9 Digital Behavioral Technology, Deep Learning, and Self-Optimization

Karola Kreitmair

9.1 Introduction

Digital behavioral technology (DBT), which includes wearables, mobile health technology, certain smartphone apps, and various neurodevices, is a rapidly expanding class of technology that increasingly permeates all areas of human life. Individuals use such technology to shape themselves physiologically, psychologically, behaviorally, and socially, in order to become healthier, more mindful, better rested, more creative, and more intelligent versions of themselves.

While previously DBT may have been used in an offline manner, allowing an individual to measure her own performance, e.g., a heart rate belt to be used during running, DBT now consists of massively interconnected sensor and logging technology that yields a comprehensive picture of the physiological, environmental, behavioral, neurological, genomic, and social dimensions of a given individual. Moreover, the quantity of data this technology produces is so enormous, that the only viable means of gleaning robust insights from these data is through deep learning and AI.

While AI can gain valuable inferences from data, it is also prone to some serious flaws that are ethically problematic. Rather than producing objective results, as it is largely perceived to do by the public [1], AI propagates biases that are inherent in the data. Moreover, deep learning, the learning architecture that many AIs used in DBT are built on, operates in a way that may be opaque to human observers. This means that humans cannot explain or control how an AI comes up with a particular solution or organization. Furthermore, algorithms can game the reward functions designed to force them to learn, and deliver useless or harmful results.

K. Kreitmair (✉)
Medical History and Bioethics, University of Wisconsin—Madison, Madison, WI, USA
e-mail: kreitmair@wisc.edu

F. Jotterand, M. Ienca (eds.), *Artificial Intelligence in Brain and Mental Health: Philosophical, Ethical & Policy Issues*, Advances in Neuroethics,
https://doi.org/10.1007/978-3-030-74188-4_9

Within a broader societal context of the reponsibilization of the individual with respect to her health and well-being, DBT is increasingly used for all manner of self-optimization. AI factors into this self-optimization work, and along with it so do such issues as algorithmic bias, opacity, and reward hacking. Unbeknownst to the DBT user, these flaws in AI may affect how she is pursuing self-optimization.

In what follows I will provide an overview of DBT and explain how it is involved in self-optimization. I will then address how AI is incorporated into DBT. Next I will look at some issues with AI that have ethical ramifications. Finally, I will consider how these issues impact the self-optimization work that the user is doing with the help of her AI-enabled DBT.

9.2 Digital Behavioral Technology

Digital behavioral technology (DBT) is a class of technologies that is used by individuals to attempt to alter some aspect of their physiological, neurological, psychological, or behavioral selves. It comprises wearables, mobile health technologies, smartphone apps, and a variety of neurotechnologies [2]. Individuals want to shape their bodies, minds, and lives in ways they (or in some cases others) deem desirable, such as, for example, being more productive at work, improving cognition, meditating better, becoming fitter and healthier, sleeping better, being more creative, being a better lover, kicking undesirable habits, and losing weight. They believe that by using DBT such goals become more achievable.

While there is variety amongst these technologies, most DBT tends to be *personal*, *digital*, and *mobile*. It is *personal*, because it tends to be used by only one individual at a time and monitors or modulates only that individual's functioning. It is *digital*, because it utilizes binary computing systems that enable powerful processing and online connectivity. Finally, it is often *mobile,* because it is small and lightweight enough that individuals can carry it around on their persons [2]. For example, DBT includes sensor-driven technologies such as the Fitbit, that keep track of a user's location and activity levels [3], as well as stimulation technologies such as the transcranial direct current stimulation (tDCS) device BrainStimulator [4] that promise to improve cognition and help with depression, chronic pain, and anxiety.

What makes all these technologies instances of digital *behavioral* technology, is that they are all technologies intended to alter an individual's behavior. Often such behavior is sought to be improved for instrumental reasons, e.g., for a particular person, eating less may lead to weight loss which may lead to better health and well-being. Sometimes changing the behavior is an end in itself, as for instance with technologies that seek to limit symptoms of obsessive compulsive disorder [5]. Note that this does not mean that what is being measured or directly affected is always behavior itself. With much of this technology, what is being measured or affected are physiological or neurological properties of the user. However, it is behavior that is sought to be affected.

9.2.1 Functionalities

DBT operates through a vast array of *functionalities* (see Table 9.1).

One of the most common functionalities is *tracking*. Tracking is the passive registering and recording of a particular dynamic feature of an individual. It is passive, because the user need not actively input data into a device. Tracked dimensions include a wide array of features, such as, for example, an individual's location, her heart rate, her ECG, her breathing volume, the composition of her sweat, her EEG, and her blood alcohol level (see Table 9.2 for a list of the dimensions that can be tracked). In general, tracked dimensions lend themselves to quantification, which makes them amenable to analysis.

Tracking is distinct from *logging*, in which an individual uses the technology to record various features. Logging is commonly used for features that are not readily amenable to quantification, such as keeping track of qualitative dimensions such as mood or satisfaction. For example, the mobile app eMoods allows the user to note daily mood, irritability, and anxiety levels [14], while the mobile app Semistry allows the user to log, classify, and rate sexual encounters and activities [15].

Vocal analysis is a further functionality of DBT. Smartphone apps such as Cogito's Companion [16] detect speech patterns that are associated with mental health conditions like bipolar disorder and depression. Vocal analysis, such as provided by Voicesense, can also be used as a predictor of individual behavior, from the likelihood of someone defaulting on a loan to whether an employee is suffering early signs of burnout [17].

Visual analysis is also a functionality of DBT. Programs like Google Lens permits users to gain information about objects they capture with their smartphone's camera [18]. Users can scan a flower, find out what kind it is, and where the nearest florist is located. Or they can scan someone's outfit and receive information on the brand of the item, as well as similar items available for purchase.

Table 9.1 Functionalities

Functionalities	
Type	Example
Tracking	*See* Table 9.2 *for categories*
Logging	Sex logging (e.g., *Semistry*)
Vocal analysis	Speech pattern recognition (e.g., *Companion*)
Visual analysis	Visual object recognition (e.g., *Google Lens*)
Gamification	Mental health app (e.g., *SuperBetter*)
Stimulation	TDCS (e.g., *Foc.Us Go Flow*)
Drug delivery	Nicotine delivery (e.g., *Chrono Therapeutics*)
Virtual reality	Haptic VR suit (e.g., *Teslasuit*)
Assistant	Sleep assistant (e.g., *Neurogixs Alpha AI*)

Table 9.2 Tracked dimensions

Tracked Dimensions	
Location	E.g., Strava [6]
Activity	E.g., Fitbit Versa 2 [3]
Sleep	E.g., Fitbit Versa 2 [3]
ECG	E.g., Qardio [7]
EEG	E.g., Muse [8]
HR	E.g., Fitbit PurePulse [9]
Respiration rate and volume	[10]
BAC	E.g., Bactrack Skyn [11]
Ingestion events	E.g., Proteus Ingestible Sensor [12]
Sweat composition	[13]

A further functionality of DBT is *gamification.* By introducing game-like elements such as competition (against others or oneself), badges, points, and levels, DBT seeks to capitalize on the appeal of games to compel users to adhere to their use of the technology. Apps like SuperBetter provide users with games that involve completing quests and defeating "bad guys" in an effort to improve mental health, including tackling anxiety, depression, chronic pain, and recovering from concussions [19].

