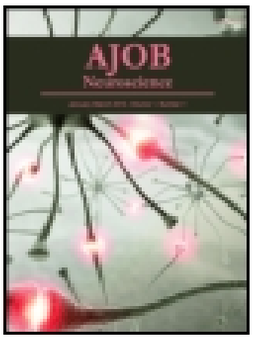

AJOB Neuroscience

ISSN: (Print) (Online) Journal homepage: https://www.tandfonline.com/loi/uabn20

Brain Data in Context: Are New Rights the Way to Mental and Brain Privacy?

Daniel Susser & Laura Y. Cabrera

To cite this article: Daniel Susser & Laura Y. Cabrera (2023): Brain Data in Context: Are New Rights the Way to Mental and Brain Privacy?, AJOB Neuroscience, DOI: 10.1080/21507740.2023.2188275

To link to this article: https://doi.org/10.1080/21507740.2023.2188275

Published online: 05 Apr 2023.

AJOB NEUROSCIENCE
https://doi.org/10.1080/21507740.2023.2188275

Taylor & Francis
Taylor & Francis Group

TARGET ARTICLE

Brain Data in Context: Are New Rights the Way to Mental and Brain Privacy?

Daniel Susser and Laura Y. Cabrera

The Pennsylvania State University

ABSTRACT

The potential to collect brain data more directly, with higher resolution, and in greater amounts has heightened worries about mental and brain privacy. In order to manage the risks to individuals posed by these privacy challenges, some have suggested codifying new privacy rights, including a right to "mental privacy." In this paper, we consider these arguments and conclude that while neurotechnologies do raise significant privacy concerns, such concerns are—at least for now—no different from those raised by other well-understood data collection technologies, such as gene sequencing tools and online surveillance. To better understand the privacy stakes of brain data, we suggest the use of a conceptual framework from information ethics, Helen Nissenbaum's "contextual integrity" theory. To illustrate the importance of context, we examine neurotechnologies and the information flows they produce in three familiar contexts—healthcare and medical research, criminal justice, and consumer marketing. We argue that by emphasizing what is distinct about brain privacy issues, rather than what they share with other data privacy concerns, risks weakening broader efforts to enact more robust privacy law and policy.

KEYWORDS

privacy; mental privacy; contextual integrity; brain data; neural data; neurotechnology

INTRODUCTION

Privacy has always been a central concern in neuroethics (Farah 2005; Farah and Wolpe 2004; Roskies 2002). Initially, privacy issues were raised mainly as long-term possibilities—existing neurotechnologies were still in their infancy and not yet efficient at accessing and interpreting brain data. But in recent years, with the development of novel neurotechnologies that allow new forms and greater amounts of data collection, privacy has become a more urgent issue in neuroethics and beyond. The potential to collect brain data more directly and with higher resolution has heightened worries about mental and brain privacy. For example, many of the recent guidelines around neurotechnology list privacy as a key issue to consider and address (Goering and Yuste 2016; Greely et al. 2018; Yuste et al. 2017). Some have argued that the specific kinds of data collected by neurotechnologies raise unique and unprecedented privacy challenges, because information about the brain is particularly revealing, it is less subject to conscious control, and there is unusual uncertainty about what can be inferred from it (Goering et al. 2021; Ienca and Andorno 2017). In order to manage the risks to individuals created by these data flows, some have proposed codifying new privacy rights, including a right to "mental privacy" (Ienca and Andorno 2017; Ienca 2021a; see also NeuroRights Initiative at https://neurorightsfoundation.org/mission).

We argue that this understanding of brain data's privacy risks stems from a particular conception of privacy—privacy understood as individual control over personal information. In recent years, developments in information privacy theory have cast doubt on "control theories," and new ways of conceptualizing privacy have emerged. In particular, we suggest the use of a conceptual framework from information ethics—Helen Nissenbaum's "contextual integrity" theory (2010)—to better understand the privacy stakes of brain data. According to this framework, privacy norms are inherently context-specific, and understanding the privacy implications of new data-driven technologies requires analyzing them in particular social contexts.

We proceed as follows: First, we explore how neuroethics scholars define privacy and why they think new neurotechnologies pose such grave privacy threats, and we ask if these concerns, and the rights proposed to remedy them, might be misplaced. Next, we introduce the theory of privacy as contextual integrity. To illustrate the importance of context, we examine

CONTACT Laura Y. Cabrera lcabrera@psu.edu Dorothy Foehr Huck and J. Lloyd Huck Chair in Neuroethics. Department of Engineering Science and Mechanics, and Rock Ethics Institute, The Pennsylvania State University, W-316 Millennium Science Complex, University Park, PA 16802, USA.

neurotechnologies and the information flows they produce in three contexts familiar from prior discussions—healthcare and medical research, criminal justice, and consumer marketing. In each case, we find that existing approaches to information privacy are adequate for theorizing the privacy implications of neurotechnologies and addressing the main worries raised by proponents of new mental privacy rights. Finally, we conclude with practical recommendations for how to understand and address privacy in the context of neurotechnologies, arguing that while neurotechnologies do raise significant privacy concerns, such concerns are—at least for now—no different from those raised by other well-understood data collection technologies, such as gene sequencing tools and online surveillance.

PRIVACY IN NEUROETHICS

The conception of privacy that predominates discussions in neuroethics derives from the information technology sphere: namely, privacy is understood as the ability to *control* the flow of personal information. Alan Westin famously defined informational privacy as our claim to determine for ourselves when, how, and to what extent information about us is communicated to others (Westin 1968). Likewise, and more recently, Moore defined it as the "right to maintain a certain level of control over the inner spheres of personal information and access to one's body" (2008, 420). Scholars in neuroethics have largely followed suit. Ienca and Andorno (2017) see privacy as "control over information about oneself." Goering et al. (2021) describe it as a "right that others not access one's personal information and personal space." And Schönau et al. (2021) describe privacy as the "right for individuals to establish their boundaries and dynamics with others."

Issues around privacy are not unique to neurotechnologies. In our homes, for example, we expect others to respect our physical privacy—that is, we expect to control who can access our property. In health and medical contexts, new technologies have created privacy challenges, for example, who should be allowed to access our genetic data, what they should be allowed to do with it, and whether consent should be required to share it. The proliferation of digital technologies more broadly, especially smartphones and wearable devices, has raised concerns about unauthorized access to the personal data they generate, such as location information and online behavioral data.

In neuroethics, recent debates have centered on worries about the ability of neurotechnologies to directly target neural processing (Goering et al. 2021; Ienca and Andorno 2017). Some have focused on concerns about unauthorized access to people's neural data[1] (Goering et al. 2021; Ienca and Andorno 2017; Schönau et al. 2021), or on issues around consent to sharing brain data, while others worry about the use of neural data to predict and manipulate people's behavior (Hallinan et al. 2013).

