Privacy and ethics in brain-computer interface research

Eran Klein, Department of Neurology, Oregon Health and Science University

Alan Rubel, Information School and Center for Law, Society and Justice, University of Wisconsin-Madison

Abstract:

Neural engineers and clinicians are starting to translate advances in electrodes, neural computation, and signal processing into clinically useful devices to allow control of wheelchairs, spellers, prostheses, and other devices. In the process, large amounts of brain data are being generated from research participants, including intracortical, subdural and extracranial sources. Brain data is a vital resource for BCI research but concerns have been raised about whether the collection and use of this data generates risk to privacy. Further, the nature of BCI research involves understanding and making inferences about device users' mental states, thoughts, and intentions. This, too, raises privacy concerns by providing otherwise unavailable direct or privileged access to individuals mental lives. And BCI-controlled prostheses may change the way in which clinical care is provided and the type of physical access caregivers have to patients. This, too, has important privacy implications for patients and caregivers. Our goal in this chapter is to examine several of these privacy concerns in light of prominent views of the nature and value of privacy. We argue that increased scrutiny needs to be paid to privacy concerns arising from Big Data and decoding of mental states, but that BCI research may also provide opportunity for individuals to enhance their privacy.

1.   Introduction: Privacy and BCI

Brain-computer interface technology has advanced considerably in the last several decades, presenting substantial opportunities for treatments (Shih, Krusienski, and Wolpaw 2012; Soekadar, Cohen, Leonardo G., and Birbaumer 2014). Research in BCI is leading to unprecedented accumulation of brain data, the need to better understand the neural structures and mental states of research participants and patients, and potential changes in how BCI users, caregivers, and researchers interact. Each of these carries important privacy considerations.

BCI can be defined as "a computer-based system that acquires brain signals, analyzes them, and translates them into commands that are relayed to an output device to carry out a desired action" (Shih, Krusienski, and Wolpaw 2012). And this volume is a testament to the fact that research and implementation of BCI is both promising and growing. But despite the fact that BCI research is advancing steadily, there is still relatively little scholarship dedicated to privacy considerations. Among those who have discussed privacy and BCI (see, e.g., Jebari 2013; Finn, Wright, and Friedewald 2013), no consensus approach has emerged. More importantly, though, commentary on BCI and privacy tends not to address *why* privacy in the context of BCI is (or is not) morally important. In a chapter devoted to privacy issues in new and emerging technologies, Finn et al describe a number of facets of privacy that BCI may affect. They "carry the potential to impact upon privacy of the person, privacy of behavior and action, privacy of communication, privacy of data and image and privacy of thoughts and feelings." For example, privacy of behavior and action may be diminished if BCI information is used to predict behaviors or as a means of gaining marketing influence. Finn et al. also suggest that communications privacy may be affected if data are vulnerable to interception. As plausible as such concerns are, there is a further question as to *why* they matter. Part of our task here is to consider not just how privacy may be affected by BCI and BCI research, but what kinds of privacy effects are morally considerable, and why.

There are several potential reasons why privacy in BCI has not been addressed more extensively, each based on some questionable assumptions. First, there may be an implicit assumption on the part of

researchers, clinicians, and policy-makers that loss of privacy or heightened risk to privacy is a worthwhile or inevitable tradeoff of developing devices with enormous potential health benefits. The promise of a BCI device that would allow someone locked-in to speak or someone who is tetraplegic to reanimate a limb or control a robotic limb is compelling and might seem to make privacy concerns trivial in comparison. Second, more BCI data is collected than is currently interpretable (Finn, Wright, and Friedewald 2013). Until more is known about what's in BCI data, concerns with privacy can seem premature (Hallinan et al. 2014). Third, implantable and non-implantable BCI research with human research participants takes place within academic institutions that have strict data protection policies and informed consent regulation that requires seeking prior consent or authorization of participants; this may lead some to believe that privacy risks are already in some sense adequately addressed.

Each of these reasons for dismissing or minimizing the importance of privacy considerations rest on questionable assumptions. (1) Recent work exploring the perspectives of end users and potential end users of BCI has shown that much is still unknown about what constitutes acceptable tradeoffs in BCI (Blabe et al. 2015; Blain-Moraes et al. 2012; Lahr et al. 2015; Huggins, Wren, and Gruis 2011; Collinger et al. 2013) (2) Waiting until better data analysis methods are available before taking privacy concerns seriously may hinder the goal of developing anticipatory guidance for privacy norms and regulations. (3) While policies and practices of informed consent are already in place in BCI research, the extent to which research informed consent achieves meaningful consent has been challenged, and BCI research may need to include atypical risks (like identity, agency, and stigma) that are not a standard part of the informed consent process (Klein 2015).

From here, the paper is organized as follows. We begin, in section 2, by outlining several key conceptions of privacy relevant to BCI. We then turn to key moral considerations with respect to privacy. In section 3 we address issues of mental privacy. In section 4 we consider opportunities that BCI presents for carving out important domains of privacy for users. And in section 5 we consider implications of large amounts of

data collected in BCI research. Finally, we conclude by offering some preliminary recommendations for next steps in addressing the kinds of privacy concerns identified.

2.   BCI and definitions of privacy

There is a substantial literature on privacy across a wide range of scholarly disciplines, and just how to define privacy remains controversial. Nonetheless, a key step in addressing privacy concerns is to clarify the concept of privacy itself. Here, we describe several important facets of privacy that reflect important threads in privacy scholarship. To begin, it is useful to think of privacy as involving three parts: (1) some person who has, or lacks, privacy, (2) some form or domain of privacy, and (3) some other person or persons who limit the first person's privacy (Rubel and Biava 2014; Blaauw 2013; Matheson 2007). Among the forms of privacy scholars have identified are informational (what others can learn about a person), physical (the degree to which others can physically access one's person), associational (whether others can control the people with whom one associates), and decisional (whether others can limit the range of important decisions one can make for oneself) (J. W. DeCew 1997; Allen 1988; Nissenbaum 2010). It is useful for our purposes to focus on three of these forms. Note that these forms of privacy do not entail anything about the conditions under which privacy is morally important (or indeed whether it is morally valuable at all). We will address those questions below.

Table 1

Types of privacy	Definition	BCI examples
Physical	The condition in which others' access to one's person (by sight, sound, touch, presence) is limited.	Ability to attend some activities of daily living with less intensive intervention by others. Physical access to skull. Being pulled aside in security screening.
Informational	The condition in which others' ability to learn about one, or to make inferences about one, is limited.	Potential for neural recording to expose thoughts, dispositions, and intentions. Unknown inferences from troves of data currently stored.

Draft: Please cite to published version in Nam, Nijholt, and Lotte (eds) Brain–Computer Interfaces Handbook: Technological and Theoretical Advances (CRC Press: 2018).

Decisional	The ability of a person to make important, intimate decisions without excessive influence or control by others.	Storage of intimate BCI conversations. Being prohibited from entering studies based on exclusion criteria regarding (e.g.) reproductive decisions.

Much privacy discussion centers on informational privacy, or the extent to which others may learn about, access information about, or make inferences regarding a person. Informational privacy is protected in research and clinical care by physician-patient confidentiality, health information privacy laws (such as HIPAA (1996) and GINA (2008) in the U.S.), general privacy laws (e.g., EU Directive 95/46/EC (1995)), human subject protections in research, and data security laws. Here it is worth noting that because BCI research often involves international collaboration, it is likely to implicate (inter alia) U.S. sectoral privacy laws and EU general data protection laws.

