# Brain Computer Interfaces: Privacy, Ethics, and Policy

Katherine Pratt

A dissertation
submitted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy

University of Washington

2019

Reading Committee:

Howard Chizeck, Chair

Ryan Calo

Chet Moritz

Program Authorized to Offer Degree:
Electrical and Computer Engineering

University of Washington

**Abstract**

Brain Computer Interfaces:
Privacy, Ethics, and Policy

Katherine Pratt

Chair of the Supervisory Committee:
Professor Howard Chizeck
Electrical and Computer Engineering

The ability to record and access neural signals is expanding, primarily driven by commercial entities developing non-invasive headsets that can control digital devices using a person's neural electrical signals. This is an excitingly novel modality for interacting with technology that will allow companies to profit from the sale of these devices. But hardware isn't the only thing of value in this exchange: companies also get access to the neural information of users. While some may choose to only use these signals for control purposes, this represents a new way that information can be elicited from someone, knowingly or unknowingly. Some may think nothing of revealing their preferred cereal or the kind of car they drive, because data brokers already collect hundreds of data points about us and interpolate many more. But what about the things that you wouldn't want to share, like a credit card number, whether you committed a crime, or perhaps even your sexual orientation? This is just one example of the many questions about privacy and perceptions of neurally-derived information that are becoming apparently for this emerging technology.

These recordings of the electrical activity in the brain fall into an unregulated area of law: while they are medically-based, they are obtained for non-medical purposes and thus do not fall under the jurisdiction of existing legislation to protect health information (HIPAA, Health Insurance Portability and Accountability Act). A general absence of federal data

privacy legislation and overall limited enforcement mechanisms on behalf of the Federal Trade Commission (FTC) means that new sources of consumer information are at risk of being exploited like other kinds of data are currently (e.g. from social media). Companies may claim that they are not selling the data they have on you, or that their practices are clearly outlined in the terms of service you accepted, but this only belies how opaque this particularly industry has become; consumers cannot say what happens to their information, but they inherently believe that it is only being used for the things they think it is (not what is actually happening to it). While lawmakers act reactively to address these kinds of known issues, it is possible to be proactive in the development and regulation of nascent technologies with commercial potential like neurally-controlled devices.

This dissertation describes a multi-disciplinary research approach to addressing the issues that arise from emerging technologies that rely on access to neural signals. It begins with a literature review of the techniques used to elicit information, as well as prior work. Then, experimental results will be presented that demonstrate the kinds of information that can come from real-time analysis that does not rely on prior classifier data. The following chapter discusses the philosophical foundation of neural privacy and presents the results of a neuroethics survey that demonstrates how individuals perceive the privacy of their neural information. Finally, the existing regulatory landscape will be discussed with respect to how neural information can be better regulated; this is done based on responses to policy-based questions from the neuroethics survey.

# TABLE OF CONTENTS

# LIST OF FIGURES

# LIST OF TABLES

# GLOSSARY

ACM: Association for Computing Machinery

ALS: Amyotrophic Lateral Sclerosi

BCI: Brain-Computer Interface

CNT: Center for Neurotechnology[1]

DBS: Deep Brain Stimulator

EEG: Electroencephalography

EMG: Electromyography

ERP: Event Related Potential

FMRI: Functional Magnetic Resonance Imaging

FTC: Federal Trade Commission

GDPR: General Data Protection Regulation

HIPAA: Health Insurance Portability and Accountability Act

IEEE: Institute of Electrical and Electronics Engineers

IRB: Institutional Review Board

PI: Perceived Information

UW: University of Washington

[1]This is an NSF Engineering Research Center based at the University of Washington that until September 2018 was known as the Center for Sensorimotor Neural Engineering, or CSNE.

# ACKNOWLEDGMENTS

It feels cliché to start with my K-12 education, but I was incredibly lucky to have teachers that inspired a life-long love of learning: Meg Scanlon and Jepson Lonquist (1st and 2nd), Dennis Frates (4th and 5th), Ken Gouveia (7th grade humanities), Carole Beedlow (AP Biology), Pat Canan (Physics and AP Physics), and Jim Dort (government). Mr Dort in particular gave me the best advice when I was trying to decide between MIT and the Air Force Academy: you can always serve in the military, but you can only go to MIT once. Thanks, Mr Dort.

