant by nudging. To nudge an agent X as regard some decision Y, is to make changes to X’s choice architecture relevant to Y such that some preferred choice is promoted without either removing any options from the table or introducing new economic incentives (Thaler and Sunstein 2009; Felsen, Castelo and Reiner 2013; Moles 2015; Levy 2017). The idea behind a nudge is that the agent (for our purposes, the smart wearable user) retains full autonomy in her decision-making, but it increases the likelihood that the agent selects the choice the nudgee wants. Nudges can be, and regularly are, employed to promote welfare or even to support the autonomy of the nudgee. Such nudges can be particularly effective when applied by wearables, thanks to the qualities of proximity and ubiquity. A smartwatch that tracks a user’s fitness data while out on a jog and then uses this data to suggest when the user should take a break is a simple but effective

example of a wearable employing a nudge that can prevent a user from overexerting themselves or aggravating a medical condition. This is clearly a case where the nudge serves to promote welfare, while at the same time not obviously infringing on the autonomy of the user. But nudges can be better than autonomy-neutral, and in some cases can actively strengthen a user's capacity for self-governance (Levy 2017; Niker 2021). Envisage the following scenario: a smoker seeks to break her addiction, and to this end she purchases a health wearable that can remind a user of the dangers of smoking, perhaps accompanying the warning with off-putting images, when it detects the user is smoking. The device is serving to support the user's autonomy by supporting their attempt to quit smoking. These nudges can also be far subtler than direct communication with the user: the layout of a user interface – colouring one option brightly while leaving the other dull, placing some qualities very visibly while placing others behind menus, etc. – can nudge users toward some choices over others. For this reason, the design of a user interface must be carefully considered, both to avoid unintended nudges and where possible to employ nudges that best support the autonomy and welfare of the user.

When we consider the challenges to the decisional autonomy of smart wearable users, there are three general categories these can fall into: (i) the risk of overchoice, (ii) the risk of de-skilling and dependency, and (iii) the possibility for sludging and overnudging. The first of these has been mentioned already and refers to the – now well studied – situation where the provision of increased options serves to reduce the user's ability to choose the option that is in fact the best fit for her authentic aims. It is vital, therefore, that wearables should strive to provide palettes of *relevant* options in a fashion usable to the user, and that user agreements (a common environment for overchoice) should be aimed more at explainability and usability than sheer transparency or providing maximum details. In terms of the reducing overchoice in the application of wearables, agentially useful epistemic accessibility will often only be accomplished through active user engagement and feedback – the best way to know what option or information will be relevant and useful to a user is to facilitate increased responsiveness to user needs. But it should also not be expected that users will have uniform needs. Different options and different information will be variably useful to different users, a common-sense fact but no less important for being so. Fortunately, where the provision of options is concerned, smart technology is well-positioned to tailor options and informational input to the needs of individual users in a dynamic fashion.

When regarding user agreements, which must be hurdled before such tailoring can be undertaken, these can often be difficult to penetrate. Some

of which employ overly jargon-riddled technical or legal language that form a barrier to comprehension. Though not strictly the result of the features of smart wearables themselves the very nature of the algorithmic processing that is essential to the effective operation of smart wearables offers a barrier to providing easy epistemic access, as though such systems may not be black boxes, they can still be very dark shades of grey (Jovanović and Schmitz 2022). Thus, though there is undoubtedly a moral reason for the drafters of these agreements to strive for explainability, this will understandably not be their only concern, and there is also a moral onus on the user to take reasonable steps to educate themselves on the legal details of the agreement they enter. How all this is to be achieved is an important discussion, but outside the focus of this piece and so we only take this opportunity to gesture towards its significance.