The extent to which BCI-derived data can be used to generate health or other personal information about a user is only beginning to be understood. For example: Can BCI reveal incidental findings of clinical significance (such as a proclivity to seizures or early tumor) or correlate with disease, such as Alzheimer's disease (Soekadar, Cohen, Leonardo G., and Birbaumer 2014)? Can neural activity patterns be used to detect attention or motivation, which are critical to the success of BCI training (Curran and Stokes 2003) and also may say something about the underlying personality of the person using the BCI? Already, BCI data is collected alongside, and could in principle be correlated with, actions or circumstances with implications for privacy, such as a desire to use the bathroom, discussions with loved ones, or outbursts of emotion. The extent to which associated BCI data could be used to infer mental states or personality characteristics is uncertain, but raises legitimate concern.

Informational privacy is distinct from (though related to) physical privacy. Physical privacy is the condition of having one's body or personal spaces protected from intrusion by others. So, for example, one may desire a degree of physical solitude or to remain free of video surveillance, regardless of whether one is concerned about information gathering. Being subjected to physical examination is an imposition

Draft: Please cite to published version in Nam, Nijholt, and Lotte (eds) Brain–Computer Interfaces Handbook: Technological and Theoretical Advances (CRC Press: 2018).

on one's physical privacy, independently of the degree to which new information is gained by the person conducting the examination.

Participation in BCI research involves the loss of some physical privacy. Placement of EEG electrodes onto the skull involves a loss of physical privacy in this regard, but the loss of physical privacy is more obviously the case with implantable BCI research where neurosurgeons must temporarily traverse or remove part of the skull in order to place electrodes into, atop, or in close proximity to the brain. This is an intrusion, albeit voluntary, of the body and hence involves a loss of physical privacy. Even after surgery, physical privacy remains an issue in so far as an implanted device facilitates easier physical access to a person's body. Participants in the Braingate trial, for instance, have a pedestal attached to their skull that serves as an access port for attaching a data recording cable during BCI experiments (Hochberg et al. 2012). Facilitated physical access is not just a feature of the lab, though, but of the presence of a device and the bodily residuals of having been implanted with it. For instance, individuals with deep brain stimulators are subject to special screening at airports. Being pulled aside from standard metal detectors or other detection devices and showing one's cranial scar or subcutaneous bulge of a battery placed under the chest wall as evidence of a pmed (personal medical electronic device) are forms of physical accessibility to which others are not subject.[1] Individuals with implantable BCI devices are likely to face similar kinds of physical privacy loss.

Decisional privacy concerns whether others can limit the range of important decisions one can make for oneself. In the United States, decisional privacy is often discussed in the context of access to birth control (*Griswald v. Connecticut*, 381 U.S. 479 (1965)), abortion (*Roe v. Wade*, 410 U.S. 113 (1973)), legal restrictions on same-sex partners (*Lawrence v. Texas*, 539 U.S. 558 (2003)), and "conscience clauses" (whether health care professionals must provide services to which they have moral objections). Some

---

[1] In the U.S., the Transportation Security Administration (TSA) issues notification cards which can be shown to screeners, but this does not exempt individuals from additional screening.

Draft: Please cite to published version in Nam, Nijholt, and Lotte (eds) Brain–Computer Interfaces Handbook: Technological and Theoretical Advances (CRC Press: 2018).

commentators object that decisional privacy is better understood as individual autonomy, but others note that legal limits on birth control, abortion, sexual partners, and actions of conscience requires surveillance and information-gathering practices, and is hence deeply entwined with informational and physical privacy (DeCew 2016).

Participation in BCI research can require that volunteers engage in or forego certain activities or make certain decisions. Most obviously this involves adhering to current medical therapy, such as taking one's medications as prescribed, in order that no additional risks of research participation are incurred. But the kinds of decisions required for research participation can also straddle the medical versus lifestyle divide. For instance, in her autobiography *On My Feet Again*, Jennifer French describes the difficult decision to forego having biological children in order to participate in a trial of an implanted neural device (functional electrical stimulation (FES)) (French 2012).[2] Establishing similar inclusion criteria in BCI research may be reasonable given lack of pregnancy safety data, but does represent a significant restriction on an important life choice. Of course, what constitutes an important life choice varies by individual. In a recent study of deep brain stimulation (DBS) for depression or obsessive compulsive disorder (OCD) conducted by one us (Klein), one research participant lamented that he had been told to avoid roller coasters so long as he had an implanted device. It is as yet unclear which life choices individuals with BCI devices will be subject to and whether BCI devices will include technology for monitoring BCI performance that will provide concomitant surveillance data on adherence to restrictions (e.g., accelerometer embedded within a BCI device that identifies a roller coaster versus an automobile ride). Hence, it seems clear that BCI research will implicate questions of informational and decisional privacy.

---

[2] The FES device was not a BCI, but similar rationale for study exclusion could be applied to BCI research.

Draft: Please cite to published version in Nam, Nijholt, and Lotte (eds) Brain–Computer Interfaces Handbook: Technological and Theoretical Advances (CRC Press: 2018).

3.   Mental Privacy

The prospect of a device capable of reading the contents of minds has long generated popular appeal and philosophical interest (Levy 2007). While far from that imagined in science fiction movies and philosophical thought experiments, BCI shows that it is possible to use thought to control devices. As early as the 1960s, research participants were able to learn to send Morse code by modulating alpha oscillations measured by EEG (Dewan 1967). More recent work in BCI demonstrates the possibility of using information gathered from scalp or surgically implanted electrodes to control prostheses, wheelchairs, and computer spellers, among other devices. The ability to decode intentions is currently limited but improving in accuracy, and the range of potential applications is expanding.

Advances in neuroimaging in recent decades helped pave the way for BCI work on decoding brain states. While computed tomography and magnetic resonance imaging have provided increasingly detailed diagnostic information about structural features of the brain, the development of functional MRI (fMRI) is a critical advance in understanding mental activity. fMRI allows for measurement of brain activity (indirectly through measures of blood flow) during certain mental processes or in conjunction with the experience of certain mental states. This has allowed researchers to determine, for instance, which visual image someone is viewing (even being able to reconstruct the image (Schoenmakers et al. 2013) or what implicit attitudes correlate with moral decision-making (Greene et al. 2001). It has also allowed for measurement of specific intentions (Haynes et al. 2007).

Intentions are of particular interest in BCI insofar as they can be used to control an output device. Measurement electrodes placed on the scalp (electroencephalography (EEG)), on the cortex (electrocorticography (Ecog)), in the cortex (penetrating cortical electrode arrays), or in deep regions of the brain (deep brain electrodes) can be used to gather data with high temporal resolution on neural processes associated with preparing for, reflecting on, or acting on intentions. How complex intentions or goals are represented in neural networks is an active area of research. The central idea is that pattern-

recognition algorithms can be used to decipher what neural activity underlies a particular intention. In other words, BCI can be used to determine what a person is thinking (or planning) at a given moment.

BCI offers the prospect of accessing other minds, in some sense, *directly*. We typically ascribe mental states – perceptions, beliefs, memories, attitudes, emotions – based on observations made at a remove, such as what people tell us they are thinking and what we observe in their body language or behavior. And we make further ascriptions of personality traits or characteristics such as sociability, honesty, procrastination, suspicion of authority, and so on. But the prospect of decoding mental states opens a new window on mental life. The animating, even if overly simplistic, metaphor is that the development of radiographs let us look past the skin to the bones and tissues beneath, and now advances in BCI will allow us to peer not just into the brain, but into the mental life buzzing within it (or less colloquially, "supervening upon it").

