Grey Matter's Grey Areas: Privacy in the Age of Brain-Computer Interfaces

Sawsan Haider

Queen's University & University of Cambridge OSRP
Cambridge, UK

Abstract

This paper explores the emerging ethical and privacy challenges posed by brain-computer interfaces (BCIs), focusing on mind-reading BCIs that decode neural activity to interpret thoughts and intentions. As BCI technology progresses from medical applications to consumer markets, the stakes for personal privacy and autonomy rise exponentially. This work examines three unique privacy dilemmas, termed the "Impulsivity Problem", the "Judgement Problem", and the "Fingerprint Problem". These issues emphasize that neural data, with its deeply personal and inextricable link to identity and thought, cannot be treated like conventional forms of information. Drawing on philosophical frameworks, particularly Foucauldian concepts of surveillance and biopower, this paper critically analyzes the potential for BCIs to create a new mode of privacy-infringing observation. To address these concerns, the study proposes a value-sensitive design (VSD) framework and provides a roadmap for ethically aligned BCI development.

This dissertation is my own work and includes nothing which is the outcome of work done in collaboration except as specified in the text. It is not substantially the same as any work that has already been submitted before for any degree or other qualification except as specified in the text. It does not exceed the agreed word limit.

Table of Contents

1. INTRODUCTION 4

2. CURRENT AND FUTURE LANDSCAPE 5

2.1 WHAT ARE BCIS? 5

2.2 BCI APPLICATIONS 6

3. BCIS AND THE FUTURE OF PRIVACY 8

3.1 SCIENCE, WITHOUT THE FICTION 8

3.2 THE PHILOSOPHY OF IT ALL 9

The Impulsivity Problem 10

The Judgement Problem 11

The Fingerprint Problem 13

4. DESIGNING VALUE-SENSITIVE BRAIN-COMPUTER INTERFACES 14

5. CONCLUSION 17

WORKS CITED 18

1. Introduction

The trouble with having an open mind, of course, is that people will insist on coming along and trying to put things in it.

— Pratchett, *Diggers*

Brain-computer interfaces (BCIs) are computer-based systems that collect, analyze, and convert brain signals into messages that are then transmitted to an output device so as to carry out an intended action. This technology creates a novel output channel in which brain impulses can interact with or control external devices, bypassing the conventional output channels of peripheral nerves and muscles (Shih et al., 2012). A BCI device identifies the user's mental state by analyzing electrophysiological brain signals and translates them into output commands that fulfil the user's goal. This technology offers a ground-breaking proposal: it attempts to "separate human communication from human muscle and to give thought the power of action", to allow paralyzed patients to control prosthetic limbs with their thoughts alone, to record our memories, and to convert abstract thoughts into speech (Parker, 2003, para. 1; see also Abdulkader et al., 2015; Nabavi, 2014). There is no doubt about it: BCIs are devices of Rube Goldbergian complexity. This, coupled with the fact that they are emerging technologies still largely confined to their developmental stages, means that there has yet to be any widely acknowledged policies or safeguards in place when it comes to their design and development (Rutger et al., 2012). As they become more prevalent, we must examine the hazards they may pose, as well as the ethics and regulations that need to be considered in order to protect the privacy, consent, autonomy, integrity and dignity of its users. A failure to do so could mean that the introduction of BCIs to the consumer market would grant companies and governments unfettered access to the brain activities- and, by extension, the thoughts- of patients and customers. One possible risk of designing BCIs without considering human values is what Martha Nussbaum calls 'the tragic question': when none of our viable options are free from moral wrong-doing (Nussbaum, 2000). Consider, for example, how some technologies (such as cloud computing platforms and wireless internet) have become so pervasive in our lives in ways that make it virtually impossible to opt out. If we are not careful, BCI users could very well run into 'the tragic question' sometime in the near future.

Make no mistake, I am not suggesting that these technologies never be developed in the first place nor am I attempting to become the watchdog of neuroscience. Instead, I hope to propose general guidelines centered around human values that will incorporate ethics into the early design phases of such

innovations, to maximize their benefit and minimize their potential hazards — both on an individual and societal level.

This paper will focus solely on the privacy implications of BCI technology. The scope of this paper will also be limited to a type of BCI known as 'mind-readers'- devices that have the ability to decode abstract thoughts and convert them into transmittable messages (Miller et al., 2016). The aim of this dissertation is to (1) identify the significant privacy-related challenges that may arise with the advent of BCI mind-readers and (2) propose value-sensitive design parameters that can help us to avoid these challenges and ensure that these technologies operate in an ethically favourable or, at the very least, ethically acceptable manner. I shall argue that mind-reading BCI devices present unique privacy challenges and as such, the neural data they collect from an individual cannot be treated like any other form of information (web-browsing activity, genetic makeup, etc.).