DBT may also function by directly *stimulating* the body or the brain. For example, transcutaneous electrical nerve stimulation devices such as the Thync Relax Pro [20] and transcranial direct current stimulation devices such as the foc.us Go Flow [21] deliver low-intensity electric currents to particular areas of the brain in an attempt to facilitate or inhibit neuronal activity in that area [22]. This is done to modulate brain functioning and improve cognition, relieve symptoms of anxiety and depression, combat cravings, and enhance meditation [23].

Another functionality of DBT is to directly *deliver drugs* that impact the body. Chrono Therapeutics for example, delivers medication such as nicotine transdermally in specific dosages timed to coincide with detected symptoms [24].

DBT also includes *virtual reality* (VR) systems. VR systems create a computer-generated environment into which the user can immerse herself. Such systems were once clunky, expensive, and required precise positioning of computers and sensors throughout a room. Today, they have shrunken to affordable untethered headsets, like Oculus Go [25]. VR can also be extended to include "tactile" or "haptic responsiveness." For example, the Teslasuit consists of a full body suit that uses nerve and muscle stimulation to generate haptic sensations for a fully bodily immersive VR experience. This allows users to feel a virtual breeze, the warmth of a virtual sun, or the touch of an avatar [26].

Finally, a new functionality of DBT is that of AI-enabled *assistant*. For example, Google Assistant, which is integrated into Google's Smartwatch, can listen and talk to the user, integrate questions with data analyzed from various other DBT, and provide the user with information and recommendations regarding a vast array of individual-specific information in a multitude of domains [27]. Examples of domain-specific AI assistants are: (1) Baby Connect [28] [29], which helps parents keep track of the quantity and quality of soiled diapers and expressed breast milk, computes the average duration of breastfeeding sessions, recommends when to nurse and when to switch breasts, and even sends information directly to Twitter; (2) Symptomate [30] [31], which analyzes the medical symptoms a user reports and uses AI to generate a differential diagnosis; (3) Neurogixs Alpha AI [32], which analyzes sleep conditions including EEG measurements and provides customized curation and recommendations; (4) CarePod from Sensory Health Systems [33], which provides lifestyle solutions for the elderly and those with limited mobility; (5) Good Morning Routine [34] and Bedtime Routine [35], which turns on/off lights, sets alarms, opens blinds, and briefs/debriefs you on the day; (6) Controlicz [36], which allows the user to speak to and give commands to smart objects and appliances around the home; (7) WorkAssist AI [37], which analyzes information collected from health and fitness wearables, mobile apps, and online behavior, to determine if the user's behavior is "beneficial or detrimental to productivity," and accordingly generate "clear instructions/recommendations to increase productivity" [38]; and (8) Girlfriend Maya [39], a chatbot who replies to a user's utterances, like "Good night, darling," with appropriately "girlfriend"-like responses.

Having identified the functionalities of DBT, let us turn to the types of users of this technology.

9.2.2 Usage Profiles

DBT is used in different circumstances and by different actors. There are thus different *usage profiles* of DBT. Some DBT is employed within the parameters of a *clinical context*, in which a healthcare provider oversees use of the technology. For instance, the Proteus ingestible sensor, which is embedded in pills and tablets, allows healthcare providers to monitor their patients' medication adherence [40]. A further usage profile of DBT is the *research context*. This includes both research that is conducted in academic settings and research by those developing DBT intended for sale, either with or without FDA approval. A third usage profile of DBT is the *direct-to-consumer* (DTC) context. As with clinical use, DBT is used here in order to "treat" the user (in some broad sense), but this is not done within the parameters of a clinical, provider–patient relationship. In this context, users simply purchase DBT products, the vast majority of which are not FDA-regulated, and apply them as they themselves see fit [2, 41]. The final usage profile is the *third-party* context. In this arena, DBT is used by a party that is not a healthcare provider, a researcher, nor the individual herself, in order to track or affect other individuals. Employee wellness programs, schools monitoring students, military applications all fall within this usage profile.

The focus of this chapter is self-optimization and so the primary usage profile with which I am concerned is DTC use. DTC DBT has gained hugely in popularity in the past 5–10 years. Reports by market research firms show that between 2010 and 2014 there was a 500% increase in the number of non-invasive neurotechnology patents filed [42]. The global wearables industry alone was valued at 32.63 billion

US dollars in 2019, and is expected to grow to over 100 billion US dollars by 2021 [43]. This growth suggests a perception on the part of the user that DTC DBT is useful in achieving one's goals [44]. The next section will discuss how this technology is implicated in the pursuit of self-optimization.

9.3 Digital Behavioral Technology and Self-Optimization

As described above, DBT can be used in the context of different usage profiles. Self-optimization is generally done within the DTC context, although it can sometimes occur in the third-party and even the clinical relationship contexts as well. Individual consumers are increasingly using DBT in order to optimize themselves. This is occurring as patients and consumers take on more of the burden of their own health and well-being. Sometimes referred to as the "democratization of healthcare," numerous observers [45] [46] have argued that technology is contributing towards the "responsibilization" of patients and consumers. Thanks to DBT, individuals now have the capability of using technology for the purpose of monitoring their own health and well-being, and use such technology to attempt to improve these [47]. This *capability* has contributed to an *expectation* that individuals now have a responsibility of improving their own health and well-being, as we will see below.

Much of the language around DBT includes this exhortation towards self-optimization. In the media, DBT is praised as a means to take on responsibility for improving oneself. Headlines like "This New Generation of Wearables Empowers People to Take Charge" [48] and claims that "wearables empower 'busy lives' to develop a more responsible approach towards themselves" [49], illustrate how DBT is perceived as increasing consumers' agency in their quest for well-being and health.

In the bioethics literature the term "e-patient," short for "empowered patient," has been coined to describe "health consumers participating fully in their medical care" [50, p. 2]. On this view of patients, individuals have an obligation to be informed about conditions and treatment options. Access to information gives patients a responsibility to take control in medical decision-making. The acquisition of this information has been facilitated by technology, specifically DBT. Thus, the "e" in "e-patient" has come to stand for "'electronic', 'equipped', 'enabled', […] 'engaged' or 'expert'" [50, p. 2]. As Schmietow & Marckmann [46] note, "[s]elf-empowerment turns into a self-obligation to be 'digitally engaged' and at the same time expresses a shift of priorities from externally induced healthcare to a more elusive health and self-management" (p. 627).

Sociologist Deborah Lupton has documented this embrace of self-optimization. In her extensive research on *self-quantification*, i.e., the phenomenon of individuals embracing the tracking functionality of DBT (see Table 9.1), she identifies a desire of improving the self as a central focus of self-tracking activities that are designed to radically expand self-knowledge [45, 51, 52]. As a subject from one of her interviews puts it: “Unless something can be measured, it cannot be improved” [45, p. 67]. Another states, “[y]ou want to be your best self. […] It’s studying yourself as an interesting topic in ways that you couldn’t study yourself before […] this is just giving you self-awareness into previously invisible aspects of your life” [45, p. 65]. Lupton describes this as a practice “of self-hood that conforms to cultural expectations concerning self-awareness, reflection and taking responsibility for managing, governing oneself and improving one’s life chances” [45, p. 68]. “Self-tracking therefore represents the apotheosis of the neoliberal entrepreneurial citizen ideal” [45, p. 68].