The prospect of decoding intentions raises concerns about informational privacy. Some of these informational privacy concerns resonate with broader concerns about "mental privacy" arising from work in neuroimaging (Farah and Wolpe 2004; Ryberg 2016).[3] There are related terms for this, including "neural privacy", "brain privacy", "cognitive privacy", "thought privacy", and "cognitive liberty" (Trimper, Wolpe, and Rommelfanger 2014; Illes and Racine 2005; Schneider, Fins, and Wolpaw 2012; Räikkä 2010). Richmond distinguishes between two uses of "mental privacy" in the literature – an older, philosophical sense in which "mental states are (descriptively) private, accessed by the subject in a distinctive first-personal way" and one more involved with civil rights debates concerned with "invasions of the (normatively) private area of the mind" (Richmond 2012, p. 186, n. 2). It is this latter use, often underwritten by a folk psychological notion that the mind contains inner monologues, judgments, eidetic memories capable of being exposed for all to see, that has been taken up in debates about wholesale mind

---

[3] It should be noted that, at present, the ability to discern mental states using fMRI is significantly more advanced than that of BCI. It is too early to say whether this difference in capabilities is a contingent or inherent feature of these two technologies. Regardless, the use of fMRI to infer mental states provides an instructive example for thinking through considerations of privacy in BCI research.

reading or brainotyping of personality (Illes and Racine 2005). Two examples of debates related to the "invasion" sense of mental privacy are lie detection (which is a matter of informational privacy) and neuromarketing (which involves use of information to influence individuals' decision-making, and hence affects decisional privacy).

Advances in neuroscience will have effects on the law.[4] One area that has attracted significant attention is the use of neuroimaging and neural recording to assist or replace testimony (Wolpe, Foster, and Langleben 2005). Memory and eye witness testimony are fallible forms of forensic evidence subject to numerous kinds of psychological biases and errors (for substantial review, see Lacy and Stark 2013). Direct neural recording through EEG (or implantable electrodes if available) may provide an alternative source of forensic information. For instance, if neural patterns are discovered that correlate with deception or feeling of guilt or recognition of people, places, or events, these patterns could be introduced as evidence. There has been significant legal and ethical resistance to using neuroscience in this way (Greely and Illes 2007). Even if not admissible in court, correlations between neural data and mental states could be used in military or anti-terrorism contexts (Tennison and Moreno 2012; Finn, Wright, and Friedewald 2013), employee surveillance (Jebari 2013), or social interactions (e.g., verifying veracity of personal information posted on a dating site).

The application of neuroscience to improve products and advertising has been termed neuromarketing (Ariely and Berns 2010). Neural recording could be used in various ways to gather data of value to companies (Jebari 2013). Data gathered on reward-processing regions of the brain (such as the nucleus accumbens) is used to fine-tune advertising to particular individuals or groups of individuals (Haynes, John-Dylan 2012). Neural data also could be gathered early in the design process to make products that better fit user needs. Further, the commercialization of devices capable of measuring neural data (e.g.,

---

[4] See, e.g., http://www.lawneuro.org/, accessed October 30, 2016.

Draft: Please cite to published version in Nam, Nijholt, and Lotte (eds) Brain–Computer Interfaces Handbook: Technological and Theoretical Advances (CRC Press: 2018).

portable EEGs), such as for gaming or wellness, may create a wide conduit for collection of neural data by companies. At present, neuromarketing represents a relatively unregulated commercial area. As techniques for measuring and finding patterns in neural activity continue to improve, legitimate concerns will be raised about whether neuromarketing technologies are akin to acceptable forms of persuasion or whether they constitute a kind of manipulation or hijacking of the subconscious.

These kinds of privacy concerns do not seem particularly relevant to the current state of BCI decoding of intentions. This is for two reasons. The first is that decoding of intentions is largely a *willed* process in which participants engage in arduous training to learn to isolate and execute an intention for a given task (e.g., to move a cursor, activate a prosthetic). This process can take weeks (or longer) to have a desired effect on the world (e.g., to control a prosthetic, computer screen, stimulation levels of implanted electrodes, etc.). Second, this process is *voluntary* and done with the consent of the participant. As (Owen 2012) notes with regard to fMRI, "[L]ike raising an arm in response to the instruction to do so, activating the brain by, say, imagining playing tennis, is a voluntary response, which can be suppressed at will….[and thus] poses no more of an ethical issue than observing that same participant outside the scanner and asking them to raise their left arm when told to do so."(Owen 2012, 84–85). A BCI motor intention task might be viewed similarly. If we confine ourselves to voluntary and deliberate decoding of intentions in a research setting - and not to talk of brain translating machines (Edwards 2012) – are there still privacy concerns?

Privacy concerns can still arise in the context of voluntary participation in BCI research if *additional* information about brain function can be gathered alongside the conduct of the primary experiment. For instance, BCI can be used to covertly monitor other cognitive processes while a research participant is engaged in BCI research task (Zander and Jatzev 2009). Participants may be wholly unaware of this additional monitoring and may even deny that they are engaging in the cognitive activity being monitored at all (e.g., subconsciously attending to some background feature of the environment). Further, passive

BCI monitoring can be combined with interventions to modulate or manipulate cognitive states. For instance, a BCI has been used to measure affective states (e.g., happiness, tranquility) and select music via algorithm that modulates these affective states (Daly et al. 2016).

Even if we confine ourselves to the decoding of motor intentions, different kinds of privacy concerns are relevant. The first might be what we call a *voyeuristic* concern. This is the concern that someone might peer into a private realm (one's private intentions) without invitation (an intrusion upon informational privacy) and the individual might feel violated. There are protections against this in the context of research. The second we might call a *maliciousness* concern. Here the concern is that knowing someone's intention could allow external actors to pry the lid off one's mental life, so to speak, and muck around with vicious intent, implicating a person's decisional privacy. Again, research protections are in place to prevent this. But there is another kind of concern here, what might be called a *collateral information* concern. The idea with the collateral concern is that additional information collected alongside targeted neural data has relevance to privacy (Haynes and Rees 2006). This is our focus here.

An obvious kind of collateral information from BCI studies of intention is incidental diagnostic findings. Electrode recording may turn up patterns of activity that indicate pathology. Consider two examples. First, changes in delta, theta, and beta band frequencies have been associated with Alzheimer's dementia (Soekadar, Cohen, Leonardo G., and Birbaumer 2014). Further, EEG in BCI studies could identify people with lower thresholds for future seizures. Second, neuroimaging needed to localize brain regions for electrode placement can identify brain pathology (e.g., tumor). Some of the ethical challenges of neuroimaging incidental findings in research are well recognized (Edwards 2012). As the BCI field develops, there will be more opportunity to correlate electrophysiologic data with signs and symptoms of clinical significance. As of now, this is rare but perhaps suggests that there should be put in place mechanisms in the future to collect this information.

An important and underappreciated privacy risk of collateral information is the effect on identity and sense of self. Take the example of someone in the Braingate trial conjuring up a motor intention ("I will try to move my right hand NOW"). Imagine what other information might be gathered in the process of learning to decode the relevant intention. Perhaps information about one's concentration or level of interest in a task might become evident (Curran and Stokes 2003). Perhaps a neural pattern portending an inability to master a BCI task (sometimes inaccurately called "BCI illiteracy") could become apparent (Blankertz et al. 2010). Could this lead to stigmatization or demoralization? Or maybe there is a neural signature when someone (e.g., a particular research assistant) walks in the room and this is found to correlate with BCI performance. Could this "extra fact" be used to motivate (or manipulate) research participants or allow research teams to infer something about an individual's attitude towards race, attractiveness, or authority?

One can get a sense of the challenge here by thinking about the close connection between intention and action. Imagine a BCI device capable of decoding intentions for operating an electronic wheelchair ("go left", "speed up", "avoid hitting the pedestrian", "stop"). It is important not just that intentions be accurately decoded and implemented, but that there be "checks" on putting intentions into action. For instance, if my nemesis walks into a room I would not want my pre-reflective intention "I really want to run over his foot" to immediately be an operative command for the wheelchair. Even if able to override the intention with another, the fact that the wheelchair moves in his direction *says something about me* ("I'm hotheaded," "not as collegial as I purport to be," or worse).