The paper is organized as follows: Section 2 describes the current BCI landscape, highlighting several examples of their current uses. It also outlines the benefits and risks of such a technology and provides predictions for what the landscape will look like in the near future. Section 3 touches on the privacy implications associated with mind-reading BCIs and specifically outlines three ethical problems that arise as a result. Section 4 provides a summary of value-sensitive design and presents a roadmap for how value-sensitive design principles can be used to (1) create ethically aligned brain-computer devices and (2) solve the ethical problems presented in Section 3.

2. Current and Future Landscape

2.1 What are BCIs?

BCIs are ever-so diverse and the industry itself is continually evolving. This makes it tricky to find consensus on a single definition of a BCI. Nevertheless, most scholars in the field agree on a few elements that must be met to constitute a brain-computer interface system. These critical elements include (1) the measurement of brain activity [i.e signal acquisition], (2) translation of signals into commands, (3) device output, and (4) real-time or near-real-time feedback to the user on whether they have achieved the desired goal. In other words, a BCI is a brain-implanted device that can directly detect and read neuronal signals and subsequently translate them into messages or executable commands that can be carried out by specific external devices (Shih et al., 2012). BCIs can be classed as non-invasive, invasive, or partially invasive depending on how they acquire their neural signals. Because of its relative safety and feasible

technological components, electroencephalography (EEG)— which collects electrical impulses from the scalp— has been the primary form of recording signals for non-invasive BCIs. Alternatively, invasive BCIs use microelectrodes implanted in the cortical layers of the brain to extract single-neuron recordings. Electrocorticographic (ECoG) recordings extracted from sensors positioned above or below the dura mater of the cortical surface are used to collect signals for partially invasive BCIs (Kawala-Sterniu et al., 2021) (Fig 1.).

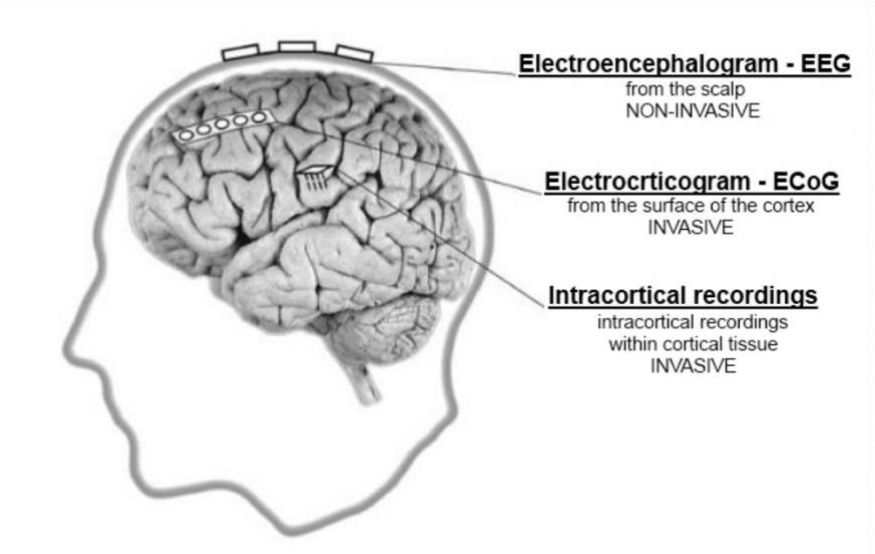

Figure 1. *Three different methods for measuring the electrical activity of the brain* (Kawala-Sterniu et al., 2021).

2.2 BCI Applications

Most early studies of brain-computer interfaces were centered around a concept dubbed 'conscious direct control', or the use of neural impulses to directly alter the state of machines such as wheelchairs, computers, and prosthetic devices (Lance et al., 2012). In 2012, for example, Hochberg et al. published the first peer-reviewed study of tetraplegia patients successfully controlling a robotic arm in three-dimensional space. By the end of the study, participants were able to grab and transport small objects placed in front of them by simply thinking about where they wanted to move their arm (Hochberg et al., 2012). Studies like these were almost entirely limited to the clinical setting and mostly focused on a user's 'intended movement'. In other words, the initial aim of BCI devices was to use electrophysiological signals to predict how a user would want to move their body. More recently, a large portion of BCI research has focused on expanding the concept of 'conscious direct control' and is attempting to develop BCI technologies that can interpret users' thought processes and intended speech.

These devices would not only allow us to decode movement but also higher-order plans and abstract thoughts. Using cortical surface recordings of the brain, researchers have been able to identify and reconstruct the visual mental images that a subject perceives while thinking (Miller et al., 2016). In a 2016 study by Miller et al., seven epilepsy patients were hooked up to ECoG sensors and were told to watch a computer screen as several greyscale images were displayed in a random sequence. These images varied in content and included pictures of houses and faces, as well as blank screens. Each user's ECoG data stream was sent to a powerful computer program that analyzed brain signals in real-time, determining what these brain signals looked like when users were seeing each image. The first two-thirds of the images were used to train the algorithm essentially telling it, 'this is what brain signals look like when someone views a house'. For the remaining one-third, the algorithm was able to predict, with 96% accuracy, what the person actually saw (Miller et al., 2016).