This conception of the self is in line with the existential notion of the self as something to be fashioned or created. In his analysis of the Nietzschean self, Anderson [53, p. 229] describes this concept of self-hood as something normative, i.e., a *task* that one is continually setting for oneself. In this way, the self is not a static component of an individual, but is constantly fashioned through the actions an individual undertakes. As Anderson notes, there appears to be a paradox in this notion of self-creation, for surely the thing being created must already exist to do the creating. But this paradox, he points out, can be dissolved if one distinguishes between a descriptive conception of the self, that carries out the plan of self-creation [54], and a normative conception of an ideal self that is the telos of one’s self-fashioning pursuit [53]. In short, self-optimization as it is embraced by users of DBT (and as it is advertised by its manufacturers) presupposes something like the existential conception of a self that is continually being created [2].

Rather than focusing on the social determinants of health and well-being, DBT shifts the onus of responsibility to the individual. Here the assumptions are that this individual ought to be both equipped and motivated to take up DBT in pursuit of optimization. As becomes clear from Lupton’s research, individuals see this technology as central in the project of fashioning themselves in the existential sense discussed above. DBT is used to craft the body and mind in a way that the individual endorses, as individuals believe that using this technology enables them to create the conditions to become the kinds of persons they want to become.

Such an outlook is grounded in a belief of what matters. It matters to be the best version of oneself, to be optimal. The idea of “working on the self” is central to how users of DBT see themselves in the world. But it is also embedded in a larger culture of productivity. Individuals who are responsible for all aspects of their goings-on populate the workplace. Productivity in the workplace is enmeshed with productivity at home. The same devices that are used for productivity at work are used in the private sphere and vice versa.

Thus, the picture that emerges is one in which DBT enables self-optimization to a point where individuals are expected to take on responsibility for all dimensions of the self. I identify three avenues of self-optimization: *information*, *parameterization of behavior*, and *direct interaction*.

9.3.1 Information

Individuals can use DBT to gain *information* about various dimensions of themselves. Specifically, DBT provides information on physiological (e.g., heart rate [9]), psychological (e.g., mood logging [14]), and neurological (e.g., EEG [55]) features of the self. Often, such information is then used by individuals to adjust behavior in ways to favorably affect these dimensions. Lupton talks about individuals achieving "knowledge, awareness, problem-solving" [56].

This kind of tracked (or logged) self-knowledge is a departure from how we standardly acquire self-knowledge. It encourages the user to gather information about the self through the processing of quantitative representation, rather than gaining self-knowledge through embodied situated unconscious cognition. This places the process of gaining information about the self on par with that of gaining information about objects external to the self. As such, the individual takes a third-person approach to herself, as she encounters her body, mind, and brain as a quantifiable object that permits of manipulation [2].

9.3.2 Parameterization of Behavior

A further use of DBT is the *parameterization of behavior*. DBT can issue signals or alarms when certain tracked values fall outside of desired parameters. For example, fitness trackers can alert an individual when her heart rate drops below a certain value. Signals can be used in this way to help users refrain from behaviors they find undesirable. Alternatively, users can also be rewarded, such as with badges or points, when values are within desired parameters. These methods give users incentives to behave well and disincentivizes poor behavior.

One technology that explicitly utilizes the principles of conditioning is the Pavlok 2 [57]. The Pavlok 2 is an aversive conditioning device, that emits small electric shocks when a user engages in behavior, e.g., nail-biting, smoking, eating sweets, sleeping too late, or spending too much time on time-wasting websites, that she is seeking to curtail. In this way, it aims to reinforce desirable behavior traits.

VR is a functionality of DBT that can also be used for behavior parameterization. VR generates alternative visual and auditory phenomenological experiences and can be extended to generate embodied virtual reality experiences including "tactile" or "haptic responsiveness." For example, the Teslasuit consists of a full body suit that uses nerve and muscle stimulation to generate haptic sensations for a fully bodily immersive VR experience [26]. One use of VR is in the addiction recovery. Patients can practice saying no to drugs in triggering environments, such as crack houses or bars [58]. While such addiction recovery is usually performed within the usage profile of a clinical relationship, there are many DTC applications that are either already being used or may be used in the future. For instance, consumers can use VR in a DTC setting to attend to their nicotine cravings. Alternatively, VR can *gamify* one's fitness routine by allowing a user to immerse herself into an alternate reality where she is boxing with a virtual opponent [59].

9.3.3 Direct Interaction

An additional use of DBT that contributes to self-optimization is *direct interaction*. For instance, DTC neurostimulation devices directly stimulate (or inhibit) parts of the brain in an attempt to impact brain function. Much of this direct interaction use is based on speculative scientific claims. For instance, technology that seeks to harness the effects of non-invasive vagus nerve stimulation to dampen the sympathetic nervous system response claims to enhance focus, promote positive thinking, and curb cravings [60]. In these cases, the user attempts to self-optimization not through conscious action, but rather through direct intervention in the relevant brain region. Various other neurostimulation technologies, such as the tDCS BrainStimulator, also operate in this way [4]. A further form of direct interaction is DBT that delivers medication. For instance, wearable devices by Chrono Therapeutics assist individuals in quitting smoking by monitoring nicotine levels in the blood and administering nicotine transdermally when individuals most require it [24].

9.4 Digital Behavioral Technology and Artificial Intelligence

DBT increasingly involves artificial intelligence (AI) in order to perform the functionalities mentioned above. The AI employed in this arena is trained through deep learning. Deep learning is a class of machine learning, in which an artificial neural network extracts patterns from data with which it is supplied. Deep learning extracts these patterns at multiple layers of abstraction, ranging from the specific to the more abstract. Given enough data and a large enough network, these networks can learn very complex patterns, such as recognizing faces from visual content, meaningful elements from natural language, and medical conditions from health data. Moreover, AI learns inductively from experience. Algorithms are iteratively updated when new data are provided. Such updating occurs without being explicitly programmed by human programmers. Rather, neural networks absorb new data and adjust connection weights between nodes in a stochastic fashion. This means both that the structure of neural networks is entirely dependent on the data that are inputted, and that there are no explicit programming rules discernable by humans [61–63].

Deep learning functions thanks to the availability of *big data*. Data scientists from IBM describe big data as being made up of four dimensions: *volume, variety, velocity*, and *veracity* [64]. The *volume* of available data is staggering. 2.3 trillion gigabytes are generated every day [65], with 90% of all currently existing data having been generated in the last 2 years. Experts believe 1.7 megabytes of data are created every second for every person on earth [65].

These data come from a *variety* of sources. Five new Facebook profiles are created every second and more than 300 million photos are uploaded to Facebook every day [66]. Every minute, 16 million text messages are sent [66]. At the same time, the number of connected wearable devices worldwide has increased from 325 million in 2016 to 722 million in 2019, with forecasts predicting this number to reach one billion by 2022 [67]. Meanwhile, the internet of things (IoT) has grown from two billion connected objects in 2006 to 200 billion in 2020 [68]. Moreover, an estimated 2.3 trillion gigabytes of electronic healthcare record data were produced in 2020 [69].