This challenge is made even more evident by decoding intentions related to communication. Though one might want a BCI for communication to be fast to facilitate fluid conversation, there is a risk of all intentions getting through, for instance, "I never intend to see you again!". The risk of not filtering or a failure to adequately filter actionable intentions is profound. Telling a stranger this may be very different than telling one's life partner, who might reasonably think: "I never thought you would be the kind of

person who would say such a thing." Neurotechnologies, like BCI, make us "answerable" for a wider range of our mental life (Richmond 2012). We are not just responsible for what we verbalize or our body language, but also for what can be decoded. Part of this may be addressed by developing BCI error detection mechanisms, but the challenge is not merely technical. What counts as an actionable intention, in part, depends on who I take myself to be and how I project myself into the future. This will be particularly hard to incorporate into a decision algorithm.[5]

It is also important to recognize that device output controlled by intentions can magnify the identity concern. DBS has been investigated as a potential treatment of medication-refractory depression and OCD. BCI control is being investigated as a way to modulate DBS to meet clinical need. For instance, an individual could increase stimulation when feeling depressed by forming a particular intention. There are potential benefits to such control, including battery conservation and enhanced patient sense of agency. Higher levels of DBS stimulation in particular brain regions have known side effects, however, including hypomania, which can lead to gambling, hypersexuality, and overspending. An individual who has control of DBS stimulation levels is likely be to viewed as responsible in part for these effects. What might before have been viewed as unfortunate side effects of DBS, may come instead to be viewed as the outcome of choice: "You must be the kind of person who wants to engage in such behavior." This implicates decisional privacy in that others may view one as accountable for DBS stimulation levels and hence responsible for underlying side effects, which puts constraints (viz., disapproval) on their ability to undergo DBS in the first place.

4. BCI and privacy opportunity

---

[5] Note that this is a case that blurs the distinction between informational and decisional privacy. There is information conveyed about persons' decision procedures, but at root the issue is about what kinds of intentions suffice to render something a decision to act.

A different way of understanding the privacy implications of BCI is as a kind of *opportunity for* privacy. Much of the literature addressing privacy in the context of clinical research, clinical care, commercial transactions, employment, law enforcement, security, and so forth addresses ways in which individuals' privacy may be diminished and the conditions under which such privacy loss is morally, legally, or socially problematic. Under this "loss model," individuals have privacy (to some degree) and others may infringe that privacy to the individuals' detriment. There are a couple of ways, however, in which the loss model is inadequate, and this has implications for BCI.

First, the loss model implies that privacy always has a positive moral value, and hence that any loss of privacy demands a justification. Compare this with other moral goods, such as freedom of expression or freedom of conscience. Any instance in which a person's ability to express herself freely, or in which her ability to exercise matters of conscience is limited is prima facie bad and demands some kind of reason. That is not to say that such limitations are unjustified—far from it. Rather, it just means that the limitations require *some* kind of justification. For example, limitations of free expression may be justifiable to prevent harm or fraud. Privacy, though, is not like that. Given the definition we offer above, once we determine that some person's privacy is diminished, there is a further question as to whether that diminution is negative at all. So, when a person walks down a grocery store aisle, his privacy regarding his food choices are diminished with respect to others in the store. But there is no sense in which this simple fact presents a moral problem or demands a justification.

To understand why this is important in the context of BCI, it's useful to consider Ruth Gavison's seminal article examining the concept and value of privacy. Similar to our approach in this paper, Gavison's view is that "privacy is a limitation of others' access to an individual" (Gavison 1980, 428). Perfect privacy, then, is a condition in which one "is completely inaccessible to others…no one has any information about X, no one pays any attention to X, and no one has physical access to X" (Gavison 1980, 428). On this

view, Gavison notes, having *too much* privacy is undesirable, and indeed people "may resent privacy that is imposed on them against their will" (Gavison 1980, 428, n. 24).

Now consider Locked-In Syndrome (LIS), a rare neurological disorder, often caused by a vascular or traumatic brain stem injury or by late-stage amyoptrophic lateral sclerosis (ALS or "Lou Gehrig's disease") (Walter 2010). Patients with LIS have voluntary motor paralysis and are hence unable to communicate by speaking or movement, though unless they have complete LIS they are usually able to move their eyes and eyelids (Walter 2010, 62).[6] Despite physical paralysis, people with LIS typically retain consciousness, self- and environmental awareness, and cognitive function.

The experience of LIS is recounted in Jean Dominique Bauby's book *The Diving Bell and Butterfly*, the content of which Bauby conveyed to a transcriber by blinking a code for letters of the alphabet (Bauby 1998). The book conveys Bauby's frustrations with being unable to share both his most basic needs and desires (e.g., physical comfort) and his more complex and nuanced thoughts. Until he can actually communicate (via blinking, a code, and a skilled transcriber) his mental states are "completely inaccessible to others…no one has any information about [his thoughts], [and] no one pays any attention to" his thoughts. He has, in other words, almost perfect informational and physical privacy regarding his conscious life, with respect to *all* others. And this is profoundly alienating. Denise Dudzinski explains that Bauby realizes that he "can be apprehended by others as an object in their world to be ignored or to be noticed," and that this affects Bauby's understanding of himself and his identity (Dudzinski 2001, 37).

Recent work in BCI has demonstrated that a patient with LIS following a brainstem stroke could communicate successfully using an electroencephalography (EEG) BCI (Sellers, Vaughan, and Wolpaw 2010), and this has generated optimism for further work on the basis of "the tremendous potential of non-invasive BCIs to cope with the unbearable condition of complete isolation from the social environment"

---

[6] LIS may be classified by severity, ranging from total immobility (including eye and eyelid movement), classic LIS (quadriplegia, anarthria, eye movement), and incomplete LIS (some degree of voluntary movement beyond eyes and eyelids) (Walter 2010).

Draft: Please cite to published version in Nam, Nijholt, and Lotte (eds) Brain–Computer Interfaces Handbook: Technological and Theoretical Advances (CRC Press: 2018).

(Chaudhary and Birbaumer 2015). Further, studies with communication devices controlled by implantable electrodes are underway.[7] The benefits of such advances are undeniable; our point here is that the benefit is constituted in substantial part by *decreasing* privacy.

A second way in which BCI presents a privacy opportunity focuses on relational aspects of privacy. In section 2, we explain that it's useful to think of privacy as a three-part relation between some person, some domain, and some other person or persons with respect to whom the first person has privacy. Specifying particular privacy relationships is key to understanding one of privacy's main values—it's effects on personal relationships. In an early article outlining moral foundations for privacy claims, James Rachels argues that privacy is vital in fostering many and varied social relationships. So, one might wish to have substantial privacy regarding one's family life with respect to one's work colleagues so as to maintain a professional distance, and one defining feature of intimate relationships is that they involve sharing important things (which privacy facilitates) (Rachels 1975). Indeed, what one shares (information, access to one's person) is a defining feature of different kinds of relationships.

Now consider patients with severe spinal cord injuries (SCI). Dreer et al. (2007) explain that because life expectancy of people with spinal cord injuries has been significantly extended due in part to improvements in medical technology, "family members often become the primary sources of assistance for various activities of daily living, such as feeding, dressing, transfers, and bowel and bladder care" (Dreer et al 2007, p. 2). In Collinger et al's study of attitudes toward BCI in persons with spinal cord injury, nearly 20% of participants with tetraplegia and over 30% of those with paraplegia received unpaid assistance with self-care activities or mobility by family members or others (Collinger et al. 2013). More generally, there is substantial evidence that many family members taking on caregiving roles encounter a range of difficulties: emotional, psychological, physical, financial. Providing care, of course, may have positive effects as well, such as fostering deeper, more meaningful relationships with loved ones, and

---

[7] See, e.g., [http://neuroprosthesis.edu; Hochberg NIH#5R01DC009899-05)].