Turning to (ii), although it is true that cognitive offloading can result in a promotion of decisional autonomy, the opposite outcome is also possible. One way in which this can occur is if a user becomes too dependent on a device, such that their own skills and decision-making ability atrophy to the point where autonomy is threatened (this is often referred to as “de-skilling” (Vallor 2015)). This is most likely to occur in situations where the use of the technology becomes unreflective or habitual, precisely the danger raised by the proximate, ubiquitous, and convenient nature of wearable technologies. There are two ways in which this sort of atrophication can prove dangerous to autonomy: a) where a dependency forms on an unreliable technology and b) where the dependency stunts the development of capacities necessary for decisional autonomy. If the technology is unreliable, then the delegation of control from the user to the technology in order to grant them greater overall control backfires. The user, if they are dependent – that is, the skill necessary to fulfil the task the technology now fulfils has atrophied away – will be left with reduced overall autonomy if the technology fails. An illustrative example involving smart wearables would be the use of an augmented reality headset for in-store product comparisons during shopping. Assume a user who has become dependent on this functionality in order to make purchasing decisions, but then experiences a failure of their device. Bereft of the guidance from the wearable, this user is now incapable of making effective (that is, in line with their authentic ends) purchasing choices – whereas this would not have been the case had they never formed the dependency. To be very clear, this is not a polemic against *any* dependency resulting from cognitive offloading: dependency on navigation technology is a boon to the autonomy of many, and since these systems are sufficiently reliable (most of the time!) we can judge that they are autonomy-promoting. The

vital takeaway from (a) is that the developers of technologies that can facilitate cognitive offloading must be thorough in assessing whether or not they are likely to result in dependencies, and if they are it is vital that the reliability of the technology be of the highest order. However, even if reliability is not a concern, there are still some dependencies that may be pernicious to decisional autonomy. If the dependency results, directly or indirectly, in the stunting or loss of a capacity necessary for decision-making, then we have *pro tanto* moral reason to oppose the dependency-inducing technology, one that will rarely if ever be overridden in the commercial realm. Though there is not yet robust evidence of this occurring with smart wearables, the first longitudinal studies on the topic have found that the growing use of, and dependency on, various digital offloading technologies correlates with a deterioration in attentional capacities among adolescents (Baumgartner et al, 2018), capacities which we identified in Section 1 as necessary for human decision-making and by extension for decisional autonomy.

Lastly, we have (iii). Following current convention, we call nudges that nudge a user against their best (or better) interests, *sludges* (Thaler 2018). Given our account of autonomy, such sludges can be autonomy reducing if they work against a user achieving their authentic aims. These sorts of interventions can take many shapes but are usually employed with the interest of increasing profits at the user's expense. Having a pair of smart glasses that consistently gives listing priority to products manufactured by the purveyor of the device even when these are not the best value is an example of a measure that can function as a sludge, inducing customers to purchase these products even when it works against their authentic aims – assuming they aim to purchase the best value product in this example. Combating sludges is often best achieved by informing users about their presence and the danger they pose. Awareness of a nudge or sludge, though not foolproof, can go a long way to helping people resist its possible effects on their decision-making.

Apart from sludges, there are two other ways in which nudges can undermine autonomy. Firstly, our aims and values can often prove very endogenous, leaving us vulnerable to being nudged away from our own authentic self-government. This is particularly true if nudges operate by bypassing our deliberative capacities (Grüne-Yanoff 2012). Secondly, nudging can serve to prevent or impair the development of capacities necessary for autonomy by cutting a user off from irreplaceable learning experiences (Blöser et al 2010; Niker et al. 2021). This is exacerbated when the nudgee is the target of many concerted nudges or the source of the nudging is unreflectively integrated into the nudgee's decision-making. One of the

best and simplest ways to combat this risk is to inform users about how they are being nudged – or will be nudged. This will likely reduce the efficacy of at least some nudges, which often work best when undetected, but this is a price that should be paid in seeking out the appropriate balance, especially in a commercial context (Sunstein 2014).