Also while I was in high school I helped found a FIRST robotics team, and I competed with them for three years. My senior year, one of the judges gave me his business card and said to look up the company Blue Origin in a few years for an internship. Two internships and a job later, Mark Hofer was the officiant at my wedding and he and his wife Heather Leach are some of the kindest and most wonderful people I know. Thank you for being such incredible friends.

When I left the military, I knew that I needed some academic experience before I could apply to grad school. Thank you Dr Adrian KC Lee, for taking a chance on an aerospace engineer to help run your lab. I learned everything I know now about coding and experimental practices from you and Dr Ross Maddox, Dr Eric Larson, Dr Mark Wronkiewicz, and soon-to-be Dr Lindsey Kishline.

When I was rejected from graduate school the first time, I almost gave up. But then I talked to Dr Kristi Morgansen, who helped create the CNT post-bac program that got my foot in the door to research. I appreciate the time that you took to talk with me, and that you developed a program that continues to provide experience for

students looking to transition to graduate school. In particular, the CNT has been a second home for me, providing space to work and volunteer opportunities. I'm grateful that you've been a funding and research support mechanism for myself and many other graduate students who are a part of the Center. Through the Center I've been able to collaborate with students in all sorts of departments. In particular, thank you to Dr Nile Wilson, Dr Jenny Cronin, Dr James Wu, Dr Dev Sarma, and Dr David Caldwell, who are fantastic people who always seemed to know the answer when I ran down from the 4th floor with a question.

For the post-bac, all I needed was a desk in a lab; my mom says that I was like a foster child who needed a home, which was eventually provided by Dr Howard Chizeck. Thank you for the chance to join the BioRobotics Lab, and for becoming my advisor. I'm so excited about the research that I've been able to do, and I wouldn't have been able to do it without your support. In the lab, Dr Maggie Thompson was a fantastic labmates and I'm so proud that I got to work and study with you. Dr Jeffrey Herron was (and still is) a great resource for all things electronic and coding.

In the course of my research, I've had the pleasure of working with many disciplines across campus. In particular, I want to thank Dr Sara Goering and soon-to-be Dr Tim Brown from the Philosophy department, as well as Dr Laura Specker-Sullivan (previously a post-doc in the department). Sara has been an invaluable resource as I've written my document, and I appreciate all of the feedback she's been able to give during the process. Tim is an amazing collaborator and I'm so glad we got to work together on my questionnaire and a few conference papers. Dr Eran Klein from OHSU also advised on the neuroethics survey, and it was helpful to compare notes on responses to our neuroethics questionnaires. I'm also indebted to Marlena Bannick from the UW Statistics program, for her assistance in analyzing my neuroethics data as a part of the UW Statistical Consulting Service.

In the Law School, the Tech Policy Lab (TPL) has been an amazing space for collaboration, discussion, and education. Thank you to Emily McReynolds, Hannah Almeter, and especially TPL co-director and committee member Ryan Calo. I use your leadership and multidisciplinary vision as examples of how research should be done around the country.

My work in ethics and philosophy has allowed me to do some amazing things. I want to thank the ACLU of Washington and the ACLU Speech, Privacy and Technology Project for providing a summer internship where I got to work hands-on with technology policy. In particular, my boss Shankar Narayan has become a fantastic friend and mentor in the policy space. I'd also like to thank New America, the Open Technology Institute, Travis Moore, and Andres Bascumbe, for accepting me to as the first TechCongess Congressional Innovation Scholar. Thank you to Congresswoman Suzan DelBene, Lauren Soltani, and the whole DelBene team for an amazing fellowship placement.

I'd like to especially thank my whole committee for supporting my research and multidisciplinary vision: Dr Howard Chizeck, Dr Chet Moritz, Ryan Calo, JD, Dr Sara Goering, Dr Tamara Bonaci, Dr Yoshi Kohno, and Dr Sarah Tuttle.

For funding, I'd like to thank the NSF GRFP, Tech Policy Lab, UW ECE Irene Peden Endowed Fellowship, CNT, and the BioRobotics Lab.

For non-academically related funding, I want to thank all of my friends who provided emotional support and nourishment during the process. I don't have the room to list all of you individually, and I'd be mortified to find out that I forgot someone. But, I do want to give a special shoutout to Jocelyn Scheintaub and Mike O'Malley for letting me park in their driveway and the semi-frequent dog daysitting they provided, and to occasional carpool buddy and fairy godmother Emily Lloyd.