Proponents of the view that neurotechnologies threaten privacy in unique and unprecedented ways tend to make a number of important assumptions concerning the use of neurotechnologies, either implicitly or explicitly. First, many start from the premise that neural data is especially revealing—that it can expose particularly intimate aspects of the person (Goering et al. 2021; Ienca and Andorno 2017). Goering et al., for example, argue that "[s]uch data is sensitive by nature because it contains information about the organ that generates the mind." Second, and closely related to the first point, many believe that neural data has a special nature compared to other forms of personal data, as it is perceived to be connected to the core of who we are in ways that other data is not.[2] Some worry about the special *directness* of brain data—the idea that "the information to be protected is not easily distinguishable from the source itself that produced the data: the individual's neural processing." (Ienca and Andorno 2017, 14). Third, many hold that the risks posed by neural data are especially acute because it is less subject to conscious control (Goering et al. 2021; Ienca and Andorno 2017). Finally, some believe neurotechnologies enable a level of access and control of brain activity that is unprecedented.

However, not everyone is convinced of the arguments previously described. Another set of scholars offer compelling reasons to resist the view that neurotechnologies raise privacy concerns distinct from those raised by other technologies (Ryberg 2017). For example, genetic information is (arguably) equally intimate and novel, and when such information is

[1]In what follows, we adopt the IEEE WG P2794 definition of "neural data" as "all biosignals of neurological origin, including those recorded directly from neural tissues, and downstream biosignals (e.g. EMG)." It is important to point out that the discussion of privacy concerns resulting from neurotechnology often time uses interchangeably the term brain data (data directly obtained from the brain) and neural data. Also of note, a few years ago discussions of privacy were mostly elicited when discussing brain imaging technologies, nowadays the focus has been mostly on discussions around brain computer interfaces and brain implants.

[2]Others view this as a kind of "neuro exceptionalism" (Bublitz 2022; Tovino 2007)—the idea that brain data has a special nature relative to other personal information.

captured, there is just as much uncertainty about what could be inferred from it down the road (Hallinan et al. 2013; Ryberg 2017). Also, claims about the "directness" of neural data may be overblown; how we choose to filter, process, and display brain data introduces a series of uncertainties that call claims about its "directness" into question. Moreover, just as brain signals can be used to distinguish between and identify individuals, so too can genetic information, thus raising similar privacy concerns (Hallinan et al. 2013). It's true that—like genetic data—some forms of neural data may provide unique and persistent representations of single individuals, brain patterns often change based on "environmental conditions or even the mood of the person at the time," and the brain itself changes as individuals grow and develop (Hallinan et al. 2013), thus raising questions about the certainty with which brain data can be used to identify individuals over time, or to capture other intimate details about them.

It is unclear that we can infer sensitive information, such as a person's thoughts and feelings, just from recordings of neural or brain activity. For brain data or neural data to reveal higher order information about the mind it must be linked with other behavioral or physiological measures—it must be interpreted to say anything meaningful. In discussions about privacy in neuroethics, there is often a leap from claims about brain data—data about the function, status or structure of the brain, directly or indirectly recorded—to worries about having direct access to the mind. In Lippert-Rasmussen's terms, from "brain-reading" (i.e., "any technique that establishes neurological properties of the brain") to "mindreading" (a "technique that establishes the psychological properties of a person, such as her thoughts, feelings, emotions, psychological dispositions, etc."). As Lippert-Rasmussen warns, however, "not all mindreading involves brain-reading" (Lippert-Rasmussen 2017).[3] We add that not all brain-reading involves mindreading: on its own, recording and monitoring neural signals without other variables providing context tells us little if anything about our mental life.

It is certainly true that neurotechnologies allow us to access the brain in new and interesting ways, without doubt raising concerns about privacy. But that does not mean such concerns are necessarily different or more worrisome than those raised by any of the other data-driven technologies with which we interact. Despite the fact that, historically, we have lacked access to brain data, humans have nevertheless been able to access other people's minds. We are not persuaded by the above arguments, that "reading the brain" or accessing the brain in these neurotechnological ways really means acquiring more intimate, morally relevant information about the mind. As Ryberg suggests (2017), we need neuromodesty when discussing neurotechnologies and privacy.

DO WE NEED MENTAL AND BRAIN PRIVACY RIGHTS?

Human rights occupy an important place in contemporary moral and legal theory, as well as in practical politics more broadly. Given their broad scope and myriad of applications, we have recently seen a variety of different programs calling for a human rights approach, including in neuroethics discussions about mental and brain privacy.

In the early 2000s an important step toward neurorights was work on the notion of "cognitive liberty" (Boire 2001; Sententia 2004).[4] Almost a decade later Farahany (2012a) and Bublitz (2013) expanded the rights-based view of cognitive liberty. Though none of these authors used the term neurorights their work laid the foundations to the current scholarship covering neurorights. In 2017, Ienca and Andorno argued for four new rights: "the right to cognitive liberty, the right to mental privacy, the right to mental integrity, and the right to psychological continuity" (Ienca and Andorno 2017). In particular, Ienca and Andorno suggested the need for "formal recognition of a right to mental privacy, which aims to protect any bit or set of brain information about an individual recorded by a neurodevice and shared across the digital ecosystem," and "the right to brain privacy [which] aims to protect people against illegitimate access to their brain information and to prevent the indiscriminate leakage of brain data across the infosphere." Since Ienca and Andorno's paper others have joined this chorus: Yuste et al. (2017), for example, argue in a comment article in Nature that "citizens should have the ability—and right—to keep their neural data private." The Chilean congress created further momentum when, on April 12, 2021, it approved an amendment to article 19 of the Chilean Constitution endorsing "the rights to physical and mental integrity

[3]If we consider the implications of the extended mind thesis, namely that the mind is not only tied to the brain (Clark and Chalmers 1998; Lippert-Rasmussen 2017), then indeed looking at our browser history in our phone might be more revealing of our minds than certain brain data.

[4]For a more detailed discussion please see Ienca 2021b.

and protection of cerebral activity and data."[5] This undertaking was part of the NeuroRights Initiative (https://neurorightsfoundation.org/, last accessed it March, 2023), born out of a workshop at Columbia University, where an interdisciplinary group of global leaders discussed the ethical principles of neurotechnology and AI. The Initiative, which is now run out of the Columbia University Neurotechnology Center and directed by Dr. Rafael Yuste, recently released a list of five neurorights, among which they listed mental privacy and proposed that "any data obtained from measuring neural activity should be kept private. Moreover, the sale, commercial transfer, and use of neural data should be strictly regulated."[6]

By now the spectrum of this debate has become much broader, with some scholars arguing that no new rights or legal reforms are needed to address brain-related informational privacy, and others arguing that while no new fundamental rights are necessary, there is still a need for context-specific and multi-level approaches (like the approach we defend here, cf. Ienca et al. 2022; Fins 2022). This has motivated scholarship on the complex legal landscape that needs to be considered in brain data governance (cf. Ienca et al. 2022). While we acknowledge the existence of these more nuanced approaches, we focus in this paper on the position, often overhyped and overrepresented in certain advocacy efforts and international debates on the issue[7], that novel fundamental rights (such as a right to mental privacy) are needed.