A similar study was conducted by Nishimoto et al. in 2011 in which researchers were able to reconstruct video clips based only on the brain activity of the people who watched them. The figure below presents some of their findings; on the left is the original video clip that was presented to the subjects and on the right is an average of 100 similar clips from YouTube, selected by a computer algorithm that matched these clips to the brain activity of participants while watching the original video (Fig 2). This technology may prove to be particularly useful for patients in a vegetative or minimally conscious state as it can allow them to communicate with the outside world. It may also allow us to records dreams and hallucinations and would allow users to re-experience these states in full consciousness (Nishimoto, et al., 2011).

Figure 2. *Presented clip vs. the clip reconstructed from brain activity* (Nishimoto, et al., 2011).

While brain-computer interfaces were initially created in the context of clinical medicine, a multitude of consumer-grade neurotechnology devices with non-clinical purposes have made their way onto the market in recent years. Emotiv and Neurosky, for example, are two companies that produce a wide range of wireless BCI headphones that can be linked to compatible smartphones and PCs (Ienca & Haselager, 2016). These devices may be used to not only assess a user's brain activity (e.g. concentration levels), but also to operate equipment remotely in order to engage in a variety of activities such as gaming,

communication, and self-monitoring. Moreover, in 2017, Facebook announced that it was in the process of developing a non-invasive, wearable BCI that would enable users to "type with their brains" (Facebook, 2017, para. 19). McCartney and colleagues projected that, in light of these trends, neurodevices will eventually replace the keyboard, touch screen, and voice command device as humans' preferred way to interact with computers (McCartney, 2015).

Rapid advancements in BCI technologies will also prove useful in military operations and training, especially in terms of the improvement and optimization of combat efficiency. For example, DARPA- the Pentagon's research arm- is currently working on a project dubbed 'Silent Talk' which intends to develop battlefield user-to-user communication whereby EEG readings transmit a soldier's 'intended speech'. This would thereby obviate the need for any vocalization or body movements when communicating with another soldier There have already been reportings of successful applications of 'silent speech' during reconnaissance and special operations situations (United States Department of Defense, 2009).

In summary, there have been rapid developments in BCI technology in recent years. If we continue at this rate, it is only a matter of time before mind-reading BCI devices enter the consumer market. With that being said, the next natural question is: are there any ethical concerns that must be addressed? I hope to use my understanding of the current and future state of BCI technology in order to predict some of the issues associated with consumer-grade mind-reading BCI devices.

3. BCIs and the Future of Privacy

3.1 Science, Without the Fiction

When Captain Kirk of *Star Trek* learns of a spy joining one of the groups onboard the spaceship *Enterprise*, he desperately wants to learn more about them and their intentions. To do this, Kirk decides that he will use one of his staff member's telepath abilities to read the minds of all those onboard the ship. However, before Kirk has the opportunity to do this, he is reminded by one of his assistants that, under the law, "the right to mental privacy is an inalienable right of all Federation citizens and shall not be abrogated without due process of law" and that "to find one guilty individual in either of those groups means there is a large probability of invading the privacy of a number of innocent people" (Mitchell, 1990, pp. 52, 150). This scenario, set in the 23rd-century, was once thought to be nothing more than science fiction. This kind of predicament, however, may become a reality much sooner than anticipated.

New scientific discoveries have created new opportunities for understanding the human brain and the recent developments in BCI technology are the first steps towards being able to read the human mind.

However, as Annabelle Lever puts it, “the power to do good is also the power to harm, so scientific advances inevitably foster as many dystopian fears as utopian hopes” (Lever, 2011, p. 1). The nature of brain-reading BCIs lends itself to the fear that an individual’s thoughts, feelings, and deepest secrets may be revealed. While the risk of adverse outcomes is not unique to mind-reading BCIs, Kenneth Foster argues that these outcomes are perhaps more profound with the latter than with other technological innovations, because neurotechnology interfaces are “intimately and fundamentally related to a person’s communication with the outside world” (Foster, 2006, p. 196). In other words, these devices likely to create new ways of hurting people, many of which will involve infringements of privacy. To better understand these unique privacy infringements, I will first identify the distinctive characteristics of BCIs that separate it from other forms of technology. I will then use ethical concepts and principles to evaluate the privacy implications of these characteristics.

3.2 The Philosophy of it All

Neuroscience and philosophy interact on many different levels. This is to be expected. While neuroscience studies the underlying brain processes that guide human behaviour, philosophy is focused on making sense of this behaviour. It can be argued, then, that the two disciplines will always be natural partners. I hope to use philosophy as a tool to help better understand the unprecedented privacy challenges we may face with the advent of consumer-grade mind-reading BCIs.