5. Vulnerable persons and the differentiated impact of smart wearables

Although the above-mentioned impacts on decisional autonomy are relevant for all users, this is nowhere truer than in the case of what we will here refer to as *persons uniquely vulnerable to autonomy infringement*. In this paper we use this term to encompass those *individuals* whose capacities for autonomous decision-making are, *in generality*, more sensitive to negative impacts. This sensitivity is the result of one or more elements of the decision-making process – as discussed in Section 1 – operating at a level sufficiently less than that normatively expected of an autonomous decision-maker. This state is multiply realisable as it can take many forms, some examples might include: a shortfall in working memory, a limited attentional capacity, a shortcoming in integrating information stimuli into situational awareness, or a limited capacity for meta-cognition. This is also to be understood multidimensionally, where shortfalls in some elements and gains in others can co-exist. These varied impacts may indeed be incommensurable on final appraisal, thus preventing the formulation of a straightforward final verdict on whether decisional autonomy has been positively or negatively impacted. Viewed ethically, the consequence of a diminishment along a dimension of a user's decisional autonomy is that they are less able to pursue and achieve their authentic ends through their own (evidence- and reasons-responsive) decision-making processes, and/or there is a higher risk of external influences having an overriding impact on these processes. That is to say, they are more *likely* to have their self-governance infringed, though of course this possibility need not transpire for them to be vulnerable. As it is far too wide a task for this piece to address every possible variation of such uniquely vulnerable persons, we choose to focus our attention on certain illustrative examples in order to better elucidate our claims. This is not to claim that these examples represent the only types or groups of persons who are uniquely vulnerable to autonomy infringement, nor that they necessarily represent the examples most worthy of consideration. Our choices here are motivated solely by the practical aim of illustrating the potential impacts of smart wearables on vulnerable persons as digestibly as possible.

On our assessment there are two primary categories of vulnerability we should pay special attention to as regards the impact on the decisional autonomy of vulnerable persons:

1. Harm vulnerability: these persons are especially vulnerable to harms and manipulations resulting from failures of self-governance and may have a reduced ability for recourse in the face of such
2. Paternalism vulnerability: as a society, we often permit violations of the decisional autonomy of vulnerable persons in the name of utility or other moral values where such violations would be intolerable for others

The first of these is self-apparent, but note that it makes sludging, dependency formation, and privacy violations of vulnerable persons *uniquely concerning*. To illustrate the second, consider that vulnerable persons are frequently "protected" by excluding them from access to certain activities or technologies. This is not to imply that all vulnerable persons are treated equally in this regard but being overprotective to the degree of paternalism is a commonality in their treatment. Given our commitment to autonomy priority, at least in the case of commercial smart wearables, we adopt a skeptical stance toward the justification of any paternalism that violates decisional autonomy.

The first example we consider is of a user with an age-related diminishment in one or more of their decision-making capacities. They remain full agents, but as a result of the diminishment in their capacities they have some limitations on their ability for effective self-governance – though this is not tantamount to concluding that they, all-things-considered, lack decisional autonomy. One possible version of this example would be of an older adult with diminishments in their situational awareness, memory, attention, and overall cognitive abilities as a result of natural, biological attrition (Wilson et al. 2002; Salthouse et al. 2003; Deary et al. 2009). Deficits in situational awareness amplifies the opportunity and the risk for decisional autonomy of those impacts that are most effective when unreflected upon – precisely what we take to be the result of the morally-relevant qualities of the class of smart wearables. Positively, it is precisely here where smart wearables can best promote the autonomy of such persons by allowing them to offload tasks, thus freeing up their comparatively more valuable cognitive resources (Lewis and Neider 2017). There is also increased scope for effective nudging, however the inclination toward paternalism – even of the libertarian variety – must be carefully balanced lest, as Schachar and Greenbaum (2019) fears may happen all too easily, the nudge becomes a shove. More unambiguously negative,

persons with age-related diminishments remain particularly vulnerable to sludging, as such influences thrive in the absence of reflection, memory, and attentional capacity.