We are not arguing that human rights approaches are inherently inappropriate. In fact, we agree that such frameworks could help influence global governance of data, pushing governments to respect, protect and fulfill privacy.[8] Human rights frameworks can provide valuable conceptual and rhetorical tools for demanding government action and accountability—in this case, for protecting personal data. Yet such an approach comes with issues regarding accountability, monitoring and enforcement. Member states (and their legal and political institutions) are generally seen as the actors responsible for upholding human rights, with local and regional institutions helping link international human rights norms to domestic actors (Engstrom 2017). Furthermore, we all as members of society are meant to play a role in ensuring human rights are upheld. All of this complicates views about who bears responsibility regarding the protection of human rights, who monitors progress and the mechanisms to enforces compliance. Some human rights scholars have pointed to the complex task of monitoring state's compliance, as often they are the ones in charge of reporting about compliance. This creates a state of affairs where there are not strong incentives to establish "forceful human rights mechanisms" to monitor and enforce human rights (Carraro 2019). States' commitments to protect even the most basic human rights continue to be inadequate, with several international human rights being under-enforced (Koh 1998). The main proponents of brain and mental privacy rights have not engaged enough with these benefits and pitfalls of rights-based approaches, instead focusing on the features of brain and mental data that might justify new rights (see also Bublitz 2022). Here we focus on this aspect of the discussion.

In our view, "NeuroRights" proponents have not sufficiently explained how these rights would differ, meaningfully, from existing informational privacy rights. For example, we generally regard "natural mind reading" as morally unproblematic. A key aspect of being a person is having a "theory of mind," which enables us to make inferences about the mental states of others (see also Lippert-Rasmussen 2017; Ryberg 2017), which—though limited—has proven to be a reasonably effective means of accessing the otherwise unobservable mental states of those around us. Furthermore, there are other less invasive and costly technologies (e.g. social media) that provide some insights into people's mental lives. Proponents of new brain and mental privacy rights argue that neurotechnologies enable more *direct* access. Yet as Ryberg (2017) and others point out, neurotechnologies process data in a variety of ways to make it intelligible to human users, from filters to amplifiers and so on, casting doubt on the idea of direct versus indirect access.

Given these challenges—to the idea that brain data is distinct from other personal information in a morally relevant way, and the notion that brain data is especially revealing of intimate information about us—there is reason for skepticism toward the claim that brain data raises unique, especially urgent privacy concerns. Some, like Ryberg (2017), argue that while individuals have a privacy interest, this is not sufficient to ground a "right to mental privacy." We agree. Moreover, it's likely that in many instances the types of brain information we are aiming to protect already

[5]https://nri.ntc.columbia.edu/projects

[6]NeuroRights Initiative, 2021

[7]We thank one of our reviewers for his/her suggestions on how to best characterize the different approaches on the current debate around neurorights.

[8]For a similar criticism see Bublitz. http://doi.org/10.1007/s12152-022-09481-3.

fit within existing legal privacy frameworks designed to govern the flow of personal information. For example, there already exist legal frameworks recognizing the right to privacy, such as the universal declaration of human rights (UDHR), the EU's General Data Protection Regulation (GDPR) (which aims at protecting individuals with regards to the processing and transfer of personal data), and in the US the Health Insurance Portability and Accountability Act (HIPPA).

Whether brain data represents a new and especially worrying threat to privacy cannot be determined simply by examining the nature of the data itself. Rather, we must understand who is using this data, how, and for what purposes. To better understand these issues, we argue that neuroethics discussions about mental and brain privacy can learn from an approach well-known in the information privacy world: Helen Nissenbaum's "contextual integrity" theory. In the next section, we provide an overview of contextual integrity and show why framing privacy concerns around brain and mental data through this lens is a more promising approach than those discussed above.

PRIVACY AS CONTEXTUAL INTEGRITY

While "control theories" of privacy remain dominant in US (and to some degree European) law and policy, the notion that information privacy amounts to individual control over personal information has come under sustained criticism by philosophers, legal scholars, and technology experts.[9] A major intervention in these debates is philosopher Helen Nissenbaum's theory of privacy as "contextual integrity." On Nissenbaum's account, privacy should not be understood primarily as a function of individual responsibility and control. Rather, context-dependent social norms govern what she terms the "appropriate flow of information," and it is just as much the responsibility of data collectors to respect these norms as it is the responsibility of individuals to ensure they are respected (Nissenbaum 2010).

Consider, for example, the norms inherent to the context of friendship. If one friend confides in another, sharing a sensitive secret, the interaction is governed by an easily intuited norm of confidentiality: clearly, it would be wrong for the confidant to disclose the information entrusted to her to a third party. This norm of confidentiality need not be made explicit—people learn it through the same process of socialization that guides other aspects of interpersonal life. The burden of responsibility for ensuring the information remains private does not fall primarily on the sharer. It is the information recipient's responsibility to abide by the relevant contextual norms—in this case, to keep it confidential.

Or consider the context of doctor-patient relationships. Patients must disclose all manner of sensitive personal information to their physicians in order to benefit from accurate, well-calibrated care—information about their bodies, sexual activity, habits, and so on. In some settings, such as mental healthcare, providers are bound by explicitly codified legal confidentiality rules. In other settings, informational norms are unspoken but nevertheless easily understood. For example, if in order to get a second opinion about a suspected diagnosis, one doctor shares with another doctor information revealed by a patient in confidence, the patient would be unlikely to object. However, if the same doctor revealed that information at a dinner party while gossiping with a friend, or—say—they sold it to advertisers, the patient would rightly be outraged.

Note that the theory of privacy as contextual integrity is philosophical—it attempts to model our conceptual and moral intuitions about privacy, to predict when new technologies that alter existing information flows are likely to cause moral disapproval or outrage (Nissenbaum 2010, 6–7). Of course, different people's intuitions do not always coincide. Contextual integrity cannot definitively resolve such disagreements, but it provides a language and conceptual structure for reasoning about and expressing them with precision. Thus, it is more an analytical than a normative theory: it aims to describe when and why new data-driven technologies provoke ethical or moral concern, in order to help guide law and policy (as well as social and technical) responses. We introduce this framework into discussions about privacy and brain data in order to help clarify the specific values and interests at stake in these debates. Viewed through the lens of contextual integrity, brain data is revealed to raise familiar privacy problems—the same problems, in fact, as other digital technologies—rather than new, unprecedented threats. As such, we argue that legal and policy responses to brain data's privacy challenges should be integrated with efforts to address related digital privacy threats, instead of being siloed into separate battles for new mental and brain privacy rights.