Now you may think that the data collected by BCIs should be regarded as ‘personal information’ and, as such, there is no reason for this data to be treated as any different from other forms of personal information. In other words, it might be reasonable to posit that this data should fall within the realm of established privacy and data protection policies. The argument here would be that if an individual has a reasonable expectation of privacy regarding the information that can be derived from their Facebook profiles, then surely, they maintain the same expectations regarding the data derived from one’s own mind. However, brain-computer interfaces— specifically those designed to ‘mind read’— are creating new, unprecedented challenges that require their own privacy protection norms. I argue that neural data cannot be treated like other forms of information such as one’s Facebook activity or location. More specifically, I believe that the introduction of these BCIs to the consumer market will result in three unique ethical problems which I call: (1) ‘the impulsivity problem’, (2) ‘the judgement problem’ and (3) ‘the fingerprint problem’. In an attempt to illustrate these dilemmas, consider this hypothetical case scenario:

Jacob is a 33-year-old accountant who was informed by his employer that beginning this month, he would have to undergo attention and productivity training by wearing a new brain-computer interface gadget that delivers neurofeedback. The device allows Jacob to do a multitude of things such as typing on his computer by just thinking of the words he wants to write. This has significantly improved his efficiency. The device is also connected to his phone and sends notifications to alert Jacob that his concentration levels have dropped below a certain point. He feels that his attention has improved since getting the device. From what Jacob understands, this device can somehow measure his brain activity and is able to discern his concentration levels and intended speech. However, Jacob is worried that the device can also read the thoughts that he does not want to publish. Last Monday, Jacob received a lecture from his boss who said that he could tell that Jacob most certainly consumed alcohol on Sunday night and that his neural activity reflected that he was still slightly intoxicated on Monday morning. While listening to the lecture, Jacob impulsively thought about how much he disliked his job. He then wondered if his boss would find out what he was just thinking and if there would be any repercussions.

This scenario (albeit somewhat futuristic) hints at the inherent difference between a company monitoring things like its employees' Facebook pages vs. monitoring their mental states.

(a) The Impulsivity Problem

The first major concern is what I've called *the impulsivity problem* which posits that the inherent impulsivity of an individual's thoughts makes it difficult for that user to predict what information a BCI will collect from them. Unlike all other forms of social media where you have to actively make a decision about what to post and what to exclude, many of our thoughts- like Jacob's complaint about his job- are impulsive in nature.

Barry Gordon, professor of neurology and cognitive science at the Johns Hopkins University School of Medicine, explains that we are cognizant of only a small portion of our brains' thinking and can only control only a mere fraction of our conscious ideas. It becomes even trickier to control thoughts when experiencing strong feelings and emotions (Gordon, 2013). This makes it almost impossible to predict what thoughts a BCI device will decode. This warrants the question: how can an individual consent to the collection of data if they are unable to control what data is being collected? This infringes on an individual's privacy because these devices may ultimately have access to information that the user never consented to sharing.

(b) The Judgement Problem

Up until now, the only way for humans to have shared our thoughts with others has been to take some sort of physical action: to speak, to move, to type out an ill-considered tweet. However, the development of neural decoding brain-computer interfaces may make it possible to breach the privacy of the human mind, and judge others not only on the basis of their actions, but also on the basis of their thoughts and inclinations. This ability to be judged for our thoughts, intentions, and unexecuted actions is a unique consequence of BCIs and is not possible with any other current form of technology (Drew, 2019). I will henceforth refer to this unique phenomenon as *the judgement problem.*

Many BCI companies are currently developing surgically implantable mind-reading BCI microchips designed for the general public in hopes of revolutionizing communication and entertainment (Abdulkader et al., 2015). With these chips, however, many users may feel as if every one of their thoughts is being collected, analyzed, shared, and used to judge them. This phenomenon draws extraordinary parallels with a philosophical metaphor known as Foucault's panopticon. Foucauldian panopticism can thus be utilized as a framework for unveiling brain-computer interfaces as a mode of privacy-encroaching surveillance.

The panopticon is a concept for a prison that was initially designed in 1787 by Jeremy Bentham. The idea behind the design was to allow a single security officer in a central tower to monitor all inmates at an institution without the inmates being aware of when they were being observed. Bentham argued that prisoners who knew that they *could* be observed at any moment would conform to the norms established by the disciplinary authority, regardless of whether they were actually being observed at that given moment. Later, Michel Foucault expanded on Bentham's ideas by using the panopticon as a metaphor for societal surveillance and disciplinary power (Gutting & Oksala, 2021).

Throughout his work, Foucault speaks of a notion he terms *biopower*. This phrase refers to the regulation of human life at the individual level. That is to say that it is "a form of power that targets those of the population" (Rogers et al., 2013, p. 34). In order to achieve this regulation, Foucault argued that governance would shift away from preventing particular behaviours and activities [through means such as violence and force] and towards processes of normalization[1]. This normalization occurs, he argues, through hierarchical observation in which an individual- who knows that they are being watched by a

[1] Normalization, as defined by Foucault, involves (1) the creation of an idealised norm of conduct (i.e the categorization of certain acts as 'normal' or 'abnormal' in the eyes of the state) and then (2) rewarding or penalising individuals for adhering to or departing from this ideal (Foucault, 1990).

higher power- will remain obedient and prevent themselves from acting out. He then goes on to explain that the most 'ideal' disciplinary system would be one in which a single gaze could supervise the entire population (Foucault, Discipline and Punish, 1995).