Draft: Please cite to published version in Nam, Nijholt, and Lotte (eds) Brain–Computer Interfaces Handbook: Technological and Theoretical Advances (CRC Press: 2018).

evidence shows that in many cases patients and caregivers value the way care provision functions in their relationship (Donelan et al. 2002). This range of experience within caregiving has led some to propose increased support and training for family caregivers to allow caregiving to be both less difficult and more rewarding (Donelan et al 2002).

BCI-controlled devices, like neuroprosthetics in SCI, will have significant effects on caregiving. Though the scope of such effects may vary by condition and device, current effects of BCI research on family and caregivers is instructive. For example, wearable EEG-based BCI devices can require extensive involvement of caregivers in daily set-up, calibration, and removal (Mak and Wolpaw 2009). In one study of patients with ALS and their caregivers, the time and effort of caregivers in BCI set-up was a significant concern (Blain-Moraes et al. 2012). Implantable BCI devices would substantially reduce BCI setup demands and thereby offer greater independence. This may explain interest in implantable devices found in studies of persons with SCI and ALS (Huggins, Wren, and Gruis 2011; Collinger et al. 2013). Blain-Mores et al (2012) quote one person with ALS as saying "'If it made it easier on the caretaker, I'd go with the implant' (Blain-Moraes et al. 2012, 520). Caregivers also recognize that BCI-controlled devices, by affording their loved ones a greater ability to manage their own activities of daily living, could reduce caregiving demands. "It can alleviate concerns that you have, it could give you, the caregiver, more time to maybe take care of things that you need done" (Blain-Moraes et al. 2012, 521). BCI-controlled devices need not yield functional independence or completely obviate the need for care provision in order to have benefits for both users and caregivers. BCI-controlled devices may allow patients and their families to structure caregiving in a way that decreases the intensity of the need, allowing for a patient to have greater physical privacy in some areas of their life (e.g., eating or toileting), and thus diminish the stress upon family caregivers. Moreover, that greater opportunity for physical privacy may allow patients and family caregivers to better appreciate the aspects of caregiving that they already value, and hence to better foster the relationship overall. In other words, to the extent that BCI-

controlled devices increase a patient's physical privacy regarding some aspects of daily living, they may help support other values—caregiver well-being and flourishing relationships.

5.   Privacy and BCI Big Data

Neural data sets from BCI research, particularly electroencephalography (EEG) or intracortical single or multi-neuron recordings constitute a kind of "big data". Big data can refer to both the complexity of analyzing large quantities of data and to large data sets themselves (Mittelstadt and Floridi 2016). BCI data sets can be "big" in both procedural and quantitative terms, with single research participant studies generating uncompressed data in gigabyte or terabyte quantities, and the capacity to dramatically increase quantity and complexity by adding electrodes or recording channels. Similar to other kinds of "big data" in neuroscience, such as neuroimaging, neurogenetics, and behavioral data, BCI data holds promise for improving diagnosis, treatment and prevention of disease, but presents challenges of collection, storage, analysis and transmission. Informational privacy is a central challenge to the rapid accretion of BCI data.

The next big advances in neuroscience will come from integration (across measuring technologies, labs, neuroscientific subdisciplines, target problems, data analysis techniques), and this integration "will require a cultural shift in the way that data are shared across labs" (Sejnowski, Churchland, and Movshon 2014). While there are countervailing pressures to data sharing in neuroscience – such as a lack of uniform attitudes toward data sharing across subdisciplines (Van Horn and Ball 2008) or unsettled publishing and data ownership norms (Poldrack and Poline 2015) – neuroscientific data sharing in general maximizes the contribution of research participants, generates new scientific questions, enhances reproducibility, improves research practices, provides a test bed for new analysis methods, reduces the cost of doing science, and protects valuable scientific resources (Poldrack and Gorgolewski 2014).

Several forces push in the direction of data sharing in BCI research. For instance, data from implantable BCI can be expensive to generate because of costs associated with building an implantable device as well as costs associated with developing surgical implantation and other kinds of related expertise. Implantable human BCI research has tended to be confined to a small number of academic centers. In addition, the cohorts within dedicated implantable BCI studies can be small, involving only several research participants at a time or can piggyback on coincident clinical procedures (e.g., temporary placement of implanted electrodes for presurgical epilepsy localization (Blakely et al. 2014).

BCI data can be shared in different ways, which raise different kinds of privacy concerns. First, BCI data can be shared informally between labs or through formal lab data-sharing agreements. Sharing of raw neural data between BCI labs can spur development and facilitate reproducibility of new analytic methods, lead to further collaborations, and reduce financial and ethical costs of redundancy within a small field. Sharing of deidentified neural data sets can raise security concerns if not encrypted or procedures are not in place to prevent reassociation with identity-compromising data (e.g., MRI images used for lead placement if not scrubbed by de-facing software). Second, BCI data can be shared through large data repositories. Repositories aggregate data sets explicitly for research purposes to allow for complex forms of data linking and mining. Neuroimaging repositories provide an example of this (Ozyurt et al. 2010) as does the Human Connectome Project (HCP). Of particular relevance to BCI, efforts have been made to collect EEGs and make these available to researchers (https://www.nedcdata.org/drupal/). Privacy concerns raised about data repositories more generally (Sorani et al. 2015; Poldrack and Gorgolewski 2014) will arise with BCI as well. Third, BCI data can be shared through new forms of scientific practice, such as crowd sourcing or coding contests using donated BCI data (for instance, the "Decoding Brain Signals" contest sponsored by Microsoft).[8] Particularly challenging questions about ownership of data (and corresponding analytic methods) can be raised by

---

[8] See, https://gallery.cortanaintelligence.com/Competition/e9d67f7668e048328b8bb3a4a81fa5e7. Accessed November 1, 2016.

Draft: Please cite to published version in Nam, Nijholt, and Lotte (eds) Brain–Computer Interfaces Handbook: Technological and Theoretical Advances (CRC Press: 2018).

novel forms of scientific activity. Who owns the data and the analytic methods in this context – the researcher ("contestant"), lab, institution, research participant, sponsor?

BCI big data raises privacy concerns related to deidentification. Research regulations require the stripping of identifiers from BCI data. There are questions about the technical adequacy of deidentification, particularly when one form of data (e.g., BCI neural data) can be combined with other data – genetic or microbiomic sequencing data, biological specimens, electronic medical records, administrative hospital data, or other forms of neural data (e.g., MRI for localizing lead placement). A more specific concern in BCI research is risk of re-identification of research participants because of the small size of BCI research studies and media publicity that such studies attract (Pelley 2012; Neergaard 2016). Publications that list research participant gender, age, or type of medical condition (e.g., spinal cord injury versus ALS) add to this risk of reidentification.

The concern about identification or re-identification of research participants has less to do with what can currently be inferred from BCI data and more to do with the potential that this data may hold. At present, what can be inferred from a neural data set of voltage spikes of a single implanted electrode, for instance, is very narrow (e.g., intending a specific body movement at time T). A log of blood pressure and heart rate readings allows for more inference at present (Undergoing a stress response? Have a pacemaker? Taking a beta blocker anti-hypertensive? Autonomic dysfunction from a disease like diabetes or Parkinson's disease?), but even here a log of deidentified readings contains little in the way of private information about personality or behavior. For instance, one cannot infer from such a log absent other information whether your blood pressure goes up when you see your spouse or a member of a different race. BCI data stripped of identifiers and context might seem to be as innocuous as a deidentified blood pressure log.