A central feature of Nissenbaum's theory is the analytical model she offers for specifying the relevant

[9]For a brief overview, see (Susser 2016). For a longer, critical discussion about the relationship between contextual integrity and control theories of privacy, see (Birnhack 2012).

norms governing the flow of information in particular social contexts. Every "context-relative informational norm" can be articulated in terms of five parameters: the information sender, its recipient, the subject (i.e., the person the information describes), the type (or "attributes") of information being transmitted, and the transmission principles governing its flow. Returning to one of the examples, above: when one friend confides in another, the confider is sender and subject (because the information is about the sender herself), the confidant is the recipient, the information type is sensitive personal information, and the transmission principle is that the information is bound by a norm of confidentiality—i.e., the information should not flow beyond the designated recipient, unless (perhaps) the sender has explicitly consented (Gupta 2013).

This last point emphasizes something worth noting: in some instances, an informational norm can approximate the control theory approach, as in the case above, where the norm of confidentiality might include exceptions if one obtains the information subject's consent (Gupta 2013). The theory of contextual integrity does not, therefore, suggest that individual control over personal information is never desirable, nor does it seek to repudiate control theories in their entirety. Rather, contextual integrity demonstrates that individual control over personal information is but one of many context-relative informational norms that govern information flows in everyday life.

New technologies—especially data-driven tools like many neurotechnologies—can alter information flows or introduce new ones. Nissenbaum describes a "decision heuristic" for determining whether such changes violate contextual integrity (Nissenbaum 2010, 148–150). First, articulate both the pre-existing informational norm and the new informational norm, by specifying the parameters discussed above. Second, compare. If the two are different, there is a prima facie violation—the new information flow has disrupted the pre-existing norm. Of course, change can be good and new informational norms might be preferable to old ones. Thus, once a prima facie violation has been flagged, one must ask if the new informational norm is harmful. Specifically, does it threaten the "values, goals, or ends" of the relevant context?

Consider, again, the doctor's office. As we described above, the flow of information between patient and physician is governed by long-standing privacy norms: one expects information disclosed to their physician will be held in confidence. Now imagine a doctor's office adopts some new technology for their practice—for example, a digital application that automates the patient intake process. Instead of filling out stacks of paper forms, when patients arrive, they simply enter their information into the app, which stores it in the cloud. Introducing this system disrupts context-relative informational norms—that is, it violates prevailing social expectations about how information ought to flow in that context—by sending the information to a new recipient. Prior to its adoption, the only recipient of patient information was the doctor's office; now, it is sent to a number of third-party intermediaries, including internet service providers, the cloud storage provider, and perhaps the company that designed the software.

This new information flow raises a red flag. Disclosing patient information to these unexpected recipients—even incidentally—constitutes a prima facie privacy violation. The question is, does the new information flow threaten the values, goals, or ends of the context? Which is to say, does disclosing patient information to third parties undermine the reason for going to the doctor's office in the first place? It is easy to imagine how it could: for doctors to care for patients successfully, to provide accurate diagnoses and tailored treatment plans, patients must be forthcoming. They must reveal sensitive personal information to their doctors that they may not have told to anyone else. An important reason why people are generally willing to do that is because they trust doctors. Learning that trust is misplaced, because information disclosed in confidence is being transmitted—unexpectedly—to digital intermediaries, could make patients less disposed to sharing medically necessary information in the future.

Thus, according to the theory of contextual integrity, bringing this new technology into the doctor's office raises privacy concerns not because it deprives patients of control over information about themselves (though it may do that). It raises privacy concerns because the technology causes information to flow in a new, unexpected way that could be detrimental to the context's purpose.

BRAIN DATA IN CONTEXT

What does all of this mean for mental and brain privacy? As we've seen, there are three broad worries about the implications of neurotechnologies for privacy: (1) the brain data these technologies generate is particularly revealing; (2) the decision to reveal or conceal that information is less subject to conscious control; and (3) there is considerable uncertainty about what could be inferred from the data, complicating decisions about whether to permit its

collection. There is little doubt that all of this makes controlling information more difficult. But contextual integrity asks us to consider different questions—namely, do neurotechnologies disrupt prevailing context-relative informational norms? And if so, do the new information flows run at cross-purposes to the values, ends, and goals of the contexts in which they appear? To explore these questions concretely, we situate them in three contexts where neurotechnologies are typically discussed: healthcare, marketing, and law enforcement.

Specifically, we explore what information is generated by the application of neurotechnologies in these contexts, who has access to it, and what it reveals. As we discussed above, calls for special attention to the privacy issues raised by neurotechnologies are motivated by the concern that data about the brain is particularly revealing about activity in the mind. Since the mind is often thought to be the seat of individual identity, personality, agency, and autonomy—the bases of moral personhood—giving others access to our minds raises deep moral questions, foremost about privacy. But it is crucial to remember that the relationship between brain data and mind data is indirect. fMRI, EEG, and related neurotechnologies collect information about ongoing activity in the brain (in the case of fMRI that "activity" is measured in relation to blood flow, in the case of EEG it is detected by measuring electrical signals) which can, in turn, be *correlated* with relevant mental phenomena (Hallinan et al. 2013). For example, it can be inferred from patterns in data about blood flow or electrical signals in the brain that a subject is likely imagining specific words, which can then be converted into text or computerized speech (Goering et al. 2021), or that they intend to move an arm, which can be used to control robotic prosthetics (Roelfsema, Denys, and Klink 2018).

Healthcare

Developers of neurotechnologies imagine a wide range of applications, from cognitive enhancement to lie detection to learning assessment, but their "major focus" has been healthcare (Eaton and Illes 2007). In clinical settings, neurotechnologies are being developed for both "assessment" and "intervention" (White et al. 2015). At present, using brain function images (such as fMRI scans) for diagnosis requires a significant degree of interpretation by clinicians. As neurotechnologies advance, the goal is to automate much of this process, lowering costs and making results more consistent (Eaton and Illes 2007). Meanwhile, non-invasive fMRI and EEG technology is driving the development of brain-computer interfaces (BCIs), which enable real-time feedback about activity in the brain (Lupu, Ungureanu, and Cimpanu 2019; White et al. 2015). BCI technology is also being used therapeutically, for treatment of "stroke recovery, paralysis, and degenerative conditions, such as amyotrophic lateral sclerosis," as well as socioemotional conditions, such as psychopathy, antisocial disorder, and schizophrenia (Bockbrader et al. 2018; Shih, Krusienski, and Wolpaw 2012; White et al. 2015, 798).