At its core, the panopticon requires two things: unequal surveillance and unequal power. While the guards may be able to see any or all of the inmates, the inmates are unable to see the guards. And while the guards may be able to punish/discipline any or all the inmates, the inmates are unable to punish the guards. This is an appropriate metaphor for brain-computer interfaces in the sense that (1) BCI users cannot see the data collected from them by the BCI company, (2) BCI users cannot see the data collected from them by the state (e.g. the National Security Agency (NSA) in the United States) and finally, (3) users may be unaware of the information that can be inspected by acquaintances/contacts. Those who have access to this information would be able to exert unequal power on the users. This information- like most other forms of information that governments/companies have on a user- may be sold to insurance companies, used to manipulate elections, or be used as blackmail in the instance of a hack (Véliz, 2020). The difference, however, is that the information acquired by brain-computer interfaces is inherently more private and intimate than that found on an individual's Facebook page or web-browsing history (Drew, 2019). With implantable mind-reading BCI chips, users may continually fear that others have access to their thoughts, desires, or unexecuted intentions and, ultimately, that they will be judged or punished for these thoughts. Gavison argues that the fear of others constantly listening to our inner narrative can ultimately force us to change the way we think:

"In such a state, there would be no private thoughts, ... we would probably try hard to suppress our daydreams and fantasies once others had access to them. We would try to erase from our minds everything we would not be willing to publish, and we would try not to do anything that would make us likely to be feared, ridiculed, or harmed. There is a terrible flatness in the person who could succeed in these attempts" (Gavison, 1980, p. 443).

Much like a panopticon can be used to control its prisoners, BCIs could be used to control its users. If this is the case, then BCIs can be considered a form of indirect mind control, one in which a user must regulate their thoughts so that they might be better received by both the government body (e.g. NSA) and by their social networking peers. The moral value of privacy is thus intricately intertwined with that of the moral value of autonomy (DeCew, 2018). In other words, autonomy is critical when it comes to personal privacy because only when an individual has a certain degree of privacy can they arrive at their own evaluations, intentions, beliefs, and ultimately, decisions. With the advent of BCI mind-reading implants, there may always be a lingering fear that there are others listening in on your thoughts. This not only

infringes on one's privacy, but also their autonomy. It means that individuals like Jacob- who aren't necessarily sure what exactly their device can do- no longer have freedom of thought and will have to constantly monitor their inner narratives for fear of repercussions like losing their job.

(c) The Fingerprint Problem

The privacy issues associated with BCIs become compounded when we consider that an individual's brain activity (and by extension, their thoughts and feelings) are biometric identifiers in-and-of-themselves (Dominguez et al., 2014; Palanippan et al., 2017). In other words, there is no anonymity when it comes to your brain data and all neural signals can be traced directly back to you. This ability to link an individual's neural signals to their identity poses a pretty significant problem for privacy.

BCI information can be traced back to you in two ways: (1) through a biometric identification framework (where neural signals are treated similar to a fingerprint or DNA sample) or (2) by linking a user's BCI device (and any subsequent information it collects) to the user's social media profile. Both of these methods are already in development. An EEG-based biometric framework for automated identification verification was developed by Palanippan and Mandic in 2007 (Palaniappan & Mandic, 2007). Facebook has also proposed a brain–computer interface system that would not only collect and analyze brain signals, but also link the data generated from them with comprehensive accounts of the user's social media activities, allowing for seamless user interactions with their systems. This would allow Facebook to form "rich links between overt actions and hitherto hidden brain activity" and will open up new dimensions of understanding human behaviour (Rainey et al., 2020, p. 2303). But, like the Facebook and Cambridge Analytica micro-targeting scandals, this data will be used for personal gain and may potentially cause social and political harm in the process. By allowing corporations to have access to their users' cognitive activities, we may be opening the floodgates of a new movement in data-driven marketing and campaigning, one that enables new, more nefarious, and perhaps harder to deflect, methods of manipulation (Rainey et al., 2020).

BCIs can therefore be considered both a biometric [i.e it can identify a user] and a content-generating technology [i.e it can acquire information on the user's behaviour]. This means that BCIs cannot be treated as if it were simply one or the other; they cannot be treated like other biometrics (such a fingerprint or DNA) nor can they be treated like any other form of communication/online sharing. Other biometrics do not reveal nearly as much information as an individual's neural activity. A fingerprint, for example, would only be able to help you identify an individual and provides no other information about them. A DNA sample, while a little more content-rich, would only be able to provide you with that

individual's ancestry, ethnicity, or their genetic risk for certain health conditions. An individual's decoded neural activity, on the other hand, may reveal all the thoughts, intentions, beliefs, dreams, and memories of a user. In this respect, it draws many parallels to the information that can be acquired from an individual's social media activity. However, BCIs can also be distinguished from all other forms of communication/online sharing. This is because with all other forms of social networking/internet browsing, it has been possible to anonymize the information you share. Accounts can be created under fake names, VPNs can be used to mask one's location, and users can choose whether or not to share content that may lead to their identification. With BCIs, however, this is not the case. Neural data contains *both* the content that the user wants to share and the user's identity. It is impossible to separate the two. As such, we must be even more careful with the sharing and protection of this data and to do this, I believe that we must implement some recommendations that are unique to BCI technology.