And finally I'd like to thank my whole family for supporting me on this wild

ride, especially my parents Bob and Columba Ingle who taught me to be curious, adventurous, and tenacious. The biggest of thanks to my husband Kevin; a third of his name is on my diploma, so it was only fitting that you helped me do a third of the work. I love you to the moon and back.

# DEDICATION

This dissertation is dedicated to everyone who has failed and stood up again, who was told that they couldn't or shouldn't, and who proved everyone wrong. We have fought so hard to get here, and I see you.

# Chapter 1
# RESEARCH MOTIVATIONS

This dissertation is a multidisciplinary look at the risk of privacy violations from emerging technologies that involve recording and processing neural signals. While it is important to quantify the current and future potential risk to consumers, there are equally important ethical and regulatory considerations that should be considered concurrently. This chapter first provides justification for why privacy is the appropriate lens upon which to base this work, the type of work accomplished, and scope of limitations.

## *1.1 The Current Privayc Landscape*

The concept of privacy has always been fluid, dependent on culture and convention. In 1890, Warren and Brandeis declared a "right to privacy" as a direct response to the proliferation of newspapers and photography [68] (see beginning of Ch 4 for more about their work). Since then, US court rulings have established privacy rights by referring to precedent in the Bill of Rights (e.g. the 4th Amendment right to privacy from *Katz v United States* [57]), but these kind of interpretations do not apply in the non-government/commercial space.[1] Here, it is up to entitites like the Federal Trade Commission (FTC) and their authority to pursue unfair and deceptive practices that cause "substantial injury" [18]. Unfortunately, understaffing and the inability to leverage sufficient monetary penalties means not all companies or actions are appropriately addressed.

So then what place does privacy, and in particular neural privacy, have in the 21s century and our current regulatory enviornment? It would seem that according to semi-regular articles in popular press, privacy is dead and we should accept our new surveillance world. This

[1] Companies can be liable if they violate statutory law, like *Rosenbach v Six Flags*, discussed in Sect 5.2.

statement is an oversimplification of a much more complex relationship between consumer and company that can be best described by the privacy paradox: an individual may profess particular privacy beliefs, but they behave differently in the real world [56]. It could be that humans are willing to compromise if the deal is right. From an economics perspective we are constantly calculating the optimal balance between privacy and anticipated gains [31], and short-term benefits may outweigh the long-term harms of a privacy-based decision [5].

If consumers are not reliable in making their own privacy choices, is it appropriate that we believe companies are acting in good faith on behalf? They claim that collecting information about us is of benefit to both parties: targeted ads and personalized user profiles increase revenue and customer satisfaction because they are able to display the right products more quickly, with one study claiming that 71% of consumers prefer tailored ads [42]. Unfortunately, consumers don't always know their information is being collected (e.g. cookies on a website), who is collecting it (company vs third party), or why it's being collected (keep track of purchase preferences, sell to data aggregators, etc). Many also don't know that the information they give out can be interpolated to infer other things about them. When particularly egregious examples of this practice are made public, like in 2018 with the news of the relationship between Facebook and Cambridge Analytica [26], consumers response tends to be overwhelmingly negative.

In light of this kind of attention, one would think legislators would be eager to pass some kind of comprehensive consumer data privacy legislation. The European Union was more proactive than the US and implemented the General Data Protection Regulation (GDPR) in May of 2018 [3] as a replacement to the 1995 Data Protection Directive [1]. However, despite numerous congressional hearings and opinion pieces calling for regulation or punitive punishments, regulations have not changed and legislation has not emerged to clarify what companies can do with customer data or to regulate the industry at large.

It is now more important than ever to discuss the ethcial and policy implications of technologies still being developed, in an attempt to learn from the previous mistakes of disregarding data privacy in the name of innovation. Brain-Computer Interfaces (BCIs) are

more established in medicine and research, but they are becoming commercialized into non-medical settings. If consumers were upset with what Facebook did with their meaningless quiz data, imagine what a tech giant can do with all of your raw neural signals.

But is there any reason to be particularly concerned about neural signals, given all the ways we express ourselves – intentionally and unintentionally – through our bodies? Each person has control over their ability to overtly share information or mental states through things like spoken/written word, choice of clothing, travel, etc. But we may also voluntarily choose to not express affection for a crush or true sexual orientation, or involuntarily suppress information like the aftermath of a prior trauma. These choices are our own variant of expressing or concealing neural information. A BCI device should not violate this control.