In healthcare contexts, clinicians have found neurotechnologies useful for a variety of purposes. Mapping activity in the brain to mental phenomena, such as "tactile, motor, language, and visual functions," can help neurosurgeons "assess surgical risk, plan surgical routes, and direct intraoperative electrophysiological procedures" (Tovino 2007, 423). It can be used to discover "atypical brain function" (Garden and Winickoff 2018, 12), like the way language is processed by patients with epilepsy (Jarosiewicz and Morrell 2021; Lo and Widge 2017; Tovino 2007, 423). It can screen for potential psychiatric conditions by pinpointing "biomarkers of mental illness" (Ienca, Haselager, and Emanuel 2018). Obviously, these inferences are highly sensitive—in the wrong hands, they could lead to stigma, discrimination, and distress. And like other sensitive information collected in healthcare settings, brain data is susceptible to both intentional and unintentional disclosure. A hospital, for example, might decide to sell patient information to data collectors, or the information could be leaked in a data breach (Ienca, Haselager, and Emanuel 2018).

Marketing

Next, consider proposed uses of neurotechnologies in consumer marketing. Just as healthcare practitioners see promise in neurotechnologies for both medical diagnosis and clinical intervention, marketers are enthusiastic about using the same kinds of devices—such as non-invasive BCIs—to understand and influence consumer decision-making. Research in "neuromarketing" has been conducted by well-known brands, such as Google, Disney, CBS, and Frito-Lay, exploring how neurotechnologies can be used to gauge consumers' subconscious reactions to new products and advertisements (Eaton and Illes 2007; Ienca and Andorno 2017; Murphy, Illes, and Reiner 2008). And neurotechnology companies are developing techniques that leverage this information about people's

preferences to "prime, imprint, or trigger" them (Ienca and Andorno 2017). With the rise of wearable, "pervasive" neurotechnologies—headsets and other devices incorporated into mobile technologies by companies like Apple and Samsung, and into entertainment and gaming platforms—advertisers have access to more and more brain data, and increasing opportunities to use it (Ienca and Andorno 2017; Penenberg 2011; Stanton, Sinnott-Armstrong, and Huettel 2017).

Like in healthcare contexts, neurotechnologies used for marketing purposes attempt to map brain phenomena to mental phenomena—they generate information about brain activity (blood flow or electrical signals), from which inferences are made about mental activity, such as consumer preferences (Ariely and Berns 2010).[10] As Ariely and Berns (2010) write, the "hope is that neuroimaging will reveal information about consumer preferences that is unobtainable through conventional methods" (284). Although preliminary studies offer some evidence in support of these inferences, they are far from conclusive. Indeed, what little research has been conducted deserves skepticism, given that "the great majority of the information is published by neuromarketing companies or academics who work in these enterprises" (Fortunato et al. 2014, 215). And while some worry that once BCIs are integrated into more consumer devices the data they generate could be collected surreptitiously, in situ, to date most neuromarketing research has been conducted in the lab, with subjects knowingly and willingly participating in the data collection process (Lim 2018).

Law Enforcement

Finally, worries about the privacy implications of neurotechnologies are perhaps most serious (if also most speculative) in the context of law enforcement, where some imagine that they "might possibly contribute to more evidence-based decisions in criminal justice, from investigation and the assessment of criminal responsibility, to punishment, rehabilitation of offenders, and the evaluation of their risk of recidivism" (Ienca and Andorno 2017, see also Ligthart et al. 2021). Thus far, the main application of neurotechnologies in criminal justice has been lie detection, with early studies showing promise over existing polygraph techniques (Eaton and Illes 2007; Holley 2009). Specifically, the developers of these tools claim that fMRI scans can detect subtle differences in blood oxygenation levels between people who are lying and people who are telling the truth (Eaton and Illes 2007), and they claim that "brain fingerprinting"—based on EEG scans—can infer from electrical activity in the brain whether someone recognizes a certain image (Holley 2009). Here again there is every reason for skepticism, but the results of fMRI-based lie detection tests have already been submitted as evidence in US and Indian courts (Holley 2009), and legal observers question whether existing laws can manage it (Farahany 2012b; Meegan 2008).

DISCUSSION

According to the theory of privacy as contextual integrity, the information flows described above certainly raise red flags. Neurotechnologies create a new type of personal information (or informational "attribute")—real-time information about blood flow or electrical signals in the brain. That alone prompts urgent questions about privacy. Given their novelty, we can't ask if these new information flows disrupt existing context-relative informational norms; we don't have norms that govern the circulation of brain data. Instead, we have to ask whether the new information flows are harmful, and if they are compatible with the values, ends, and purposes of the social contexts in which we find them. Framing the question this way is helpful, because it allows us to see that the novelty of brain data does not necessarily create novel privacy problems. To the contrary, new technologies enabling new kinds of information flows is a problem with which we are all too familiar.

In healthcare contexts, for example, privacy advocates have long voiced concerns about electronic health records (EHR) and the digitization of medical data more broadly, which have caused information historically confined to paper charts or private, closed computer systems to flow widely—to cloud storage systems and third-party data processors, academic and industry researchers, insurance companies, and others (Powles and Hodson 2017; Wetsman 2021). Viewed through the lens of contextual integrity, these technologies threaten privacy because they cause information to flow to new recipients, violating norms of doctor-patient confidentiality. Moreover, these data flows could harm data subjects if the information was

[10]Some argue that the term "neuromarketing" encompasses a more expansive set of techniques than this, including, amongst other things, the measurement of "physiological aspects such as perspiration, electrical conductivity of the skin, hormonal and neurotransmitter changes, movement and dilation of the pupil, movements of muscles (body and face), to even the understanding of complex cognitive aspects, such as the functional activity of specific regions of the brain by means of the analysis of different markers such as electrical waves, cerebral metabolism and its blood flow" (Fortunato et al. 2014).

leaked in a breach, if researchers used it for purposes to which the data subjects did not consent, or if insurance companies used it to justify raising premiums. As we saw in the previous section, information flows that violate norms of doctor-patient confidentiality can undermine the trust patients place in doctors. Less trusting patients are likely to be less forthcoming, making it more difficult for doctors to provide high quality care.

Likewise in marketing and law enforcement contexts. Much has been written about the rise of "surveillance capitalism"—a digital economy driven by targeted advertising (Zuboff 2019). Digital platforms generate information about every person's web browsing behavior, purchasing patterns, social media use, and any other activities that can be tracked by the computers, smart phones, wearable devices, and other sensors they interact with, and that data is fed through an ecosystem of data aggregators, analytics firms, marketing companies, and advertising platforms to personalize and target digital ads. Many of these data practices violate contextual integrity, not because they generate new kinds of information, but because information about activity in one context flows inappropriately to another. For example, information about personal finances might be used to target ads. Similarly, personal data collected by private platforms often finds its way to law enforcement, enabling new kinds and greater degrees of government surveillance, raising further concerns about data collected in one context being disclosed to recipients in another (Morrison 2021).