Our second example case is that of a person with an autonomy-impairing disability. To be clear, not all persons with disabilities will necessarily have reduced autonomy, either in a particular dimension of autonomy or all-things-considered. But certain persons with disabilities will have this experience, where they have a limitation in one or more dimensions of autonomy. As this remains too wide a grouping to be illustrative, let us use the particular case of Simon Wheatcroft, which we have already touched on. Simon lost his sight to a degenerative eye disease while a teenager but remained an avid long-distance jogger. Due to his condition, he found himself severely limited pursuing this dearly held interest of his, having to stick to well identified running paths and requiring a sighted guide runner when participating in large city marathons. As we described, using a wearable device capable of providing guidance through haptic cues, Simon was able to extend the range of his agency, and thereby pursue his authentic ends that had previously seemed unattainable. Though his story so far has undoubtedly been one to celebrate, individuals in Simon's situation remain uniquely vulnerable along two dimensions of decisional autonomy. Firstly, there is the risk of forming a dependency on these technologies supports, and if this couples with an atrophication of the ability to perform long-distance jogs unaided, it could leave Simon in a position of reduced autonomy if the technology fails. And secondly, those in Simon's position are vulnerable to overnudging and sludging by those who design and supply them with the technologies on which their extended agency depends.

The final example, which we will consider in more depth than the preceding two, is that of *formative agents*, whose capacities necessary for decisional autonomy are still nascent and developing. Though not only applicable to them, children and adolescents are usually considered the paradigmatic examples of such near-autonomous agents (Graf et al. 2013). This is also recognized and enshrined in the UN Convention on the Rights of the Child, which holds the *evolving capacities* of the child to be a core notion (UN-CRC, Art. 5). There is undoubtedly irreducible vagueness as to when precisely a formative agent comes into their own as fully autonomous, a dynamic that we can clearly see with the border between childhood, adolescence, and adulthood. Given this, we do not aim here to specify precise points of transition but take talk of children and adolescents to be sufficiently intuitive to grasp for our illustrative (and not intended to be exhaustive) purposes.

As children develop throughout the stages of childhood and adolescence and their capacities evolve over time, so their resilience towards potentially negative influences grows. This marks a fundamental difference to our two other examples. The younger the child, the more they need to be protected from such risks. The inchoate state of their development results in the preferences and ends of younger children being far more endogenous – and so vulnerable to nudges and sludges – than adults. Also, as children's capacities for agency are evolving with their age, the impact of interventions that promote, reinforce, stunt, or deform these capacities is uniquely amplified. As children grow older, protection can gradually be reduced to attempting to avoid serious risks and focus on the child's growing autonomy and ability to cope with nudges and sludges. This is reflected in detail in *The Intelligent Risk Management Model* developed by the German Centre for Child Protection on the Internet (2015). The model is based on an age-related concept designed both to protect children and adolescents but also to support them in developing coping strategies and skills. Parents and other guardians play a crucial role in this process, facing two general duties that can at times conflict: The duty to ensure the well-being of the child or adolescent and the duty to promote and respect the autonomy of the child or adolescent, so that they can learn and practice how to use their autonomy-enabling capacities, which will often involve allowing them to "make their own mistakes". A guiding principle in this conflict should be the "best interest of the child as a primary consideration" (UN-CRC, Art. 3). Accordingly, a violation of the child or adolescent's autonomy should only be justified where it is in their best interest, e.g., to protect them from *severe* or *unforeseeable* harm, while still allowing them the space to practice and fail. Based on the assumption that younger children are not able to oversee the consequences of disclosing private data, stronger restrictions to the use of such children's data could either be exercised by their parents or placed within the device. While parents might be inclined to infringe their children's privacy by being overprotective, safety-by-design built into the device or service could even support the child or adolescent's acquisition of data literacy and free their cognitive capacity. Smart wearables for children and/or adolescents have a high potential to extend their agency, nonetheless, designing wearables for children and/or adolescents needs to take into account when the informational input oversteps the balance of freeing versus locking cognitive capacities. In order to promote children and adolescents in their process of learning and development, certain tasks should not be taken out their hands by a wearable: for example, a smartwatch providing continual and immediate informational input could inhibit the development of attentional capacities. Additional-

ly, for adolescents, tracking weight and other body data with smart wearables may help them gain autonomy over their own decisions regards sports activities and nutrition through informational input, freeing cognitive capacity, and nudging. But at the same time, it might make them more dependent on statistical norms and puts them at risk of socially-mediated sludges, such as pressure from their peer group toward unhealthy behaviours.