But there are important differences. The first is that we don't know how much one *could* possibly infer from a log of neural data given advances in decoding algorithms. We have a pretty good sense of the

range of possible inferences from deidentified and decontextualized cardiac data; the same cannot be said of brain data. There may be a lot more *in there* to be decoded. At present, we just don't know. And, what could be in there are patterns that are revelatory of important – and private – aspects of the self. So, maybe BCI data could be decoded not only to indicate a stress response (as a cardiac log might), but also to associate it with some emotion or personality characteristic (e.g., shame, animus, jealousy, sexual desire). These kinds of personality and behavioral inferences would have significant implications for privacy. BCI also may contain more fine-grained information. For instance, given a data set from a BCI speller it might be possible to decode what an individual was saying by reverse engineering the linguistic content from the data. Decoding that a research participant thought "left" or "thirsty" or "bored" might be uninteresting, but thoughts of the sort "I (no longer) love my husband" or "This is boring…and my life is empty" might reveal thoughts that are certainly of a private character. Thus, the *potential* of BCI data matters for privacy. BCI data is stored not only because it is valuable *now*, but because it is of uncertain, yet potentially large, value in the future. But it is this potential value that also carries with it potential privacy risks.

Now, these privacy concerns reflect some of the issues that we have outlined in previous sections. One of the key issues in the large-scale collection of BCI data is that the repercussions (if any) would be years in the future, yet decisions about collection will be made in the present. And those future privacy issues cannot effectively be addressed by current research informed consent regimes. In a literature review of ethical concerns in biomedical big data, Mittelstadt and Floridi (2016) outline several key concerns about informed consent to collection, analysis, and use of big data that are applicable to BCI. One is that the notion of consent as an ability to control personal information is based on an assumption that people can reasonably infer what data will reveal. But of course, even researchers don't know this, and hence any research participant's conjectures are at least as limited. Likewise, the idea behind collecting and storing large volumes of BCI data is the potential to develop algorithms to find unforeseen connections; this, too, makes meaningful, explicit consent implausible (Choudhury et al. 2014).

Draft: Please cite to published version in Nam, Nijholt, and Lotte (eds) Brain–Computer Interfaces Handbook: Technological and Theoretical Advances (CRC Press: 2018).

One might argue that *broad* consent—to collect and use data for any purpose—is possible. (Mittelstadt and Floridi 2016, citing Clayton 2005; Ioannidis 2013). That, however, may conflict with the paradigm that informed consent be based on individual, autonomous decisions. Refinements to broad consent, such as tiered consent – letting research participants select among options for how data will be used – or opt-out consent – presuming broad consent unless subjects explicitly designate unacceptable uses for data – may be of limited value. For instance, the boundaries between different areas of and uses for BCI may be too porous to support tiers or specific opt-outs in practice. e.g., to help people with SCI reanimate a limb but not to contribute to BCI-control of exoskeletons for potential military uses. More generally, current consent paradigms consider potential ramifications for the individual while ignoring how such research may affect groups to which individuals are a part. So, if BCI currently targets certain populations (e.g., spinal cord injury, ALS) – which it does – then this research may have effects on the group (e.g. discrimination, stigmatization) that are not considered as part of the consent process. It needs to be noted that many of the challenges for informed consent described here are not unique to BCI. The important point to recognize is that while informed consent processes in BCI are important and valuable, they do not provide ready solutions to problems of privacy in BCI research. More work will need to be done.

6.   Conclusion

Summing up, we can see a number of ways in which privacy is important in BCI. BCI implicates physical, informational, and decisional privacy. More importantly, BCI raises a number of morally significant privacy issues, including mental or cognitive privacy, potential for privacy opportunity, and unknown, downriver effects of large-scale collection, analysis, and use of BCI research data, the effects of which neither researchers nor research participants can discern now.

Our focus here has been on exploring and delineating privacy issues raised by BCI research and their ethical implications. Getting clear on what privacy means in BCI research and ways in which it is (or is

not) valuable is an important step in the development of strategies for addressing privacy concerns. There are models for this further step of developing privacy policies and practices from which the BCI community can benefit. For instance, the NIH Genomic Data Sharing (GDS) Policy provides an example of how to address data privacy concerns through the granting process (National Institutes of Health 2014). In addition, the INCF Task Force on Neuroimaging Datasharing provides an example of how an emerging neuroscientific field (i.e., neuroimaging) can develop and articulate privacy and data sharing standards (Poline et al. 2012). And the Genetic Information Nondiscrimination Act (GINA 2008) in the U.S. is an example of a legal and regulatory framework for protecting privacy related to genetic information. Even with such models, more work will need to be done to develop privacy policies tailored to the particular features of BCI research data and practice.

But we have had to put aside several important debates. One is that different modes of brain data collection (ECog, EEG, intracortical electrodes) currently yield brain data of different quality, quantity, cost, and reproducibility. While these differences may have privacy implications, it is difficult to anticipate how developments in technology will affect these and, as such, for simplicity we have considered these together as a source of BCI data. Another is the philosophical question of whether decoding neural activity to infer mental states is conceptually coherent or misguided. There is a rich debate in the philosophy of mind about the relation between neural activity and mental states, such as that about the theory of extended mind, that also has implications for BCI (Heersmink 2011). While the extension of one's mind out into a neuroprosthetic, for instance, may have implications for privacy (is confiscating a prosthetic an invasion of physical privacy?), we think that such questions can be tabled for now. Lastly, BCI raises concern not just about privacy but about *security* (or "neurosecurity") and hacking of data and devices (Denning, Matsuoka, and Kohno 2009; Bonaci, Calo, and Chizeck 2014; Ienca and Haselager 2016). That is clearly an important and far-reaching issue; we've focused on privacy here because security is important either for reasons that are independent of privacy (e.g., harms to patients or compromise of devices) or *because of* the importance of privacy. We take security to be

significant in part because of the importance of privacy, and we assume that to the extent that privacy is valuable, security is all the more important.

References

Allen, Anita L. 1988. *Uneasy Access: Privacy for Women in a Free Society*. Totowa, N.J.: Rowman & Littlefield.

Ariely, Dan, and Gregory S. Berns. 2010. "Neuromarketing: The Hope and Hype of Neuroimaging in Business." *Nature Reviews Neuroscience* 11 (4): 284–292.

Bauby, Jean-Dominique. 1998. *The Diving Bell and the Butterfly: A Memoir of Life in Death*. Vintage Books.

Blaauw, Martijn. 2013. "The Epistemic Account of Privacy." *Episteme* 10 (Special Issue 02): 167–77. doi:10.1017/epi.2013.12.

Blabe, Christine H, Vikash Gilja, Cindy A Chestek, Krishna V Shenoy, Kim D Anderson, and Jaimie M Henderson. 2015. "Assessment of Brain–machine Interfaces from the Perspective of People with Paralysis." *Journal of Neural Engineering* 12 (4): 043002. doi:10.1088/1741-2560/12/4/043002.

Blain-Moraes, Stefanie, Riley Schaff, Kirsten L. Gruis, Jane E. Huggins, and Patricia A. Wren. 2012. "Barriers to and Mediators of Brain–computer Interface User Acceptance: Focus Group Findings." *Ergonomics* 55 (5): 516–525.

Blankertz, Benjamin, Claudia Sannelli, Sebastian Halder, Eva M. Hammer, Andrea Kübler, Klaus-Robert Müller, Gabriel Curio, and Thorsten Dickhaus. 2010. "Neurophysiological Predictor of SMR-Based BCI Performance." *NeuroImage* 51 (4): 1303–9. doi:10.1016/j.neuroimage.2010.03.022.

Blakely, Tim M., Jared D. Olson, Kai J. Miller, Rajesh PN Rao, and Jeffrey G. Ojemann. 2014. "Neural Correlates of Learning in an Electrocorticographic Motor-Imagery Brain-Computer Interface." *Brain-Computer Interfaces* 1 (3–4): 147–157.

Bonaci, Tamara, Ryan Calo, and Howard J. Chizeck. 2014. "App Stores for the Brain: Privacy & Security in Brain-Computer Interfaces." In *Ethics in Science, Technology and Engineering, 2014 IEEE International Symposium on*, 1–7. IEEE. http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=6893415.