From this vantage point, the threats to privacy posed by neurotechnologies begin to look fairly commonplace. Like many of the other digital technologies transforming our lives, they transmit data about our bodies and behaviors to new, unexpected recipients who can use it for purposes that may or may not serve us. Indeed, while proponents of new mental and brain privacy rights tend to emphasize what is unique about neurotechnology, on closer inspection these technologies and the data flows they generate seem to threaten *the very same* harms as other digital tools.[11]

As we've seen, proponents of new mental and brain privacy rights justify them by claiming that (1) brain data is especially revealing about people's personalities and preferences, (2) it is especially resistant to conscious control, and (3) it is especially uncertain what could be inferred from the data. First, the same is true of the physical and behavioral data collected by other digital devices. Much like "neuromarketers," digital advertisers claim that the behavioral data they collect offers a special window into people's hidden preferences and desires. Which is to say, they too suggest reliable inferences can be made about individual preferences and desires on the basis of measurable data, only they focus on data about our outward activities and behavior rather than data about blood flow or electrical signals in the brain.

Second, because digital tracking has become so pervasive, and since much of our behavior is habitual and unconscious, it is equally difficult to consciously control what an observer might learn about us through physical surveillance as it is to control what they might learn from brain scans (Zuboff 2019). (Practically speaking, unless BCIs become as ubiquitous as smart watches it will be much *more* difficult to consciously conceal behavioral data than it is to conceal brain data.) For example, Facebook attracted criticism for developing predictive analytics tools that—according to leaked internal company documents—can infer from their activity on the platform when teenagers are feeling particularly vulnerable ("anxious," "overwhelmed") (Susser, Roessler, and Nissenbaum 2019). Third, as privacy and technology scholars have long argued, there is enormous uncertainty about what could be inferred from any particular piece of personal information once it is combined with other information and analyzed in the aggregate (Barocas and Nissenbaum 2009; Solove 2013).

Indeed, many of the privacy issues neuroethicists have highlighted are precisely the same as those raised in broader discussions about digital tracking. For example, Ienca, Haselager, and Emanuel (2018) worry about the risk of "brain leaks"—i.e., the "unauthorized disclosure of brain information" resulting from cyberattacks on hospital databases or third-party cloud storage systems—and about law enforcement demanding access to brain data held by private firms (808). Concerns expressed by Goering et al. (2021), that neurotechnologies could be used to manipulate people by "writing information into the brain," mirror broader worries about the use of digital technologies to manipulate people through digital nudges, targeted advertisements, and "dark patterns." (Goering et al. 2021, 8; Susser, Roessler, and Nissenbaum 2019). And the fear, generally, that current privacy and data protection laws haven't kept pace with technology and can't provide sufficient protection against the harms threatened by neurotechnologies is, of course, shared by many who focus on privacy and digital technology

[11]Kasper Lispert-Rasmussen offers further reasons for doubting that brain data raises unique privacy challenges, beyond those we put forward below. See (Lippert-Rasmussen 2017).

more broadly (Hallinan et al. 2013; Ienca, Haselager, and Emanuel 2018; Rainey et al. 2020).

To be clear, our aim is not to deflate worries about the privacy implications of brain data. Neurotechnologies are already creating new streams of sensitive information, which—like all health data—ought to be carefully protected. And if future neurotechnologies live up to the promises made about them the privacy risks they bring in tow will only increase. Rather, our aim has been to suggest that we can better conceptualize and understand those risks by applying the latest theories in information privacy, such as the theory of contextual integrity. Doing so, we can see that the privacy concerns neurotechnologies raise are urgent but not *novel*, justifying attention and action, but not new, unique rights. Indeed, there is reason to think that focusing on what little distinguishes the privacy harms threatened by neurotechnologies from those threated by other digital tools, instead of on what they share in common, could undermine efforts to defend privacy, rather than strengthen them.

CONCLUSION

Privacy is under threat everywhere. We are monitored and tracked, online and offline, by governments and private firms (Zuboff 2019). Information about where we go, what we do, and with whom is collected from our smartphones and sold to data brokers, advertisers, and the police (The Editorial Board 2019). Viewed through the lens of contextual integrity the situation is especially dire, as information gathered in one context is routinely transferred, analyzed, and put to use in others, often in violation of context-specific social norms designed to govern these information flows. Financial information finds its way into health research. Health information is used for commercial advertising. Location information guides the targeting of political messages.

How, then, do worries about mental and brain privacy fit into this larger picture? Proponents of special mental and brain privacy rights suggest that neurotechnologies, such as brain-computer interfaces, threaten unique new privacy harms; as such, there is a need for unique new rights to protect against them. As we hope to have demonstrated in this paper, these arguments are unfounded. Brain data may be new, but the privacy concerns brain data raises are all too familiar. Focusing attention on what distinguishes brain data from other types of personal information, rather than their commonalities, is unhelpful for two reasons—one theoretical, the other practical.

Theoretically, doing so leaves valuable conceptual and normative tools on the table. As we have seen through the example of contextual integrity, theories of information privacy—beyond mental and brain privacy—have advanced considerably in recent years, offering insights that can deepen and enrich related discussions in neuroethics. Privacy is deeply tied to social context: understanding whether, why, and to what extent brain data threatens privacy requires accounting for the actors involved in its transmission, the informational norms governing its flow, and the ends, values, and purposes of the contexts in which it is collected, analyzed, and put to use. The flow of brain data, to and from medical researchers, raises different concerns than its flow to and from advertisers or the police. Trying to understand brain data's privacy implications solely in terms of the kind of information it is, without attending to context, is to rely on an outmoded theory of privacy.

Practically, cleaving mental and brain privacy from information privacy more broadly divides privacy advocacy. Worries about mental and brain privacy ought to be framed as further reasons to shore up the privacy rights we ostensibly already have, and which desperately need defending, rather than reason to fashion new ones. Emphasizing the supposedly "special" nature of brain data is a kind of "neuroessentialism" (Reiner 2011). And calling for new rights does not come without cost. As Bublitz (2022) argues, drawing from critical discussions around human rights more broadly, adding to the list of rights can create "inflationary" pressure that devalues all of them: "If every important interest or legitimate concern became a matter of human rights, they may lose their distinction, significance, and effectiveness" (3).

Neurotechnologies and the brain data they proliferate threaten real privacy harms. But what is the nature of this harm, and how do we defend against it? To answer these questions, we ought not to fetishize the brain, focusing on what makes brain data different from the myriad other forms of personal information digital technologies circulate—data collected from our smartphones, genetic data, and the like. Neuroethics should learn from these related efforts and advocates for mental and brain privacy should join with them in common cause.