6. Concluding Remarks

With the likely tremendous growth in smart wearable use over the coming years, it behoves us to take the opportunity to access the moral impacts that may accompany it. Here we have sought to unpack how commercial wearables will influence the decisional autonomy of users, with a special focus on persons uniquely vulnerable to autonomy infringement. We argue that there are several unique perils and opportunities for decisional autonomy that arise from the unique qualities of smart wearables and their capacity to facilitate cognitive offloading. What is more, these perils and opportunities, insofar as they originate from the same qualities and capacity, cannot be “designed away” – they will always demand ethical engagement and reflection in order to produce the most favourable balance between the morally desirable benefits on offer and the morally worrisome outcomes. Those examples of vulnerable persons we discuss are all particularly vulnerable to the possible infringements on decisional autonomy, but also stand to uniquely benefit from some of the opportunities. In light of this, the developers and purveyors of these technologies are under moral obligation to weigh these considerations in the design, proliferation, and support of smart wearables, and should pay special attention to the cases of children, seniors, and persons with non-age-related autonomy impairments.

References

Altman, I. (1975): *The environment and social behavior: privacy, personal space, territory, crowding*. Monterey, CA: Cole Publishing Company.

Ashworth, L. and Free, C. (2006): Marketing Dataveillance and Digital Privacy: Using Theories of Justice to Understand Consumers' Online Privacy Concerns. *Journal of Business Ethics*, 67, 107–123.

Belanger, F. and Crossler, R. E. (2011): Privacy in the Digital Age: A Review of Information Privacy Research in Information Systems. *MIS Quarterly*, 35(4), 1017-1041.

Blöser, C., Schöpf, A., and Willaschek, M. (2010): Autonomy, experience, and reflection. On a neglected aspect of personal autonomy. *Ethical Theory and Moral Practice*, 13(3), 239–253.

Bower, M. and Sturman, D. (2015): What are the educational affordances of wearable technologies? *Computers & Education*, 88, 343-353.

Carter, J. A. (2018): Virtue Epistemology, Enhancement, and Control. *Metaphilosophy*, 49(3), 283-304.

Centre for Child Protection on the Internet (2015): *The Intelligent Risk Management Model*. URL: <https://childrens-rights.digital/hintergrund/index.cfm/topic.279/ky.1497>

Deary, I. J., Corley, J., Gow, A. J., Harris, S. E., Houlihan, L. M., Marioni, R. E., Penke, L., Rafnsson, S. B., and Starr, J. M. (2009): Age-associated cognitive decline. *British Medical Bulletin*, 92(1), 135–152.

Debatin, B. (2011): Ethics, Privacy, and Self-Restraint in Social Networking. In Trepte S and Reinecke L (Eds.), *Privacy Online* (pp. 47-61), Berlin: Springer-Verlag.

Dian, J. F., Vahidnia, R., and Rahmati, A. (2020): Wearables and the Internet of Things (IoT), Applications, Opportunities, and Challenges: A Survey. In *IEEE Access*, 8, 69200-69211. DOI: 10.1109/ACCESS.2020.2986329.

Dunn, T. and Risko, E. F. (2016): Toward a Metacognitive Account of Cognitive Offloading. *Psychology, Medicine, Computer Science*, 40(5), 1080-1127.

Durso, F. T., Rawson, K. A., and Girotto, S. (2007): Comprehension and Situation Awareness. In Durso F, Nickerson R, Dumais S, Lewandowsky S, and Perfect T (Eds.), *Handbook of Applied Cognition: 2nd Edition* (pp. 163-193), Hoboken, NJ: Wiley.

Edwards, W. (1954): The reliability of probability-preferences. *The American Journal of Psychology*, 67(1), 68-95.

Endsley, M. R. (1995): Measurement of situation awareness in dynamic systems. *Human factors*, 37(1), 65-84.

European Union. (2016). *Regulation (EU) 2016/679 of the European Parliament and of the Council 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)*. In: Official Journal of the European Union, L 119, 04 May 2016, S. 1-88.

European Commission (2021). *Proposal for a Regulation of the European Parliament and of the Council laying down harmonised rules on Artificial Intelligence (Artificial Intelligence Act) and amending certain Union Legislative Acts*. COM(2021) 206 final. Brussels. Available at: <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52021PC0206> [accessed 12 May 2022]

Feinberg, J. (1983): Autonomy, Sovereignty, and Privacy: Moral Ideals in the Constitution. *Notre Dame Law Review*, 58(3), 445-492.