Different kinds of DBT generate different kinds of data, from location information to EEG, from photographic content to blood alcohol levels. Given the diversity of sources, combining the variety of data can yield extremely well-rounded representations of individuals and populations. The only way such large quantities of data can be processed in time is through powerful transaction processing systems (TPS). This is captured in the *velocity* of the four Vs.

Finally, *veracity* refers to the quality (accuracy and applicability) of the data. Much of the data that are available are of poor quality. They are inaccurate, incomplete, and even inconsistent. Data often needs to be cleaned so that so-called "dirty data" are kept from accumulating [70].

Thanks to the availability of AI, today's DBT user need no longer be satisfied with an n-of-1 trial, where she tracks or logs her own physiological, neurological, or emotional goings-on and pores over them in an attempt to discover behaviors that are conducive to self-optimization. Today's DBT is "smart" in that the wealth of data that are produced by the DBT is combined with enormous amounts of data produced by other devices, including data integrated across different devices and platforms, in order to be fed into powerful machine learning programs that use deep learning to glean insights from the combined data. Thus, the data from an individual's wrist-worn fitness tracker, or an individual's EEG device, or an individual's location are just a small fraction of an ocean of information that reveals much more powerful insights about aspects of people's selves (their bodies, minds, environment), than any individual could glean on her own.

There are multiple ways in which AI can be integrated with DBT. First and foremost, DBT uses AI for *data analytics*. As noted above, the vast quantity of data requires deep learning in order to glean usable insights. In turn, deep learning makes such data incredibly valuable. Companies learn a considerable amount about individuals' health, behavior, activities, and beliefs from such AI-aided data analysis [71]. These insights not only reveal patterns at the population level, but also the individual level [72]. In addition to companies using this information in their own product and service development, it is standardly sold to other companies, such as marketing and health insurance companies. Moreover, with the increasing involvement of "Big Tech" in the healthcare arena, companies like Google are mining medical records [73], to expand their reach.

A further way AI may be integrated with DBT is for *restorative purposes*. Deep learning can be harnessed to allow individuals who are blind to regain some of the capabilities of sighted people. For example, AIServe [74] combines computer vision and AI to provide users with a wearable device that analyzes the environment, detects different elements in the surroundings, such as bikes, cars, and people, and then gives voice navigation instructions to the user. Deep learning is also employed in hearing aids. Thanks to AI, devices can learn what kind of environments a user tends to be in and learn to filter out desirable sounds from non-desirable ones, e.g., an interlocutor's voice from background noise [75]. AI-enhanced hearing aids can also directly translate language, transcribe what is being heard and said by the user, monitor brain functioning, and interface with a smart assistant [75].

Beyond restoration, AI-enabled DBT may also be used for *enhancement* purposes. Neurostimulation devices, such as the BRAINtellect 2, may be worn during sleep and employ AI to translate the user's brain waves into engineered music-like sound waves that are believed to enhance memory, learning, and wellness [76].

AI is increasingly employed to improve *VR* and augmented reality *experiences*. Thanks to deep learning, user preferences and virtual environment layouts can be updated in real time. For example, the Hololens2 is a deep learning enabled mixed reality system that generates constantly updated realities for the user which allow users to visually and tactilely interact with objects in their environment [77]. Hololens2 can also create avatars that deliver the words an individual is speaking in a different language [78].

AI is also used to enhance DBT to include *assistants*. In a previous section, I outline the various assistant functionalities. Given the self-optimization purpose of DBT, AI-enabled assistants are employed for health and wellness recommendations. AI is used to integrate vast quantities of data from DBT and other sources to provide recommendations to the user. Companies like LifeQ, for instance, use deep learning to develop models and algorithms that translate data from wearables and mobile apps into usable information. LifeQ provides individualized information for consumers on how to modify their behavior in real time to achieve their health and well-being goals, for insurance companies and corporate wellness programs on the health risks of individuals, and for clinicians on the treatment options, progress, and compliance of their patients [79]. Other assistants go beyond health and wellness, integrating advice from both "life and work." Galahad AI has introduced a virtual personal assistant, VYou, that helps individuals forgo short-term temptations in favor of long-term goals. Explicitly set up to allow users to engage in self-optimization, VYou "empower[s] people to better manage the most important and challenging aspects of their personal lives including their time, health, relationships, and money" [80].

Clearly, there are many ways in which AI is integrated into DBT. Particularly interesting from the perspective of self-optimization are smart assistants that advise and guide users on their quest to better themselves. Such assistants may even be enhanced via "affective computing"—a form of computing that allows assistants to deliberately involve and influence emotions and other affective phenomena [81].

Users gain beliefs about themselves through AI-enabled DBT that go beyond data captured by the technology. The technology takes on the role of guide and advisor, and can impact a user's reflective beliefs of self-worth. It can prompt a user to feel good about herself, for instance when a DBT assistant explicitly commends her for her behavior, or implicitly when the user recognizes that her behavior is within desirable parameters. An individual can also be made to feel poorly by DBT, either by being explicitly admonished by an assistant, or by finding herself failing at remaining within measurable parameters. Affective computing may increase this effect.

Of note, deep learning not only affects the insights that are gleaned from this technology, but in turn also affects the very way DBT functions. The AI-driven insights from data are fed into the functioning of the devices themselves. For instance, as a technology learns patterns from the behavior of a particular user along with the behavior of all the other users, it updates the parameters of what counts as "normal" and even "desirable" behavior.

9.5 Problems with Artificial Intelligence

As described, AI is a driving force for this technology. But there are well-documented flaws with deep learning. Learning algorithms are not static, yielding results through mathematical models. Rather they are iterative, constantly updating themselves as a result of the inputted data [1]. As already noted, inputted data are increasingly emerging from more and more sources, with the internet of things contributing to the well-spring of big data. This iterative process is vulnerable to systemic bias. If the data on which an algorithm is trained and which it uses to update itself contain bias, then the insights that such algorithms yield will be systematically biased. This can shift parameters of what may count as "normal," "correct," or "desirable."

Such algorithmic bias has been seen in healthcare decision-making—where black patients' needs are underestimated compared to those of white patients [82], and job hiring—where women's resumes were systematically scored lower than those of comparable men [83].

A further concern about AI is algorithmic opacity. Deep learning is not built on explicit theoretical rules inputted by human users. Rather, the internal workings of such an AI system are a "black box," that operates in whatever way yields the appropriate results. This means that users have no way of knowing how or why a certain system generates a particular result [84]. It also means that beyond pointing to an algorithm that has yielded correct results in other instances, no justification can be provided for a particular result in a particular instance [85].

A widely cited example of problematic algorithmic opacity is the Correctional Offender Management Profiling for Alternative Sanctions (COMPAS) tool. This is a decision support tool intended to help US judges to predict the likelihood of a defendant's risk of recidivism [86]. Judges can use a defendant's risk score to determine an appropriate length and severity of sentence. However, an investigation by ProPublica showed that COMPAS systemically overestimates the risk of recidivism for black offenders and underestimates the risk of recidivism for white offenders, even though race is not a data point that is available to the deep learning algorithm [87, 88]. But since COMPAS's inner workings are hidden from users, defendants have no recourse for arguing that in their case, the algorithm's results are unjustified. Moreover, since programs such as COMPAS are proprietary and thus subject to trade secrecy, defendants cannot even find out what inputted data were used in determining their recidivism score [89].