ACKNOWLEDGEMENTS

We thank participants in the Penn State Bioethics Colloquium, the Law & Technology Workshop at the Tel Aviv University Buchmann Faculty of Law, and the 4th

Annual Symposium on Applications of Contextual Integrity at Cornell Tech for their engagement with the arguments and ideas presented in earlier versions of the paper.

FUNDING

The author(s) reported there is no funding associated with the work featured in this article.

ORCID

Laura Y. Cabrera http://orcid.org/0000-0002-6220-7096

REFERENCES

Ariely, D., and G. S. Berns. 2010. Neuromarketing: The hope and hype of neuroimaging in business. *Nature Reviews Neuroscience* 11 (4):284–92. doi:10.1038/nrn2795.

Barocas, S., and H. Nissenbaum. 2009. On notice: The trouble with notice and consent. In Proceedings of the Engaging Data Forum: The First International Forum on the Application and Management of Personal Electronic Information, 1–7.

Birnhack, M. D. 2012. A quest for a theory of privacy: Context and control. *Jurimetrics* 51 (4):447–79.

Bockbrader, M. A., G. Francisco, R. Lee, J. Olson, R. Solinsky, and M. L. Boninger. 2018. Brain computer interfaces in rehabilitation medicine. *Physical Medicine &Rehabilitation* 10 (9S2):S233–S243. doi:10.1016/j.pmrj.2018.05.028.

Boire, R. G. 2001. On cognitive liberty. *Journal of Cognitive Liberties* 2:7–22.

Bublitz, J.-C. 2013. My mind is mine!? Cognitive liberty as a legal concept. In *Cognitive Enhancement*, ed. H. Franke, vol. 1, 233–264. Berlin: Springer.

Bublitz, J. C. 2022. Novel neurorights: From nonsense to substance. *Neuroethics* 15 (1):7. doi:10.1007/s12152-022-09481-3.

Carraro, V. 2019. Promoting compliance with human rights: The performance of the United Nations' universal periodic review and treaty bodies. *International Studies Quarterly* 63 (4):1079–93. doi:10.1093/isq/sqz078.

Clark, A., and D. Chalmers. 1998. The extended mind. *Analysis* 58 (1):7–19. https://www.jstor.org/stable/3328150.

Eaton, M. L., and J. Illes. 2007. Commercializing cognitive neurotechnology—the ethical terrain. *Nature Biotechnology* 25 (4):393–7. doi:10.1038/nbt0407-393.

Engstrom, P. 2017. Human rights: Effectiveness of international and regional mechanisms. *Oxford Research Encyclopedia of International Studies*. 22; Accessed 10 Mar. 2023. doi:10.1093/acrefore/9780190846626.013.214.

Farahany, N. A. 2012a. Incriminating thoughts. *Stanford Law Review* 64:351.

Farah, M. J. 2005. Neuroethics: The practical and the philosophical. *Trends in Cognitive Sciences* 9 (1):34–40. doi:10.1016/j.tics.2004.12.001.

Farahany, N. A. 2012b. Searching Secrets. *University of Pennsylvania Law Review* 160:70.

Farah, M. J., and P. R. Wolpe. 2004. Monitoring and manipulating brain function: New neuroscience technologies and their ethical implications. *The Hastings Center Report* 34 (3):35–45. doi:10.2307/3528418.

Fins, J. J. 2022. The unintended consequences of Chile's neurorights constitutional reform: Moving beyond negative rights to capabilities. *Neuroethics* 15(3): 1–11. doi:10.1007/s12152-022-09504-z

Fortunato, V. C. R., J. De Moura, E. Giraldi, and J. Henrique Caldeira De Oliveira. 2014. A review of studies on neuromarketing: Practical results, techniques, contributions and limitations. *Journal of Management Research* 6 (2):201. doi:10.5296/jmr.v6i2.5446.

Garden, H., and D. Winickoff. 2018. Issues in neurotechnology governance. *OECD Science, Technology and Industry Working Papers* 2018/11. Vol. 2018/11. doi:10.1787/c3256cc6-en.

Goering, S., E. Klein, L. S. Sullivan, A. Wexler, B. Agüera y Arcas, G. Bi, J. M. Carmena, J. J. Fins, P. Friesen, J. Gallant, et al. 2021. Recommendations for responsible development and application of neurotechnologies. *Neuroethics* 14 (3): 365–86. doi:10.1007/s12152-021-09468-6.

Goering, S., and R. Yuste. 2016. On the necessity of ethical guidelines for novel neurotechnologies. *Cell* 167 (4):882–5. doi:10.1016/j.cell.2016.10.029.

Greely, H. T., C. Grady, K. M. Ramos, W. Chiong, J. Eberwine, N. A. Farahany, L. S. M. Johnson, B. T. Hyman, S. E. Hyman, K. S. Rommelfanger, et al. 2018. Neuroethics guiding principles for the NIH BRAIN initiative. *The Journal of Neuroscience* 38 (50):10586–8. doi:10.1523/JNEUROSCI.2077-18.2018.

Gupta, U. C. 2013. Informed consent in clinical research: Revisiting few concepts and areas. *Perspectives in Clinical Research* 4 (1):26–32. doi:10.4103/2229-3485.106373.

Hallinan, D., P. Schütz, M. Friedewald, and P. D. Hert. 2013. Neurodata and neuroprivacy: Data protection outdated? *Surveillance & Society* 12 (1):55–72. doi:10.24908/ss.v12i1.4500.

Holley, B. 2009. It's all in your head: Neurotechnological lie detection and the fourth and fifth amendments. *Developments in Mental Health Law* 28 (1):25.

Ienca, M. 2021b. On neurorights. *Frontiers in Human Neuroscience* 15:701258. doi:10.3389/fnhum.2021.701258.

Ienca, M., and R. Andorno. 2017. Towards new human rights in the age of neuroscience and neurotechnology. *Life Sciences, Society and Policy* 13 (1):5. doi:10.1186/s40504-017-0050-1.

Ienca, M., J. J. Fins, R. J. Jox, F. Jotterand, S. Voeneky, R. Andorno, T. Ball, C. Castelluccia, R. Chavarriaga, H. Chneiweiss, et al. 2022. Towards a governance framework for brain data. *Neuroethics* 15 (2):1–14. doi:10.1007/s12152-022-09498-8.

Ienca, M., P. Haselager, and E. J. Emanuel. 2018. Brain leaks and consumer neurotechnology. *Nature Biotechnology* 36 (9):805–10. doi:10.1038/nbt.4240.

Ienca, M. 2021a. Common human rights challenges raised by different applications of neurotechnologies in the biomedical field. *Report commissioned by the Council of Europe*. https://rm.coe.int/report-final-en/1680a429f3.

Jarosiewicz, B., and M. Morrell. 2021. The RNS system: Brain-responsive neurostimulation for the treatment of epilepsy. *Expert Review of Medical Devices* 18 (2):129–38. doi:10.1080/17434440.2019.1683445.