4. Designing Value-Sensitive Brain-Computer Interfaces

Although scholars have been able to identify the potential privacy issues associated with brain- computer interface technologies, the practical solutions for preserving privacy have yet to developed. To address these issues, I believe that we need to use a framework known as value-sensitive design (VSD). VSD is an established method of designing technology that considers human values and does so in a systematic and holistic way (Friedman et al., 2020).

VSD has been applied to a number of technologies, including energy systems, mobile phone usage, architecture projects, and augmented reality systems, to name a few (Longo et al., 2020). However, there has yet to be any discussion on how VSD might be applied to tackle the unique privacy challenges of BCI mind-reading technologies. I propose that in order to avoid (or at the very least, minimize) the privacy challenges associated with BCIs, we need to introduce elements in the design process that can steer these devices away from being function-centered and towards being human-centered. We must figure out how to make sure that BCI technology aligns to principles that we (as a society) respect, embrace, and prioritize. We especially need to consider how we can support and propagate values like privacy.

To do this, I have laid out a few BCI design elements that must be taken into consideration throughout the design process. These design elements will attempt to address the three aforementioned ethical problems and are proposed under the assumption that they will be government-regulated policies used to keep corporations/researchers in-check.

(1) Ensuring On/Off User Controls

BCI users should be able to control when their devices are on or off. Ideally, this deactivation switch should be located on the BCI system's sensor. This ensures that the disabling of the device occurs as early as possible in the collection/translation process, such that the device cannot even collect the data if it is turned off. If the sensor were to consistently gather data from the user, it will result in increased susceptibility or unauthorized data collection.

(2) Providing Visual Cues that Clearly Indicated when the Device is Recording

BCI users should be able to quickly recognize when their device is on. This upholds fundamental principles of informed consent that state that users should understand when their data is being collected. Many technology companies have incorporated noticeable visual signals into their devices. Most laptops, for example, have a built-in indicator light that is used to notify users when their cameras are on. BCI devices should be no exception to this rule.

(3) Encouraging Non-Invasive Designs

If possible, BCI users should be able to easily remove their device. By physically removing the device, the user can guarantee that none of their data is being collected or intercepted without their consent. With invasive or semi-invasive implants, there is always a chance that biohackers (or even the manufacturers) may non-consensually access brain data, regardless of whether or not the user had turned the device off. However, it should be acknowledged, that certain BCIs may require more advanced/invasive techniques in order to successfully carry out this function. In cases such as these, this 'non-invasive' rule would be overridden. Elements (1), (2), and (3) attempt to solve 'the judgement problem' by allowing users the freedom to ensure that the only information being collected by the device is information that they are ok with being shared or judged for. These three recommendations also address 'the impulsivity problem' because they allow users to be more cognizant of when their device is recording. These users can thereby act accordingly to turn the device off (or take it off entirely) during moments of high stress/emotion when they are less likely to be in full control of their thoughts (heated arguments, stressful meeting with an employer etc.).

(4) Implementing BCI-Decoder-Filter To Enshrine Purpose Limitation

BCI devices should only decode the information that is absolutely necessary for their function. The heart of BCI technologies is the ability to process raw brain data in order to extract meaningful information. However, not every device needs to- nor should be able to- extract all types of information. *Muse*, for example, is a company that produces EEG-reading headbands that measure concentration levels. It,

therefore, has no need for extracting information about an individual's movements or emotions (Przegalinska et al., 2020). To ensure that BCI devices only collect the information needed for them to function, I suggest the implementation of a 'decoding-filter'. This filter would mean that the BCI system could only decode signals needed for specific functions. The proper technical term for these filters is 'channel selection algorithms'. EEG data, for example, is often collected from over 100 areas of the brain and channel selection algorithms are used to identify optimal EEG channels for a given function. The suggestion, then, is to only decode and transmit the information that passes through these optimal channels, as deduced by the algorithm (Alotaiby et al., 2015). In the future, these channel selection algorithms may become so advanced that they would be able to recognize impulsive thoughts and prevent them from being decoded. This would solve 'the impulsivity problem'. If we do not end up developing filters that are capable of solving the impulsivity problem, then recommendations (1), (2), and (3) would suffice in terms of lessening the problem.

(5) Minimizing Encryption Risks

BCIs must meet a certain standard of security. The majority of existing BCI applications have unfettered access to users' raw EEG data and employ minimal security measures. In order to protect users' neural data, we can develop strategies based on modern cryptographic techniques. Researchers are already looking at the capabilities of homomorphic and functional encryption as methods by which to protect user privacy (Takabi, 2016). By ensuring proper encryption protocols are in place, we minimize the risk of brain hacking. This becomes even more important when you consider the fact that this raw data cannot be anonymized and is essentially a biomarker (similar to that of a fingerprint).