A further concern regarding AI is reward hacking [1]. The principle way in which so-called AI "agents" (i.e., algorithms) learn is through reinforcement learning. As in Skinnerian behaviorist psychology, correct behavior is sought to be reinforced. In reinforcement learning, an agent performs an action and is provided with feedback for that action from the environment. The feedback after each action is provided as a “reward” or a “punishment,” where the size of the reward or punishment is a measure of how close the agent’s action came to the correct action [90]. The goal of the AI agent is to maximize its rewards and minimize punishment. This way, through trial-and-error, the agent iteratively hones in on the correct result. However, this behaviorist approach can be gamed. An often-cited example is a cleaning robot who is rewarded for minimizing the amount of dirt it sees. Such a robot can simply close its eyes and thus receive a reward, even though this obviously does not fulfill the purpose intended by the human designer [91].

The problems of algorithmic bias and reward hacking are tenacious. Even when a given instance of bias or hacking is rectified, this provides no guarantee that a further instance will not arise. The opacity with which deep learning operates makes any attempt to correct structural biases or failures difficult.

9.6 Possible Effects of Problems with AI on DBT’s Role in Self-Optimization

Given AI’s prevalence in DBT, the aforementioned problems with AI may well manifest themselves in the self-optimizing function of DBT. I will look at how algorithmic bias, algorithmic opacity, and reward hacking affect the three avenues of self-optimization identified previously, i.e., direct interaction, information, and parameterization of behavior.

As mentioned above, neurostimulation devices use *direct interaction* for enhancement purposes. Tracked EEG sleep data is processed through deep learning algorithms in order to generate and emit music-like waves intended to improve wellness and cognition. Meanwhile devices like the Neuvana stimulate the Vagus nerve in order to enhance mental acuity, promote positive thinking, or curb cravings [60]. As AI is increasingly used by such devices, the risk of algorithmic bias becomes more significant. What the AI recognizes as a desirable EEG state will depend on a big data set sourced from many users over many individual instances. But any biases that inhere in the incoming data will be reflected in the parameters that the algorithm identifies as “normal” or “desirable.” Perhaps users of this technology are more likely than the general population to aim for mental acuity or heightened concentration—a supposition that is plausible given that such individuals are more likely to employ such technology in the first place. As a result, parameters “endorsed” by the AI may be skewed. This is problematic, because an EEG state that is in fact representative of the general population would be identified as inadequate or undesirable. With the technologies discussed here, this may result in individuals receiving neurostimulation at times or in ways that are not in fact beneficial. Egregious instances of unwarranted neurostimulation may be rather noticeable and thus may quickly be weeded out. But slight shifts in desirable parameters may be less obvious and thus more insidious as they may contribute to a skewing of what is considered normal. Beyond neurostimulation, such an effect can be problematic within the self-optimization modes of *information* and *parameterization of behavior*.

Self-optimization through *information* relies heavily on the data analytics action of the AI. Algorithmic bias that corrupts such analytics may lead to a number of problematic consequences. As described previously, vocal analysis is used to identify a range of conditions, including which users suffer from depression or are likely to default on a loan. Algorithmic bias in such functionalities can lead to individuals being categorized as suffering from a condition when in fact they are not, or being categorized as not suffering from a condition when in fact they are. This could result both in users being discriminated against, and in users not receiving the assistance that the technology is intended to provide. If we consider the use of vocal analysis in determining something like an individual's likelihood to default on a loan, it is easy to see how algorithmic bias could yield unfair results on the basis of non-praiseworthy and non-blameworthy properties, much like was the case with the COMPAS tool in criminal contexts. Vocal analysis might indirectly discriminate against individuals with certain properties. This can occur even if these properties are ommitted from the data, because they are considered too sensitive. For instance, an individual's socioeconomic status (SES) may not be a data point that is available to an algorithm, because it might be discriminatory. However, individuals of low SES reliably have higher rates of smoking, of being exposed to secondhand smoke, and of being sick from smoking-related diseases than individuals of higher SES [92]. Whether an individual smokes or is exposed to secondhand smoke is reflected in that individual's vocal qualities, and is thus useable data for analytic algorithms. Consequently, even if the data available to an AI is free of any mention of SES, the algorithm may still carve up the population along SES dimensions, thanks to correlated proxy features. Moreover, because of algorithmic opacity, affected parties may have no way of knowing how an algorithm arrives at a particular categorization, and thus cannot assess whether such results are justified in a given case.

As noted previously, vocal analysis is also used for restorative purposes, for instance in AI-enabled hearing aids that filter out voices from background noise. Certain vocal features are reliably correlated with racial groups [93]. Algorithmic bias may lead to certain kinds of voices, possibly the voices of certain racial groups, being less audible to users of smart hearing aids than others. If we assume that in general people prefer talking to people whom they can understand well, then such bias might subtly affect the kinds of people with whom users choose to spend time. What's more, this effect may happen unbeknownst to users, who may not be able to point to why they are engaging in conversation with one group of people rather than another.

Alternatively, with AI-enabled translation performed by such smart hearing aids, thanks to algorithmic bias, translated utterances may contain words that inadequately express the speaker's intended semantic content. While this may merely be frustrating if it occurs in a morally neutral way, it is worrisome if utterances exhibit unintended discriminatory or offensive language. The latter has been shown by Microsoft's chatbot Tay to be a real risk [94]. Not only might this portray users as unfairly racist, sexist, or prejudiced in some other way, with prolonged use it may even erode a user's beliefs about her own views, as I have argued elsewhere [95].

The clearest example of how the flaws in AI may clash with the self-optimization function of DBT is within the mode of *parameterization of behavior*. As already mentioned, algorithmic bias can cause parameters of what counts as "normal" or "desirable" to shift. This is particularly the case with AI-enabled tracking technology, which is designed to identify appropriate values or value-ranges of trackable dimensions (see Table 9.2) and alert, praise, punish, or reward the user on the basis of her adherence to those values or value-ranges. We see this, for instance, in the Pavlok 2 that administers a small electric shock to the user if she falls outside of certain parameters, for instance if she exercises for too short of a duration. An AI-enabled version of such behavior-parameterizing tracking technology will obviate the need for human input on what constitute acceptable parameters, e.g., an appropriate duration for exercising, by arriving at such parameters on the basis of big data, including from users of the technology and a wide array of other sources. Then, this DBT can simply punish a user if she deviates from what the algorithm has deemed as desirable.

Thanks to algorithmic bias, however, the parameters towards which the DBT is steering the user may not actually correspond to parameters that should be sought after. Moreover, because of algorithmic opacity, a user may not be able to see why parameters are what they are, and thus why she falls short. This might result in users feeling unjustifiably discouraged about their performance, or conversely unwarrantedly accomplished. The ramifications of such effects will vary. For some users falsely appearing to fall outside of parameters or erroneously appearing to fall within them will have little consequence. For others who perhaps place great importance on being within certain bounds, wrongful categorization may be unduly burdensome. Users who exhibit compulsive behaviors, such as individuals with eating or exercising disorders, already employ DBT to contribute to their disorders [96, 97]. Skewing parameters of what counts as "normal" or "desirable" could exacerbate this harmful phenomenon.