Felsen, G., Castelo, N., and Reiner, P. B. (2013): Decisional enhancement and autonomy: Public attitudes towards overt and covert nudges. *Judgment and Decision Making*, 8(3), 202–213.

Fuller, C. S. (2019): Is the market for digital privacy a failure? *Public Choice*, 180, 353-381.

Graf, W. D., Nagel, S. K., Epstein, L. G., Miller, G., Nass, R., and Larriviere, D. (2013): Pediatric neuroenhancement: Ethical, legal, social, and neurodevelopmental implications. *Neurology*, 80(13), 1251-1260.

Grune-Yanoff, T. (2012): Old wine in new casks: libertarian paternalism still violates liberal principles. *Social Choice and Welfare*, 38, 635–645.

Hamilton, K. A. and Yao, M. Z. (2018): Cognitive Offloading and the Extended Digital Self. In M. Kurosu (Ed.), *Human-Computer Interaction. Theories, Methods, and Human Issues*. HCI 2018. Lecture Notes in Computer Science, vol 10901. Springer, Cham.

Heersmink, R. and Carter, J. A. (2020): The philosophy of memory technologies: Metaphysics, knowledge, and values. *Memory Studies*, 13(4), 416-433.

IDTechEx Research (2021): *Wearable Technology Forecasts: A comprehensive review of market opportunities across all wearable electronic devices, from smartwatches to skin patches, AR, VR & MR to hearables, smart clothing to smart eyewear, and more*. URL: <https://www.idtechex.com/en/research-report/wearable-technology-forecasts-2021-2031/839>

Jovanović, M. and Schmitz, M. (2022): Explainability as a User Requirement for Artificial Intelligence Systems. *Computer*, 55(2), 90-94.

Köhler, S., Roughley, N., and Sauer, H. (2017): Technologically blurred accountability? Technology, responsibility gaps and the robustness of our everyday conceptual scheme. In Ulbert C, Finkenbusch P, Sondermann E, and Debiel T (Eds.), *Moral Agency and the Politics of Responsibility* (pp. 51-68). London: Routledge.

Kop, M. (2021): EU Artificial Intelligence Act: The European Approach to AI. *Transatlantic Antitrust and IPR Developments*. Stanford Law School.

Kutscher, N. (2015): *Internet ist gleich mit Essen - Empirische Studie zur Nutzung digitaler Medien durch unbegleitete minderjährige Flüchtlinge*. URL: https://www.researchgate.net/publication/287209029_Internet_ist_gleich_mit_Essen_-_Empirische_Studie_zur_Nutzung_digitaler_Medien_durch_unbegleitete_minderjährige_Flüchtlinge

Levy, N. (2017): Nudges in a post-truth world. *Journal of Medical Ethics*, 43, 495–500.

Lewis, J. E. and Neider, M. B. (2017): Designing Wearable Technology for an Aging Population. *Ergonomics in Design*, 25(3), 4-10.

Moles, A. (2015): Nudging for Liberals. *Social Theory and Practice*, 41(4), 644-667.

Nagel, S. K. and Reiner, P. B. (2018): Skillful Use of Technologies of the Extended Mind Illuminate Practical Paths Toward an Ethics of Consciousness. *Frontiers in Psychology*, 9, 1251.

Niknejad, N., Ismail, W. B., Mardani, A., Liao, H., and Ghani, I. (2020): A comprehensive overview of smart wearables: The state of the art literature, recent advances, and future challenges. *Engineering Applications of Artificial Intelligence*, 90, 103529.

Niker, F., Felsen, G., Nagel, S. K., and Reiner, P. B. (2021): Autonomy, Evidence-Responsiveness, and the *Ethics of Influence*. In M. Blitz and J.C. Bublitz (Eds.), *Neuroscience and the Future of Freedom of Thought*. Hampshire: Palgrave-Macmillan.