Thus, I make the following assertions:

- BCIs have the ability to violate a person's agency because elictation and analysis of neural signals bypasses a person's control over what they share with the outside world. This includes conscious and unconscious actions; even unconscious actions (such as a physical tell in a poker game) reveal information about the person (even if it's simply that you are witholding further information).

- Individuals should be able to exercise the same control over their personal information during their use of a BCI as they do with other aspects of their lives.

- There is a privacy violation when the signals are obtained or analyzed for purposes not intended by the individual, because they no longer maintain the ability to mediate their own information. Recordings of neural signals should only be obtained and maintained for the purposes of completing the BCI-related task.

With these limits defined, it is possible to frame research questions like determining the methods that can be used to violate them (Ch 3), or rules that can be implemented to protect the consumer (Ch 5). These are discussed next.

### *1.2 Research Statement*

This research is dividied into three main subject areas:

*1. Examine the kind of risk associated with elicitation of private information using a BCI.*

- Development of an human subject experiment that elicits neural signals in response to subject-selected digits and attempts to calculate the subject's chosen digit without any prior training data.

- Analysis of experimental results, particularly for the accuracy of this elicitation method as would be used in a side-channel attack strategy.[2]

*2. Establish if consumers, mobility-impaired and non mobility-impaired, have preconceived beliefs of privacy as related to information derived from neural signals*

- Discussion of existing literature on privacy to establish the factors relevant to the term neural privacy

- Presentation of survey results that evaluates how individuals perceive their neural information

*3. Discuss the shortfalls of existing legislative and regulatory authority as it relates to neural signals and derived information, along with potential solutions.*

- Analysis of topics with existing legislation and regulation that are relevant to the discussion about neural security

- Presentation of survey results that demonstrate which entities should be involved at each stage of of the BCI development lifecycle.

---

[2]Definition and further information provided in Ch 3

### *1.3 Limitation of Scope*

This is a broad area of research that crosses a multitude of disciplines. To refine the scope for this dissertation, several assumptions and limitations were defined.

*Commercial*

This dissertation will focus on BCI-devices that can be purchased without a prescription and for non-medical uses. This market already exists, and "... current spending on neuro-technology by for-profit industry is already US$100 million per year, and growing fast" [70]. This market share will only increase as BCIs are applied to more use cases, and as the signal to noise ratio and accuracy increase with improved signal processing methods and electrode design.

As will be discussed later in Ch 5, one of the few data privacy protections to exist in the US covers health care information: HIPAA (Health Insurance Portability and Accountability Act). Ostensibly this means that the information collected by medically necessary BCIs will be subject to this legislation. Additionally, experiments that are carried out by researchers through an academic institution are required to comply with human subject review board processes that approve experimental protocols. This includes deidentification of collected data and duration that the information is kept.

*Non-Invasive*

Ch 3 will discuss further the differences between invasive (surgical) and non-invasive neural recording. However, since there are no commercially-available implanted electrode systems, this dissertation will focus on devices which can be purchased without the need for surgery or a pre-existing medical condition.

*US-legal and regulatory framework*

Data privacy is a construct that is uniquely defined in every country, particularly whether or not there is an established or constitutional right to privacy. While large markets in the European Union and Asia may help drive consumption of BCI devices, this dissertation is using a US-based legal framework to discuss understandings of privacy, as well as potential regulatory responses.

# Chapter 2

# NEURAL ENGINEERING BY DESIGN

Technology is becoming pervasive in ways that are impacting our lives. Everything from autonomous cars to individual gene sequencing gives rise to questions of who has access, what is it used for, and what are the consequences. But who is responsible for examining these impacts and guiding development to respond how a novel device will impact a user's privacy? In 2010 Kenneth Bamberger and Deirdre Mulligan began to look at how privcy is perceived and understood at organizations by chief privacy officers and their equivalents [9][10]. This work was continued later in the decade by Ari Ezra Waldman, who interviewed executives like chief privacy officers as well as those building and develooping the technology [66]. It is apparent from this literature that there is a disconnect between what a company says it is doing with respect to privacy and the engineers below them. Waldman in particular found ambivalence and even hostility towards those who were brought in to advise or work with teams on privacy, because concepts like ethics and privacy are not part of core engineering curriculums and thus perceived as unimportant to the overall design. Rather than see these individuals as a threat, engineers must not consider these inputs as negative critiques of their work. It is long past due that in their work, engineers consider the unique perspectives of the humanities and law.