Beyond the DTC context, such skewing of parameters may be problematic in the third-person context. When tracking technology is used by corporate wellness programs to dole out rewards and punishments [98], or by insurance companies to determine premiums [99], miscategorizations of users may lead to injustices in the same vein that have occurred with the COMPAS tool. Moreover, just as with COMPAS, the presence of algorithmic opacity makes redressing any such injustices hard if not impossible.

Finally, smart assistants are becoming widespread amongst AI-enabled DBT. As noted earlier, there are life assistants designed to help users with a huge array of tasks, including determining whether their behavior is beneficial or detrimental to productivity [37], whether their behavior will help them achieve their health and well-being goals [79], and whether they are successfully forgoing short-term temptations in favor of long-term goals [80]. Such DBT also assists insurance companies and corporate wellness programs in determining the health risks of individuals [79].

However, such smart assistants may be vulnerable to reward hacking. Assistants are “rewarded” for giving “good” “instructions and recommendations” [38] to users, where “good” is determined by running algorithms that mine big data from a plethora of sources. But honing in on what exactly constitutes “good” instructions and recommendations may be algorithmically burdensome. It may, for some algorithmically opaque reason, be easier to generate instructions and recommendations that can appear as “good,” but are actually not. This is an instance of reward hacking. When the AI manages to find such “imposter” results, it receives its reward and has thus fulfilled its purpose. Of course, while this may be beneficial from the perspective of the AI, it is not beneficial at all for the human user. But again, weeding out such reward hacking issues is aggravated by algorithmic opacity. As humans rely more on their smart assistants to guide them in their pursuits of healthier, happier, more productive lives, issues of reward hacking will become more and more serious, as they threaten to undermine individuals’ autonomy in fashioning themselves. Similarly, when corporate wellness reward and punishment structures, as well as insurance premiums, depend on a user’s adherence to assistant-provided instructions and recommendations, reward hacking can contribute to an unjust distribution of burdens and benefits.

9.7 Conclusion

AI-enabled DBT has enormous potential to affect the way users engage in activities of self-optimization. Individuals are increasingly taking on the burden of engineering their own health, wellness, and productivity in explicitly engaged and active roles. This surging involvement is predicated on a belief that technology now exists to facilitate this effective self-fashioning. Algorithmic bias, algorithmic opacity, and reward hacking undermine this pursuit, often in ways that are unknown to the user. If individuals truly are to be empowered, these issues with deep learning must be addressed.

References

1. Osoba OA, IV WW. An Intelligence in our image: the risks of bias and errors in artificial intelligence. Rand Corporation; 2017.
2. Kreitmair KV. Dimensions of ethical direct-to-consumer neurotechnologies. AJOB Neurosci. 2019;10(4):152–66. https://doi.org/10.1080/21507740.2019.1665120.
3. Fitbit official site for activity trackers and more. n.d. https://www.fitbit.com/home. Accessed 15 Dec 2019.
4. The Brain Stimulator tDCS Devices—Shop and stimulate your life today! n.d. https://thebrainstimulator.net/. Accessed 15 Dec 2019.
5. Brunelin J, Mondino M, Bation R, Palm U, Saoud M, Poulet E. Transcranial direct current stimulation for obsessive-compulsive disorder: a systematic review. Brain Sci. 2018;8(2):37. https://doi.org/10.3390/brainsci8020037.
6. Strava. Run and cycling tracking on the social network for athletes. n.d. https://www.strava.com/. Accessed 15 Feb 2020.

7. Qardio. Qardio official store. Qardio. n.d. https://store.getqardio.com/. Accessed 15 Feb 2020.
8. Muse—Meditation made easy. Muse. n.d. https://choosemuse.com/. Accessed 15 Feb 2020.
9. Fitbit PurePulse® continuous wrist-based heart rate. n.d. https://www.fitbit.com/purepulse. Accessed 9 Jan 2020.
10. Chu M, Nguyen T, Pandey V, Zhou Y, Pham HN, Bar-Yoseph R, Radom-Aizik S, Jain R, Cooper DM, Khine M. Respiration rate and volume measurements using wearable strain sensors. Npj Digit Med. 2019;2(1):1–9. https://doi.org/10.1038/s41746-019-0083-3.
11. BACtrack Skyn™—The World's 1st wearable alcohol monitor. n.d. BACtrack Skyn. https://skyn.bactrack.com/. Accessed 15 Feb 2020.
12. Many drugs could come equipped with ingestible sensors. Pharmacy Times. n.d. https://www.pharmacytimes.com/contributor/timothy-aungst-pharmd/2018/05/many-drugs-could-come-equipped-with-ingestible-sensors. Accessed 15 Feb 2020.
13. Bariya M, Nyein HYY, Javey A. Wearable sweat sensors. Nat Electron. 2018;1(3):160–71. https://doi.org/10.1038/s41928-018-0043-y.
14. LLC Y. Mood tracker by eMoods—Advanced Journal & Reporting Insights for Bipolar. EMoods - easy mood charting for bipolar/manic depression and other mood disorders. n.d. https://emoodtracker.com. Accessed 15 Dec 2019.
15. Semistry. Sexual statistics and measurements. n.d. https://www.semistry.com/. Accessed 15 Dec 2019.
16. Cogito spins out Companion app to detect depression in voice—Business Insider. n.d. Retrieved December 15, 2019, from https://www.businessinsider.com/goldman-salesforce-cogito-companion-detects-depression-in-voice-2018-12. Accessed 15 Dec 2019.
17. Voicesense. n.d. https://www.voicesense.com/. Accessed 15 Dec 2019.
18. How to use google lens to identify objects using your smartphone. Digital Trends. n.d. https://www.digitaltrends.com/mobile/how-to-use-google-lens/. Accessed 15 Dec 2019.
19. Roepke AM, Jaffee SR, Riffle OM, McGonigal J, Broome R, Maxwell B. Randomized controlled trial of superbetter, a smartphone-based/internet-based self-help tool to reduce depressive symptoms. Games Health J. 2015;4(3):235–46. https://doi.org/10.1089/g4h.2014.0046.
20. Thync Relax Pro review. Wareable. 2017. https://www.wareable.com/wearable-tech/thync-relax-pro-review.
21. Go flow 4mA tDCS and so-tDCS stimulator by Foc.us. Focus - take charge. n.d. https://foc.us/go-flow/. Accessed 9 Feb 2020.
22. Brunoni AR, Nitsche MA, Bolognini N, Bikson M, Wagner T, Merabet L, Edwards DJ, Valero-Cabre A, Rotenberg A, Pascual-Leone A, Ferrucci R, Priori A, Boggio P, Fregni F. Clinical research with transcranial direct current stimulation (tDCS): challenges and future directions. Brain Stimul. 2012;5(3):175–95. https://doi.org/10.1016/j.brs.2011.03.002.
23. Landhuis E. Do D.I.Y. brain-booster devices work? Scientific American. n.d. https://www.scientificamerican.com/article/do-diy-brain-booster-devices-work/. Accessed 18 Dec 2019.
24. Chrono Therapeutics. Chrono therapeutics. n.d. https://www.chronothera.com. Accessed 29 Nov 2019.
25. Oculus Go: standalone VR headset. Oculus. n.d. https://www.oculus.com/go/. Accessed 18 Dec 2019.
26. Teslasuit does full-body haptic feedback for VR. Engadget. n.d. https://www.engadget.com/2016/01/06/teslasuit-haptic-vr/. Accessed 29 Nov 2019.
27. Google assistant. n.d. https://assistant.google.com/explore. Accessed 18 Dec 2019.
28. Software S. Baby connect: baby tracker and log for android, iPhone, iPad, Kindle and for the web. n.d. https://www.babyconnect.com/. Accessed 18 Dec 2019.
29. Baby connect. Google assistant. n.d. https://assistant.google.com/services/a/uid/0000006d500622ec?hl=en-US. Accessed 18 Dec 2019.
30. Symptomate – Check your symptoms online. n.d. https://symptomate.com. Accessed 18 Dec 2019.
31. Symptomate. Google assistant. n.d. https://assistant.google.com/services/a/uid/000000a04396478b?hl=en-US. Accessed 18 Dec 2019.