Ometov, A., Shubina, V., Klus, L., Skibińska, J., Saafi, S., Pascacio, P., Flueratoru, L., Gaibor, D. Q., Chukhno, N., Chukhno, O., Ali, A., Channa, A., Svertoka, E., Qaim, W. B., Casanova-Marqués, R., Holcer, S., Torres-Sospedra, J., Casteleyn, S., Ruggeri, G., Araniti, G., Burget, R., Hosek, J., and Lohan, E. S. (2021): A Survey on Wearable Technology: History, State-of-the-Art and Current Challenges. *Computer Networks*, 193, 1-37.

Pateman, C. (2002). Self-Ownership and Property in the Person: Democratization and a Tale of Two Concepts. *The Journal of Political Philosophy*, 10(1), 20-53.

Reiner, P. B. and Nagel, S. K. (2017): Technologies of the Extended Mind: Defining the Issues. In Illes J & Hossain S (Eds.), *Neuroethics: Anticipating the Future* (pp. 108-122). Oxford Scholarship Online.

Risko, E. F. and Gilbert, S. J. (2016): Cognitive Offloading. *Trends in Cognitive Sciences*, 20(9), 676-688.

Rousseau, R., Tremblay, S., Banbury, S., Breton, R., and Guitouni, A. (2009): The role of metacognition in the relationship between objective and subjective measures of situation awareness. *Theoretical Issues in Ergonomics Science*, 11(1-2), 119-130.

Salthouse, T. A., Atkinson, T. M., and Berish, D. E. (2003): Executive functioning as a potential mediator of age-related cognitive decline in normal adults. *Journal of Experimental Psychology: General*, 132, 566-594.

Sunstein, C. R. (2014): Nudging: A Very Short Guide. *Journal of Consumer Policy*, 37, 583-588.

Seneviratne, S., Hu, Y., Nguyen, T., Lan, G., Khalifa, S., Thilakarathna, K., Hassan, M., and Seneviratne, A. (2017): A Survey of Wearable Devices and Challenges. In *IEEE Communications Surveys & Tutorials*, 19, 4, 2573-2620. DOI: 10.1109/COMST.2017.2731979.

Sisson, P. (2017): Beyond the finish line: how technology helped a blind athlete run free at the New York Marathon. *The Verge*, Nov 6. URL: <https://www.theverge.com/2017/11/6/16610728/2017-new-york-marathon-blind-runner-wearworks-wayband-simon-wheatcroft>

Thaler, R. H. (2018): Nudge, not sludge. *Science*, 361(6401), 431.

Thaler, R. H. and Sunstein, C. R. (2009): *Nudge: Improving decisions about health, wealth, and happiness*. New Haven: Yale University Press.

Thrasher, J. (2019): Self-ownership as personal sovereignty. *Social Philosophy and Policy*, 36(2), 116-133.

Vallor, S. (2015): Moral Deskilling and Upskilling in a New Machine Age: Reflections on the Ambiguous Future of Character. *Philosophy & Technology*, 28, 107-124.

Vijayan, V., Connolly, J. P., Condell, J., McKelvey, N., and Gardiner, P. (2021): Review of Wearable Devices and Data Collection Considerations for Connected Health. *Sensors*, 21(16).

Viseu, A. (2003): Simulation and augmentation: issues of wearable computers. *Ethics and Information Technology*, 5(1), 17-26.

Wickens, C. D., Helton, W. S., Hollands, J. G., and Banbury, S. (2021): *Engineering psychology and human performance*. Routledge.

UN General Assembly (20 November 1989): Convention on the Rights of the Child. *Treaty Series*, vol. 1577, URL: <https://www.ohchr.org/en/professionalinterest/pages/crc.aspx>

Wilson, R. S., Beckett, L. A., Barnes, L. L., Schneider, J. A., Bach, J., Evans, D. A., and Bennett, D. A. (2002): Individual difference in rates of change in cognitive abilities in older persons. *Psychology and Aging*, 17(2), 179-193.

Xue, Y. (2019): A review on intelligent wearables: Uses and risks. *Human Behaviour & Emerging Technology*, 1, 287-294.