There are countless examples in the popular press about how the faults of a product are only revealed after it has been released to the public. When the augmented reality game Pokemon Go was released, game players were entering private property as part of the game and playing it in places that might be considered inappropriate, like the National Holocaust Museum [44]. Earlier this year the makers of the game proposed a settlement and changes to the game to protect private property and remind game players to be respectful of others

[44]. Would the initial release have been different if someone on the team had asked about property law? Hopefully the answer to this is yes.

The situation described above is emblematic of how reactive industry and lawmakers have become to the issue of privacy. The relatively unregulated (compared to elsewhere like the EU) area of data privacy has meant that action is only taken to address problems as they come to light. There is a chance with emerging technologies to be proactive, and this dissertation is an example of how to do this with neural technologies. The quantitative findings in Ch 3 are important to the study of neural elicitation, but they are incomplete without understanding the context in which the technology will be used. That is why the neuroethics and policy survey are important: how will future users perceive loss of neural information, and what should be done about it.

It is important to note here that the aim of the experiments in Ch 3 is not to document a novel neural signal detection method; there is ample existing literature on analysis techniques. Instead, the goal is to look at what kinds of information can be obtained from neural information. This is from the perspective of an advertiser or malicious entity who knows that this methodology exists: what can they learn using it? For those who choose to build upon this research, it is imperative to understand the consequences of your findings. The results in this dissertation are provided with discussions of the current landscape of ethics and policy to demonstrate the kinds of conversations that are necessary. As the technology progresses and further uses are discovered, researchers need to be cognizant of who can use this research and for what purposes.

# Chapter 3

# NAÏVE ELICITATION OF INFORMATION

The ability to record electricty from the brain was first published in 1924 by Hans Berger, who has come to be known as the father of electroencephalography (EEG).[1] Originally most useful as a tool for medicine and research, Jacques Vidal demonstrated how such electrial activity could control a computer, coining the term Brain-Computer Interface (BCI) [65]. Since then the breadth of research possibilities has exploded, with Google Scholar returning approximately 1,710,000 academic articles related to the search term "Brain Computer Interface."[2]

Interest in non-invasive EEG has grown because of its commercial potential, which will be discussed in depth in Sect 3.2. These systems are more user-friendly than other brain imaging techniques like fMRI (Functional Magnetic Resonance Imaging), which is expensive, immovable, and requires very compliant subject. And, the absence of surgery lowers the barrier to usage. The EEG signals themselves are temporally precise, but because they must pass through many layers (brain tissue, cerebrospinal fluid, skull, scalp, hair, and air) from source to electrode the signals are diffused — we know when they happen, but we aren't as good at knowing exactly where they came from. This time-related property is one aspect of EEG that makes it conducive for BCI tasks, and one class of these signals is used in this research. It is described next.

---

[1] Berger is also now known to have been complicit and involved with many Nazi-related projects [71].

[2] As of May 20, 2019.

### *3.1 Signal of Interest: the P300*

In 1964 and 1965, two papers were published outlining a type of brain signal that appeared in response to rarely-occurring stimuli that appear in a sequence of unrelated stimuli [16][58]. It is now known that there are a series of different types of these "event related potentials" (ERPs) that can be seen in response to cues such as incorrect grammar, unexpected syntax, and mistakes. They are visually depicted in Fig 3.1.

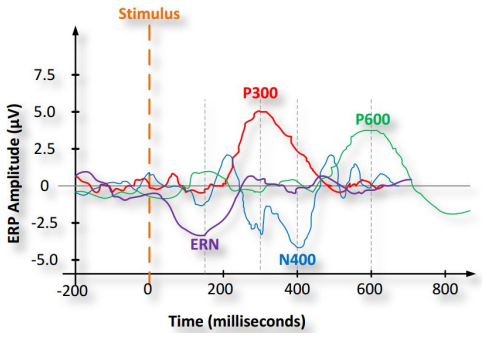


Figure 3.1: Different types of ERPs; image from dissertation of Dr Tamara Bonaci [14]

While their precise purpose in the attention pathways of the brain is not entirely defined [45], these signals are robust and predictable enough that they can be used to control a BCI. For example, a paralyzed individual can look at a screen of flashing characters, and over time an algorithm can determine the letter of attention and select it like pressing a key [21].