To further minimize the risk of hacking, the BCI system should ideally be a 'closed' one (Takabi, 2016). A closed system is one where a user's information does not leave the device (i.e it does not get sent to a database or get stored in the cloud). This, however, is not always possible as most BCI devices aim to transmit information to (1) external devices such as a user's smartphone or (2) another individual's device (i.e collaborative games or user-to-user BCI communication). In scenarios like these, the recommendation would be to only transmit data that has already been translated into functional information. This means that a user's brain signals would never have to leave the user's own BCI device. Instead, what gets transmitted is non-identifiable information such as words and instructions. This thereby reduces the chance that a user's brain signals will get intercepted or sold to third-party organizations. By ensuring that these brain signals never leave the user, it helps minimize the implications of 'the fingerprint problem' by ensuring that the transmitted data can be anonymized and will not easily be tracked back to the user.

(6) Providing Receiver-contextualized Explanations and Transparent Purposes

Any system's goals must be clearly stated. That is, a system's processes should be fully understood by the user. BCI technologies are already being used in an increasingly widespread manner. The potential ubiquity of this technology has increased the need for explainability and transparency in their operations and goals. BCI systems should be explained in a way such that their operations and objectives are receiver-contextualized so as to prevent any potential harms that may result from opaque objectives and operations. This would help to resolve 'the judgement problem' by attempting to level the 'unequal surveillance' that may exist between BCI companies/governments and users. BCI users understanding how and when they are being surveilled is akin to telling prisoners of the panopticon how and when they are being watched. Ultimately, this knowledge will result in (1) less fear and paranoia in BCI users about the data that is being collected about them and (2) grants them greater control over the information they share.

These six aforementioned recommendations can be used like a checklist by BCI designers/manufacturers to ensure that they are doing all that they can to address the three problems outlined in Section 3. As such, these recommendations can and should be considered a set of jointly sufficient conditions. This is to say that while there might be other ways to approach these challenges, these conditions (if all met to the best of the designers' capabilities) can act as an industry standard that (1) prove that these designers are doing all that they can to address 'the impulsivity problem', 'the judgement problem', and 'the fingerprint problem' and (2) ensures that the BCIs they create will be- at the very least- ethically acceptable with respect to privacy.

5. Conclusion

The widespread availability of low-cost, scalable, and simple-to-use neuroapplications has the potential to open up previously unimagined possibilities and make neurotechnology deeply ingrained in our daily lives. The ethical implications of this technology, especially those pertaining to privacy, remain largely unexplored. I argue that when it comes to mind-reading BCI devices, there are three ethical problems that may arise: (1) 'the impulsivity problem', (2) 'the judgement problem' and (3) 'the fingerprint problem'. The major takeaway from these three dilemmas is that neural data collected from BCIs cannot and should not be treated like other, more conventional, forms of informational data. I propose that, in view of the disruptive shift that neurotechnology is causing in the digital ecosystem, we need to consider human-centered design elements that will ensure that we are doing all that we can to build ethically acceptable BCI devices.

Works Cited

Abdulkader, S., Atia, A., & Mostafa, M.-S. (2015). Brain computer interfacing: Applications and challenges. *Egyptian Informatics Journal*, 213-230.

Alotaiby, T., El-Samie, F., Alshebeili, S., & Ahman, I. (2015). A review of channel selection algorithms for EEG signal processing. *Journal on Advances in Signal Processing*, 66.

DeCew, J. (2018, March 21). *Privacy*. Retrieved from Stanford Encyclopedia of Philosophy: https://plato.stanford.edu/entries/privacy/

Dominguez, O., Mills, B., & Carpenter, K. (2014). Connectotyping: Model Based Fingerprinting of the Functional Connectome. *PLOS ONE*.

Drew, L. (2019, July 24). *The ethics of brain–computer interfaces*. Retrieved from Nature: https://www.nature.com/articles/d41586-019-02214-2

Facebook. (2017, April 19). *F8 2017: AI, Building 8 and More Technology Updates From Day Two*. Retrieved from Facebook: https://about.fb.com/news/2017/04/f8-2017-day-2/

Foster, K. (2006). Engineering the Brain. In J. Illes, *Neuroethics: Defining the Issues in Theory, Practice and Policy* (pp. 185-200). Oxford: Oxford University Press.

Foucault, M. (1990). *The History of Sexuality, Volume I: An Introduction.* New York: Vintage .

Foucault, M. (1995). *Discipline and Punish.* New York: Vintage Books.

Friedman, B., Kahn, P., & Borning, A. (2020). Value Sensitive Design and Information Systems. In P. Zhang, *Human-Computer Interaction in Management Information Systems: Foundations* (pp. 1-27). New York City: M.E Sharpe.

Gavison, R. (1980). Privacy and the Limits of Law. *The Yale Law Journal*, 421-471.