32. NEUROGIXS. n.d. http://neurobeat-a.com/product/product.php?ptype=view&prdcode=1811280001&catcode=10000000&page=1&searchopt=&searchkey=. Accessed 18 Dec 2019.
33. Care pod. Google assistant. n.d. https://assistant.google.com/services/a/uid/000000e0fa0c2605?hl=en-US. Accessed 8 Jan 2020.
34. Good morning routine. Google assistant. n.d. https://assistant.google.com/services/a/uid/0000004f0e04d1da?hl=en-US. Accessed 8 Jan 2020.
35. Bedtime routine. Google assistant. n.d. https://assistant.google.com/services/a/uid/0000001807f0256f?hl=en-US. Accessed 8 Jan 2020.
36. Controlicz—The voice of your Domoticz Home Automation System. n.d. https://www.controlicz.com/. Accessed 8 Jan 2020.
37. Flow Harmonics. n.d. https://flowharmonics.com/works.html. Accessed 8 Jan 2020.
38. Flow Harmonics. n.d. https://flowharmonics.com/people.html. Accessed 8 Jan 2020.
39. Girlfriend Maya. Google assistant. n.d. https://assistant.google.com/services/a/uid/000000efa718ac50?hl=en-US. Accessed 8 Jan 2020.
40. Belknap R, Weis S, Brookens A, Au-Yeung KY, Moon G, DiCarlo L, Reves R. Feasibility of an ingestible sensor-based system for monitoring adherence to tuberculosis therapy. PLoS One. 2013;8(1):e53373. https://doi.org/10.1371/journal.pone.0053373.
41. Illes J. Neuroethics: anticipating the future. Oxford University Press; 2017.
42. SharpBrains. Market report on pervasive neurotechnology: a groundbreaking analysis of 10,000+ patent filings transforming medicine, health, entertainment and business. 2018. https://sharpbrains.com/pervasive-neurotechnology/. Accessed 24 May 2021.
43. Grand View Research. Market analysis report: wearable technology market size, share & trends analysis report by product (wrist-wear, eye-wear & head-wear, foot-wear, neck-wear, body-wear), by application, by region, and segment forecasts, 2020–2027. 2020. https://www.grandviewresearch.com/industryanalysis/wearable-technology-market. Accessed 26 May 2021.
44. Pfeiffer J. Quantify-me: consumer acceptance of wearable self-tracking devices. 16. 2016.
45. Lupton D. The quantified self. Wiley; 2016.
46. Schmietow B, Marckmann G. Mobile health ethics and the expanding role of autonomy. Med Health Care Philos. 2019;22(4):623–30. https://doi.org/10.1007/s11019-019-09900-y.
47. Sharon T. Self-tracking for health and the quantified self: re-articulating autonomy, solidarity, and authenticity in an age of personalized healthcare. Philos Technol. 2017;30(1):93–121. https://doi.org/10.1007/s13347-016-0215-5.
48. McGillin F. This new generation of wearables empowers people to take charge. The Doctor Weighs In. 2017. https://thedoctorweighsin.com/this-new-generation-of-wearables-empowers-people-to-take-charge/.
49. Day1Tech. How Wearables are transforming the Healthcare Industry. Medium. 2019. https://medium.com/@Day1Tech/how-wearables-are-transforming-the-healthcare-industry-e103f7985c8.
50. Meskó B, Radó N, Győrffy Z. Opinion leader empowered patients about the era of digital health: a qualitative study. BMJ Open. 2019;9(3):e025267. https://doi.org/10.1136/bmjopen-2018-025267.
51. Sumartojo S, Pink S, Lupton D, LaBond CH. The affective intensities of datafied space. Emot Space Soc. 2016;21:33–40. https://doi.org/10.1016/j.emospa.2016.10.004.
52. Lupton D. How do data come to matter? Living and becoming with personal data. Big Data Soc. 2018;5(2):2053951718786314. https://doi.org/10.1177/2053951718786314.
53. Anderson RL. What is a Nietzschean self? 1. Oxford University Press. 2012. https://www.oxfordscholarship.com/view/10.1093/acprof:oso/9780199583676.001.0001/acprof-9780199583676-chapter-9. Accessed 15 Apr 2020.
54. Sartre J.-P. Being and nothingness (Barnes HE, Trans.; Original ed. edition). Washington Square Press; 1993.
55. Advanced EEG technology—backed by science. EMOTIV. n.d. https://www.emotiv.com/our-technology/. Accessed 9 Jan 2020.
56. Lupton D. 'It's made me a lot more aware': a new materialist analysis of health self-tracking. Media Int Aust. 2019;171(1):66–79. https://doi.org/10.1177/1329878X19844042.