Chaudhary, Ujwal, and Niels Birbaumer. 2015. "Communication in Locked-in State after Brainstem Stroke: A Brain-Computer-Interface Approach." *Annals of Translational Medicine* 3 (Suppl 1). http://www.ncbi.nlm.nih.gov/pmc/articles/PMC4437925/.

Choudhury, Suparna, Jennifer R. Fishman, Michelle L. McGowan, and Eric T. Juengst. 2014. "Big Data, Open Science and the Brain: Lessons Learned from Genomics." *Frontiers in Human Neuroscience* 8. doi:10.3389/fnhum.2014.00239.

Clayton, Ellen Wright. 2005. "Informed Consent and Biobanks." *The Journal of Law, Medicine & Ethics* 33 (1): 15–21. doi:10.1111/j.1748-720X.2005.tb00206.x.

Collinger, Jennifer L., Michael L. Boninger, Tim M. Bruns, Kenneth Curley, Wei Wang, and Douglas J. Weber. 2013. "Functional Priorities, Assistive Technology, and Brain-Computer Interfaces after Spinal Cord Injury." *Journal of Rehabilitation Research and Development* 50 (2): 145.

Curran, Eleanor A., and Maria J. Stokes. 2003. "Learning to Control Brain Activity: A Review of the Production and Control of EEG Components for Driving Brain–computer Interface (BCI) Systems." *Brain and Cognition* 51 (3): 326–336.

Daly, Ian, Duncan Williams, Alexis Kirke, James Weaver, Asad Malik, Faustina Hwang, Eduardo Miranda, and Slawomir J. Nasuto. 2016. "Affective Brain–computer Music Interfacing." *Journal of Neural Engineering* 13 (4): 046022. doi:10.1088/1741-2560/13/4/046022.

DeCew, Judith. 2016. "Connecting Tort, Fourth Amendment, and Constitutional Privacy." In *Privacy, Security, and Accountability*, edited by Moore, Adam D, 73–88. New Jersey: Rowman & Littlefield.

DeCew, Judith. 1997. *In Pursuit of Privacy: Law, Ethics, and the Rise of Technology*. Ithaca, N.Y.: Cornell University Press.

Denning, Tamara, Yoky Matsuoka, and Tadayoshi Kohno. 2009. "Neurosecurity: Security and Privacy for Neural Devices." *Neurosurgical Focus* 27 (1): E7.

Dewan, Edmond M. 1967. "Occipital Alpha Rhythm Eye Position and Lens Accommodation." *Nature* 214: 975–977.

Donelan, Karen, Craig A. Hill, Catherine Hoffman, Kimberly Scoles, Penny Hollander Feldman, Carol Levine, and David Gould. 2002. "Challenged To Care: Informal Caregivers In A Changing Health System." *Health Affairs* 21 (4): 222–31. doi:10.1377/hlthaff.21.4.222.

Dreer, Laura E., Timothy R. Elliott, Richard Shewchuk, Jack W. Berry, and Patricia Rivera. 2007. "Family Caregivers of Persons With Spinal Cord Injury: Predicting Caregivers at Risk for Probable Depression." *Rehabilitation Psychology* 52 (3): 351. doi:10.1037/0090-5550.52.3.351.

Dudzinski, Denise. 2001. "The Diving Bell Meets the Butterfly: Identity Lost andRe-Membered." *Theoretical Medicine and Bioethics* 22 (1): 33–46. doi:10.1023/A:1009981213630.

Edwards, Sarah JL. 2012. "Protecting Privacy Interests in Brain Images: The Limits of Consent." In *I Know What You're Thinking: Brain Imaging and Mental Privacy*, edited by Sarah Richmond, Sarah JL Rees, and Sarah JL Edwards, 245–60. Oxford: Oxford University Press.

European Union. 1995. Council Directive 95/46, on the Protection of Individuals with Regard to the Processing Personal Data and on the Free Movement of Such Data, 1995 O.J. (L281) 31 (EU).

Farah, M. J., and P. R. Wolpe. 2004. "Monitoring and Manipulating Brain Function: New Neuroscience Technologies and Their Ethical Implications." *Hastings Cent Rep* 34: 34–45.

Finn, Rachel L., David Wright, and Michael Friedewald. 2013. "Seven Types of Privacy." In *European Data Protection: Coming of Age*, 3–32. Springer. http://link.springer.com/chapter/10.1007/978-94-007-5170-5_1.

French, Jennifer. 2012. *On My Feet Again: My Journey out of the Wheelchair Using Neurotechnology*. San Francisco: Neurotech Press.

Gavison, Ruth. 1980. "Privacy and the Limits of Law." *The Yale Law Journal* 89 (3): 421–71. doi:10.2307/795891.

Genetic Information Nondiscrimination Act of 2008. (GINA) 2008. *U.S.C.* Vol. 42.

Greely, Henry T., and Judy Illes. 2007. "Neuroscience-Based Lie Detection: The Urgent Need for Regulation." *Am. JL & Med.* 33: 377.

Greene, Joshua D., R. Brian Sommerville, Leigh E. Nystrom, John M. Darley, and Jonathan D. Cohen. 2001. "An fMRI Investigation of Emotional Engagement in Moral Judgment." *Science* 293 (5537): 2105–2108.

Hallinan, Dara, Michael Friedewald, Philip Schütz, and Paul de Hert. 2014. "Neurodata and Neuroprivacy: Data Protection Outdated?" *Surveillance & Society* 12 (1): 55.

Haynes, John-Dylan. 2012. "Brain Reading." In *I Know What You're Thinking: Brain Imaging and Mental Privacy*, edited by Sarah Richmond, Geraint Rees, and Sarah JL Edwards, 29–40. Oxford: Oxford University Press.

Haynes, John-Dylan, and Geraint Rees. 2006. "Decoding Mental States from Brain Activity in Humans." *Nature Reviews Neuroscience* 7 (7): 523–534.

Haynes, John-Dylan, Katsuyuki Sakai, Geraint Rees, Sam Gilbert, Chris Frith, and Richard E. Passingham. 2007. "Reading Hidden Intentions in the Human Brain." *Current Biology* 17 (4): 323–328.

Health Insurance Portability and Accountability Act. (HIPAA) 1996. *U.S.C.* Vol. 29.

Hochberg, Leigh R., Daniel Bacher, Beata Jarosiewicz, Nicolas Y. Masse, John D. Simeral, Joern Vogel, Sami Haddadin, et al. 2012. "Reach and Grasp by People with Tetraplegia Using a Neurally Controlled Robotic Arm." *Nature* 485 (7398): 372–375.

Huggins, Jane E., Patricia A. Wren, and Kirsten L. Gruis. 2011. "What Would Brain-Computer Interface Users Want? Opinions and Priorities of Potential Users with Amyotrophic Lateral Sclerosis." *Amyotrophic Lateral Sclerosis* 12 (5): 318–324.

Ienca, Marcello, and Pim Haselager. 2016. "Hacking the Brain: Brain–computer Interfacing Technology and the Ethics of Neurosecurity." *Ethics and Information Technology*, 1–13.

Illes, J., and E. Racine. 2005. "Imaging or Imagining? A Neuroethics Challenge Informed by Genetics." *Am J Bioeth* 5 (2): 5–18.

Ioannidis, John P. A. 2013. "Informed Consent, Big Data, and the Oxymoron of Research That Is Not Research." *The American Journal of Bioethics*, March. http://www.tandfonline.com/doi/abs/10.1080/15265161.2013.768864.

Jebari, Karim. 2013. "Brain Machine Interface and Human Enhancement–An Ethical Review." *Neuroethics* 6 (3): 617–625.

Klein, Eran. 2015. "Informed Consent in Implantable BCI Research: Identifying Risks and Exploring Meaning." *Science and Engineering Ethics*, 1–19.