Gordon, B. (2013, March 1). *Can We Control Our Thoughts? Why Do Thoughts Pop into My Head as I'm Trying to Fall Asleep?* Retrieved from Scientific American: https://www.scientificamerican.com/article/can-we-control-our-thoughts/

Greely, H. (2009). Law and the Revolution in Neuroscience: An Early Look at the Field. *Akron Law Review*, 687-716.

Gutting, G., & Oksala, J. (2021, June 21). *Michel Foucault*. Retrieved from The Stanford Encyclopedia of Philosophy: https://plato.stanford.edu/entries/foucault/

Hochberg, L., Bacher, D., Jarosiewicz, B., & Masse, N. (2012). each and grasp by people with tetraplegia using a neurally controlled robotic arm. *Nature*, 372-375.

Ienca, M., & Haselager, P. (2016). Hacking the brain: brain–computer interfacing technology and the ethics of neurosecurity. *Ethics and Information Technology*, 117-129.

Kawala-Sterniu, A., Browarska, N., & Al-Bakri, A. (2021). Summary of over Fifty Years with Brain-Computer Interfaces—A Review. *Brain Sciences* , 11-43.

Kozel, A., Johnson, K., Mu, Q., Grenesko, E., Laken, S., & George, M. (2005). Detecting Deception Using Functional Magnetic Resonance Imaging. *Biological Psychiatry*, 605-613.

Lance, B., Kerick, S., Ries, A., Oie, K., & McDowell, K. (2012). Brain–Computer Interface Technologies in the Coming Decades. *Proceedings of the IEEE*, 1585-1599.

Lever, A. (2011). Neuroscience v. Privacy? A Democratic Perspective. In S. Edwards, S. Richmond, & G. Rees, *I Know What You Are Thinking: Brain Imaging and Mental Privacy.* Oxford: Oxford University Press.

Longo, F., Padovano, A., & Umbrello, S. (2020). Value-Oriented and Ethical Technology Engineering in Industry 5.0: A Human-Centric Perspective for the Design of the Factory of the Future. *Applied Sciences*, 2-25.

McCartney, R., Yuan, J., & Bischof, H.-P. (2015). Gesture Recognition with the Leap Motion Controller. *Rochester Institute of Technology Scholar Works*.

Miller, K., Schalk, G., Hermes, D., Ojemann, J., & Rao, R. (2016). Spontaneous Decoding of the Timing and Content of Human Object Perception from Cortical Surface Recordings Reveals Complementary Information in the Event-Related Potential and Broadband Spectral Change. *PLOS Computational Biology*.

Mitchell, V. (1990). *Enemy Unseen.* New York City: Simon and Schuster.

Nabavi, S., Fox, R., & Proulx, C. (2014). Engineering a memory with LTD and LTP. *Nature*, 348-352.

Nishimoto, S., Vu, a., Naselaris, T., Benjamini, Y., Yu, B., & Gallant, J. (2011). Reconstructing visual experiences from brain activity evoked by natural movies. *Current Biology*, 1641-1646.

Nussbaum, M. C. (2000). The Costs of Tragedy: Some Moral Limits of Cost-Benefit Analysis. *The Journal of Legal Studies*, 1005-1036.

Palaniappan, R., & Mandic, D. (2007). EEG-based biometric framework for automated identification verification was developed by Palanippan et al. in 2007. *The Journal of VLSI Signal Processing Systems for Signal, Image, and Video Technology volume*, 243-250.

Parker, I. (2003, January 12). *Reading Minds*. Retrieved from The New Yorker: https://www.newyorker.com/magazine/2003/01/20/reading-minds

Przegalinska, A., Ciechanowski, L., Magnuski, M., & Gloor, P. (2018). Muse Headband: Measuring Tool or a Collaborative Gadget? *Collaborative Innovation Networks*, 93-101.

Rainey, S., Martin, S., Christen, A., Megevand, P., & Fourneret, E. (2020). Brain Recording, Mind-Reading, and Neurotechnology: Ethical Issues from Consumer Devices to Brain-Based Speech Decoding. *Science and Engineering Ethics volume*, 2295-2311.

Rogers, A., Castree, N., & Kitchin, R. (2013). Biopolitics. In N. C. Alisdair Rogers, *A Dictionary of Human Geography* (p. 34). Oxford: Oxford University Press.

Rutger, V., Steines, D., & Szibbo, D. (2012). Ethical Issues in Brain–Computer Interface Research, Development, and Dissemination. *Journal of Neurologic Physical Therapy*, 94-99.

Shih, J., Krusienski, D., & Wolpaw, J. (2012). Brain-Computer Interfaces in Medicine. *Mayo Clinic Proceedings* , 268–279.

Takabi, H. (2016). Firewall for Brain: Towards a Privacy Preserving Ecosystem for BCI Applications. *IEEE Conference on Communications and Network Security (CNS)*, 1-2.

United States Department of Defense. (2009). *Department of Defense Fiscal Year (FY) 2010 Budget Estimates.* Washington D.C: United States Department of Defense.

Véliz, C. (2020). *Privacy is Power.* London: Penguin (Bantam Press).