57. The Science. n.d. Pavlok. https://pavlok.com/science/. Accessed 29 Nov 2019.
58. Strauss G. Virtual reality: a powerful new tool for addiction treatment centers. Medium. 2019. https://blog.limbix.com/virtual-reality-a-powerful-new-tool-for-addiction-treatment-centers-a9dc7437fdd9.
59. Top 15 best VR fitness games for a total body workout. n.d. https://www.vrfitnessinsider.com/top-15-best-vr-fitness-games-total-body-workout/. Accessed 29 Nov 2019.
60. Science: Vagus nerve stimulation - Neuvana. n.d. https://Neuvanalife.Com/. https://neuvanalife.com/science/. Accessed 5 Feb 2020.
61. LeCun Y, Bengio Y, Hinton G. Deep learning. Nature. 2015;521(7553):436–44. https://doi.org/10.1038/nature14539.
62. Helbing D. Societal, economic, ethical and legal challenges of the digital revolution: from big data to deep learning, artificial intelligence, and manipulative technologies. ArXiv:1504.03751 [Physics]. 2015. http://arxiv.org/abs/1504.03751.
63. Luxton DD, Anderson SL, Anderson M. Chapter 11—Ethical issues and artificial intelligence technologies in behavioral and mental health care. In: Luxton DD, editor. Artificial intelligence in behavioral and mental health care. Academic Press; 2016. p. 255–76. https://doi.org/10.1016/B978-0-12-420248-1.00011-8.
64. Infographic: the four V's of big data. IBM Big data and analytics hub. n.d. https://www.ibmbigdatahub.com/infographic/four-vs-big-data. Accessed 29 Nov 2019.
65. Becoming a data-driven CEO. Domo. n.d. https://www.domo.com/solution/data-never-sleeps-6. Accessed 29 Nov 2019.
66. Marr B. How much data do we create every day? the mind-blowing stats everyone should Read. Forbes. n.d. https://www.forbes.com/sites/bernardmarr/2018/05/21/how-much-data-do-we-create-every-day-the-mind-blowing-stats-everyone-should-read/. Accessed 29 Nov 2019.
67. Global connected wearable devices 2016–2022. Statista. n.d. https://www.statista.com/statistics/487291/global-connected-wearable-devices/. Accessed 29 Nov 2019.
68. The growth in connected IoT devices is expected to generate 79.4ZB of data in 2025, according to a new IDC forecast. IDC: The Premier Global Market Intelligence Company. n.d. https://www.idc.com/getdoc.jsp?containerId=prUS45213219. Accessed 29 Nov 2019.
69. The digital universe of opportunities: rich data and the increasing value of the internet of things sponsored by EMC. n.d. https://www.emc.com/leadership/digital-universe/2014iview/index.htm. Accessed 7 Feb 2019.
70. Big data volume, variety, velocity and veracity. InsideBIGDATA. 2013. https://insidebigdata.com/2013/09/12/beyond-volume-variety-velocity-issue-big-data-veracity/.
71. Emerging Issues Task Force, International Neuroethics Society. Neuroethics at 15: the current and future environment for neuroethics. AJOB Neurosci. 2019;10(3):104–10. https://doi.org/10.1080/21507740.2019.1632958.
72. Martinez-Martin N, Insel TR, Dagum P, Greely HT, Cho MK. Data mining for health: staking out the ethical territory of digital phenotyping. Npj Digit Med. 2018;1(1):68. https://doi.org/10.1038/s41746-018-0075-8.
73. Copeland R. WSJ news exclusive. Google's 'project nightingale' gathers personal health data on millions of Americans. Wall Street J. 2019. https://www.wsj.com/articles/google-s-secret-project-nightingale-gathers-personal-health-data-on-millions-of-americans-11573496790.
74. AiServe. Computer vision navigation. n.d. https://www.aiserve.co/. Accessed 19 Nov 2019.
75. Livio AI. StarkeyPro. n.d. https://starkeypro.com/products/wireless-hearing-aids/livio-ai. Accessed 19 Nov 2019.
76. Braintellect—B2v2. Sleep deeply and wake refreshed. n.d. https://braintellect.com/. Accessed 20 Nov 2019.
77. Making the HoloLens 2: Advanced AI built Microsoft's vision for ubiquitous computing. Innovation Stories. 2019. https://news.microsoft.com/innovation-stories/hololens-2-shipping-to-customers/.
78. Demo: The magic of AI neural TTS and holograms at Microsoft Inspire 2019. n.d. https://www.youtube.com/watch?time_continue=32&v=auJJrHgG9Mc. Accessed 22 Nov 2019.
79. Solutions – LifeQ. n.d. https://www.lifeq.com/solutions/. Accessed 22 Nov 2019.

80. Vyou. Galahad. n.d. https://galahadai.com/vyou/. Accessed 29 Nov 2019.
81. MIT Media Lab: Affective Computing Group. n.d. https://affect.media.mit.edu/. Accessed 22 Nov 2019.
82. Obermeyer Z, Powers B, Vogeli C, Mullainathan S. Dissecting racial bias in an algorithm used to manage the health of populations. Science. 2019;366(6464):447–53. https://doi.org/10.1126/science.aax2342.
83. Amazon scraps secret AI recruiting tool that showed bias against women. Reuters. 2018. https://www.reuters.com/article/us-amazon-com-jobs-automation-insight-idUSKCN1MK08G.
84. Burrell J. How the machine 'thinks': understanding opacity in machine learning algorithms. Big Data Soc. 2016;3(1):205395171562251. https://doi.org/10.1177/2053951715622512
85. Paudyal P, William Wong BL. Algorithmic opacity: making algorithmic processes transparent through abstraction hierarchy. Proc Hum Fact Ergon Soc Annu Meet. 2018;62(1):192–6. https://doi.org/10.1177/1541931218621046.
86. Corbett-Davies S, Pierson E, Feller A, Goel S. A computer program used for bail and sentencing decisions was labeled biased against blacks. It's actually not that clear. Washington Post. n.d. https://www.washingtonpost.com/news/monkey-cage/wp/2016/10/17/can-an-algorithm-be-racist-our-analysis-is-more-cautious-than-propublicas/. Accessed 10 Feb 2020.
87. Julia Angwin JL. Machine bias [Text/html]. ProPublica. 2016. https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing.
88. Jeff Larson JA. How we analyzed the COMPAS recidivism algorithm [Text/html]. ProPublica. 2016. https://www.propublica.org/article/how-we-analyzed-the-compas-recidivism-algorithm.
89. Wexler, R. Opinion. When a computer program keeps you in jail. The New York Times. 2017. https://www.nytimes.com/2017/06/13/opinion/how-computers-are-harming-criminal-justice.html.
90. Arulkumaran K, Deisenroth MP, Brundage M, Bharath AA. A brief survey of deep reinforcement learning. IEEE Signal Process Mag. 2017;34(6):26–38. https://doi.org/10.1109/MSP.2017.2743240.
91. Amodei D, Olah C, Steinhardt J, Christiano P, Schulman J, Mané D. Concrete problems in AI safety. ArXiv:1606.06565 [Cs]. 2016. http://arxiv.org/abs/1606.06565.
92. CDCTobaccoFree. Cigarette and tobacco use among people of low socioeconomic status. Centers for Disease Control and Prevention. 2019. https://www.cdc.gov/tobacco/disparities/low-ses/index.htm.
93. Baugh J. Linguistic profiling. Black Linguist. 2005; https://doi.org/10.4324/9780203986615-17.
94. Garcia M. Racist in the machine: the disturbing implications of algorithmic bias. World Polic J. 2016;33(4):111–7. https://doi.org/10.1215/07402775-3813015.
95. Kreitmair KV. Commentary: neuroprosthetic speech: pragmatics, norms, and self-fashioning. Camb Q Healthc Ethics. 2019;28(4):671–6. https://doi.org/10.1017/S0963180119000616.
96. Levinson CA, Fewell L, Brosof LC. My Fitness Pal calorie tracker usage in the eating disorders. Eat Behav. 2017;27:14–6. https://doi.org/10.1016/j.eatbeh.2017.08.003.
97. Plateau CR, Bone S, Lanning E, Meyer C. Monitoring eating and activity: links with disordered eating, compulsive exercise, and general wellbeing among young adults. Int J Eat Disord. 2018;51(11):1270–6. https://doi.org/10.1002/eat.22966.
98. Ajunwa I, Crawford K, Ford JS. Health and big data: an ethical framework for health information collection by corporate wellness programs. J Law Med Ethics. 2016;44(3):474–80. https://doi.org/10.1177/1073110516667943.
99. Strap on the Fitbit: John Hancock to sell only interactive life insurance. Reuters. 2018. https://www.reuters.com/article/us-manulife-financi-john-hancock-lifeins-idUSKCN1LZ1WL.