Lacy, Joyce W., and Craig E. L. Stark. 2013. "The Neuroscience of Memory: Implications for the Courtroom." *Nature Reviews Neuroscience* 14 (9): 649–58. doi:10.1038/nrn3563.

Lahr, Jacob, Christina Schwartz, Bernhard Heimbach, Ad Aertsen, Jörn Rickert, and Tonio Ball. 2015. "Invasive Brain–machine Interfaces: A Survey of Paralyzed Patients' Attitudes, Knowledge and Methods of Information Retrieval." *Journal of Neural Engineering* 12 (4): 043001.

Levy, Neil. 2007. *Neuroethics*. Cambridge, UK ; New York: Cambridge University Press.

Mak, Joseph N., and Jonathan R. Wolpaw. 2009. "Clinical Applications of Brain-Computer Interfaces: Current State and Future Prospects." *Biomedical Engineering, IEEE Reviews in* 2: 187–199.

Matheson, David. 2007. "Unknowableness and Informational Privacy." *Journal of Philosophical Research* 32: 251–67.

Mittelstadt, Brent Daniel, and Luciano Floridi. 2016. "The Ethics of Big Data: Current and Foreseeable Issues in Biomedical Contexts." *Science and Engineering Ethics* 22 (2): 303–41. doi:10.1007/s11948-015-9652-2.

National Institutes of Health. 2014. "Supplemental Information to the National Institutes of Health Genomic Data Sharing Policy." https://gds.nih.gov/pdf/supplemental_info_GDS_Policy.pdf.

Neergaard, Lauran. 2016. "Obama Shakes Mind-Controlled Robot Hand Wired to Sense Touch." *US News & World Report*, October 13. http://www.usnews.com/news/news/articles/2016-10-13/paralyzed-man-feels-touch-through-mind-controlled-robot-hand.

Nissenbaum, Helen Fay. 2010. *Privacy in Context: Technology, Policy, and the Integrity of Social Life*. Stanford, Calif: Stanford Law Books.

Owen, Adrian. 2012. "When Thoughts Become Actions: Neuroimaging in Non-Responsive Patients." In *I Know What You're Thinking: Brain Imaging and Mental Privacy*, edited by Sarah Richmond, Sarah JL Rees, and Sarah JL Edwards, 73–88. Oxford: Oxford University Press.

Ozyurt, I. Burak, David B. Keator, Dingying Wei, Christine Fennema-Notestine, Karen R. Pease, Jeremy Bockholt, and Jeffrey S. Grethe. 2010. "Federated Web-Accessible Clinical Data Management within an Extensible Neuroimaging Database." *Neuroinformatics* 8 (4): 231–249.

Pelley, Scott. 2012. "Paralyzed Woman Uses Mind-Control Technology to Operate Robotic Arm." *60 Minutes*. http://www.cbsnews.com/news/paralyzed-woman-uses-mind-control-technology-to-operate-robotic-arm/.

Poldrack, Russell A., and Krzysztof J. Gorgolewski. 2014. "Making Big Data Open: Data Sharing in Neuroimaging." *Nature Neuroscience* 17 (11): 1510–1517.

Poldrack, Russell A., and Jean-Baptiste Poline. 2015. "The Publication and Reproducibility Challenges of Shared Data." *Trends in Cognitive Sciences* 19 (2): 59–61.

Poline, Jean-Baptiste, Janis L. Breeze, Satrajit Ghosh, Krzysztof Gorgolewski, Yaroslav O. Halchenko, Michael Hanke, Christian Haselgrove, et al. 2012. "Data Sharing in Neuroimaging Research." *Frontiers in Neuroinformatics* 6: 9. doi:10.3389/fninf.2012.00009.

Rachels, James. 1975. "Why Privacy Is Important." *Philosophy & Public Affairs* 4 (4): 323–33. doi:10.2307/2265077.

Räikkä, Juha. 2010. "Brain Imaging and Privacy." *Neuroethics* 3 (1): 5–12.

Richmond, Sarah. 2012. "Brain Imaging and the Transparency Scenario." In *I Know What You're Thinking: Brain Imaging and Mental Privacy*, edited by Sarah Richmond, Sarah JL Rees, and Sarah JL Edwards, 185–204. Oxford: Oxford University Press.

Rubel, Alan, and Ryan Biava. 2014. "A Framework for Analyzing and Comparing Privacy States." *Journal of the Association for Information Science and Technology* 65 (12): 2422–31. doi:10.1002/asi.23138.

Ryberg, Jesper. 2016. "Neuroscience, Mind Reading and Mental Privacy." *Res Publica* . doi:10.1007/s11158-016-9343-0

Schneider, Mary-Jane, Joseph J. Fins, and Jonathan R. Wolpaw. 2012. "Ethical Issues in BCI Research." *Brain-Computer Interfaces: Principles and Practice*, 373–83.

Schoenmakers, Sanne, Markus Barth, Tom Heskes, and Marcel van Gerven. 2013. "Linear Reconstruction of Perceived Images from Human Brain Activity." *NeuroImage* 83: 951–961.

Sejnowski, Terrence J., Patricia S. Churchland, and J. Anthony Movshon. 2014. "Putting Big Data to Good Use in Neuroscience." *Nature Neuroscience* 17 (11): 1440–1441.

Sellers, Eric W., Theresa M. Vaughan, and Jonathan R. Wolpaw. 2010. "A Brain-Computer Interface for Long-Term Independent Home Use." *Amyotrophic Lateral Sclerosis* 11 (5): 449–455.

Shih, Jerry J., Dean J. Krusienski, and Jonathan R. Wolpaw. 2012. "Brain-Computer Interfaces in Medicine." In *Mayo Clinic Proceedings*, 87:268–279. Elsevier. http://www.sciencedirect.com/science/article/pii/S0025619612001231.

Soekadar, Surjo R., Cohen, Leonardo G., and Niels Birbaumer. 2014. "Clinical Brain-Machine Interfaces." In *Cognitive Plasticity in Neurologic Disorders*, edited by Tracy, Joseph I., Hamstead, Bejamin M., and Sathian, K., 347–62. Oxford: Oxford University Press.

Sorani, Marco D., John K. Yue, Sourabh Sharma, Geoffrey T. Manley, Adam R. Ferguson, Shelly R. Cooper, Kristen Dams-O'Connor, et al. 2015. "Genetic Data Sharing and Privacy." *Neuroinformatics* 13 (1): 1–6.

Tennison, Michael N., and Jonathan D. Moreno. 2012. "Neuroscience, Ethics, and National Security: The State of the Art." *PLoS Biol* 10 (3): e1001289.

Trimper, John B., Paul Root Wolpe, and Karen S. Rommelfanger. 2014. "When 'I' Becomes 'We': Ethical Implications of Emerging Brain-to-Brain Interfacing Technologies." *Frontiers in Neuroengineering* 7. http://www.ncbi.nlm.nih.gov/pmc/articles/PMC3921579/.

Van Horn, John Darrell, and Catherine A. Ball. 2008. "Domain-Specific Data Sharing in Neuroscience: What Do We Have to Learn from Each Other?" *Neuroinformatics* 6 (2): 117–121.

Walter, Sven. 2010. "Locked-in Syndrome, BCI, and a Confusion about Embodied, Embedded, Extended, and Enacted Cognition." *Neuroethics* 3 (1): 61–72. doi:10.1007/s12152-009-9050-z.

Wolpe, P. R., K. R. Foster, and D. D. Langleben. 2005. "Emerging Neurotechnologies for Lie-Detection: Promises and Perils." *The American Journal of Bioethics : AJOB* 5 (2): 39–49.

Zander, T. O., and S. Jatzev. 2009. "Detecting Affective Covert User States with Passive Brain-Computer Interfaces." In *2009 3rd International Conference on Affective Computing and Intelligent Interaction and Workshops*, 1–9. doi:10.1109/ACII.2009.5349456.