Koh, H. 1998. How is international human rights law enforced? *Indiana Law Journal* 74 (3):1397–417. https://www.repository.law.indiana.edu/cgi/viewcontent.cgi?article=2279&context=ilj.

Ligthart, S., T. Douglas, C. Bublitz, T. Kooijmans, and G. Meynen. 2021. Forensic brain-reading and mental privacy in European Human Rights Law: Foundations and challenges. *Neuroethics* 14 (2):191–203. doi:10.1007/s12152-020-09438-4.

Lim, W. M. 2018. Demystifying neuromarketing. *Journal of Business Research* 91:205–20. doi:10.1016/j.jbusres.2018.05.036.

Lippert-Rasmussen, K. 2017. Brain privacy, intimacy, and authenticity: Why a complete lack of the former might undermine neither of the latter! *Res Publica* 23 (2):227–44. doi:10.1007/s11158-016-9344-z.

Lo, M.-C., and A. S. Widge. 2017. Closed-loop neuromodulation systems: Next-generation treatments for psychiatric illness. *International Review of Psychiatry* 29 (2):191–204. doi:10.1080/09540261.2017.1282438.

Lupu, R. G., F. Ungureanu, and C. Cimpanu. 2019. Brain-computer interface: Challenges and research perspectives. In *2019 22nd International Conference on Control Systems and Computer Science (CSCS)*, 387–394. Bucharest, Romania: IEEE. doi:10.1109/CSCS.2019.00071.

Meegan, D. V. 2008. Neuroimaging techniques for memory detection: Scientific, ethical, and legal issues. *The American Journal of Bioethics* 8 (1):9–20. doi:10.1080/15265160701842007.

Morrison, S. 2021. Here's how police can get your data—even if you aren't suspected of a crime. *Vox*, July 31, 2021. https://www.vox.com/recode/22565926/police-law-enforcement-data-warrant.

Murphy, E. R., J. Illes, and P. B. Reiner. 2008. Neuroethics of neuromarketing. *Journal of Consumer Behaviour* 7 (4–5):293–302. doi:10.1002/cb.252.

Nissenbaum, H. 2010. *Privacy in context: Technology, policy, and the integrity of social life*. Stanford, California: Stanford Law Books.

Penenberg, A. 2011. NeuroFocus uses neuromarketing to hack your brain. *Fast Company*, August 8, 2011. https://www.fastcompany.com/1769238/neurofocus-uses-neuromarketing-hack-your-brain

Powles, J., and H. Hodson. 2017. Google deepmind and healthcare in an age of algorithms. *Health and Technology* 7 (4):351–67. doi:10.1007/s12553-017-0179-1.

Rainey, S., K. McGillivray, S. Akintoye, T. Fothergill, C. Bublitz, and B. Stahl. 2020. Is the European data protection regulation sufficient to deal with emerging data concerns relating to neurotechnology? *Journal of Law and the Biosciences* 7 (1):lsaa051. doi:10.1093/jlb/lsaa051.

Reiner, P. B. 2011. The rise of neuroessentialism. In *The Oxford Handbook of Neuroethics*, eds. J. Illes and B. J. Sahakian. Oxford, UK: Oxford University Press.

Roelfsema, P. R., D. Denys, and P. C. Klink. 2018. Mind reading and writing: The future of neurotechnology. *Trends in Cognitive Sciences* 22 (7):598–610. doi:10.1016/j.tics.2018.04.001.

Roskies, A. 2002. Neuroethics for the new millennium. *Neuron* (1)35:21–3. doi:10.1016/s0896-6273(02)00763-88.

Ryberg, J. 2017. Neuroethics and brain privacy: Setting the stage. *Res Publica* 23 (2):153–8. doi:10.1007/s11158-016-9340-3.

Schönau, A., I. Dasgupta, T. Brown, E. Versalovic, E. Klein, and S. Goering. 2021. Mapping the dimensions of agency. *AJOB Neuroscience* 12 (2–3):172–86. doi:10.1080/21507740.2021.1896599.

Sententia, W. 2004. Neuroethical considerations: Cognitive liberty and converging technologies for improving human cognition. *Annals of the New York Academy of Sciences* 1013:221–8. doi:10.1196/annals.1305.014.

Shih, J. J., D. J. Krusienski, and J. R. Wolpaw. 2012. Brain-computer interfaces in medicine. *Mayo Clinic Proceedings* 87 (3):268–79. doi:10.1016/j.mayocp.2011.12.008.

Solove, D. J. 2013. Privacy self-management and the consent dilemma. *Harvard Law Review* 126:1880–903.

Stanton, S. J., W. Sinnott-Armstrong, and S. A. Huettel. 2017. Neuromarketing: Ethical implications of its use and potential misuse. *Journal of Business Ethics* 144 (4):799–811. doi:10.1007/s10551-016-3059-0.

Susser, D. 2016. Information privacy and social self-authorship. *In Techné: Research in Philosophy and Technology* 20 (3):216–39. doi:10.5840/techne201671548.

Susser, D., B. Roessler, and H. Nissenbaum. 2019. Online manipulation: Hidden influences in a digital world. *Georgetown Law Technology Review* 4:1–45.

The Editorial Board. 2019. Total surveillance is not what America signed up for. *The New York Times*, December 22, 2019. https://www.nytimes.com/interactive/2019/12/21/opinion/location-data-privacy-rights.html.

Tovino, S. A. 2007. Functional neuroimaging information: A case for neuro exceptionalism? Scholarly Works. 76. https://scholars.law.unlv.edu/facpub/76

Westin, A. F. 1968. Privacy and freedom. *Washington and Lee Law Review* 25 (1):166–70. https://scholarlycommons.law.wlu.edu/wlulr/vol25/iss1/20.

Wetsman, N. 2021. Hospitals are selling treasure troves of medical data—what could go wrong? *The Verge*. June 23, 2021. https://www.theverge.com/2021/6/23/22547397/medical-records-health-data-hospitals-research.

White, S. W., J. A. Richey, D. Gracanin, M. A. Bell, S. LaConte, M. Coffman, A. Trubanova, and I. Kim. 2015. The promise of neurotechnology in clinical translational science. *Clinical Psychological Science* 3 (5):797–815. doi:10.1177/2167702614549801.

Yuste, R., S. Goering, B. Agüera y Arcas, G. Bi, J. M. Carmena, A. Carter, J. J. Fins, P. Friesen, J. Gallant, J. E. Huggins, et al. 2017. Four ethical priorities for neurotechnologies and AI. *Nature* 551 (7679):159–63. doi:10.1038/551159a.

Zuboff, S. 2019. *The age of surveillance capitalism: The fight for a human future at the new frontier of power*. First edition. New York: Public Affairs.