Because it is a reliable and robust signal that can be elicited for both auditory and visual stimuli [17], there is abundant literature about using the P300 response as a lie detector using the Guilty Knowledge Test [36][49][51]. In these kinds of studies, subjects are asked to "steal" or hide a picture of an object. Then they are shown a series of images, a small percentage of which are of the object they have puportedly taken. Since the reactions to seeing these images is unconscious, it should be difficult to trick the test; in fact, one study showed that deception tactics increased P300 response (in amplitude) [50]. However, this same group has also published results that demonstrate masking effect by having subjects

perform various actions in response to the non-targets that they knew about before the experiment started [52]. This kind of preparation is admirable, but may not be feasible in all cases: you might not know what they're probing you for, or how many irrelevant probes you need to respond to.

Despite this puported infalliability, these methods still aren't allowed as evidence in court in the United States. For one thing, the technology is still so new that its true long-term efficacy has not been documented. The MacArthur Foundation Research Network on Law and Neuroscience has a prepared document on its website that lists propr court precedent on why fMRI cannot be admitted in court: it is too new and unproven [40].[3] In the US, it is also illegal in most cases for a private employer to require a lie detector test as a condition of employment, which presumably would include a neural lie detector test [39].

There are also vastly differing error rates that can be attributed to factors like experimental design [13]. Until there is a standardization of methods, these kinds of studies will presumably remain academic in nature only.[4]

### *3.2 Range of Applications*

The range of BCI devices and diversity of companies offering them is increasing every year. Where previously such technology would only have been used in a medical setting, it is now clear that there are applications for entertainment and augmentation. Events like the Consumer Electronics Show[5] and Experiential Technology Conference[6] thrive on unveiling the latest and greatest. The video game company Valve even had a presentation at this year's Game Developers Conferene about how they're working on BCI technology [60]. And, there's

---

[3]There is also a question of if this kind of technique would violate a suspect's right against self-incrimination.

[4]There won't be an in-depth discussion of military uses of this kind of technology, except for a passing reference in Sect 3.2.

[5]`https://www.ces.tech/`

[6]`http://www.xtechexpo.com/`

already a company called Neurable[7] advertising a combined virtual reality/EEG headset that allows users to control a virtual environment using neural signals. The first reliable BCI device released with an accessible range of programs and games would be revolutionary and the gateway to making this technology as pervasive as the XBox and Playstation.

Marketing companies have already realized the power of neural data. The Nielsen company, for example, acquired NeuroFocus in 2011[8] and subsequently in 2015 named itself the largest consumer neuroscience organization[9]. Images like the one in Fig 3.2 show the benefits of use to trades from media to advertising[10].

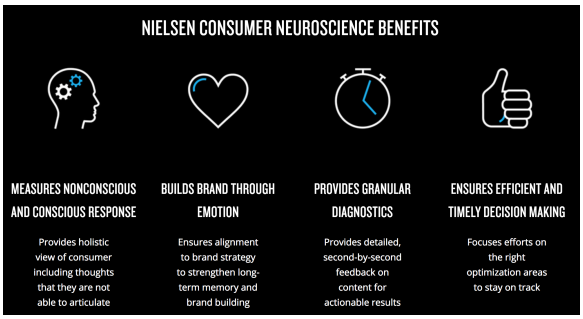


Figure 3.2: Screenshot of Nielsen's advertised benefits of consumer neuroscience

Granted, the current use of Nielsen's neuroscience research is limited to participants who volunteer to take part in product studies, and to devices that individuals choose to use. On the other hand, devices like a Deep-Brain Stimulator (DBS, or sending electrical current directly into the brain) are surgically implanted and medically necessary for some individuals. Currently they are only prescribed for motor disorders (Parkinson's, essential tremor), and is being investigated for other chronic conditions such as Tourette's, depression, pain, and ALS (Amyotrophic Lateral Sclerosis). But, they have also been found effective

[7] http://neurable.com/

[8] http://www.nielsen.com/us/en/press-room/2011/nielsen-acquires-neurofocus.html

[9] http://sites.nielsen.com/newscenter/nielsen-creates-the-worlds-largest-consumer-neuroscience-organization-with-acquisition-of-innerscope-research/

[10] http://www.nielsen.com/us/en/solutions/capabilities/consumer-neuroscience.html

in treating disorders such as Tourette's and depression. Researchers are investigating how to make these devices more effective, which includes the capability to sense neural signals as well as stimulate neural tissue. In some experiments, patients have a cortical strip, or small set of electrodes, implanted just under their skull so that further research can be done toward volitional control of a DBS system [15][61]. Individuals with a range of neurological conditions may soon be walking around with their neural signals continuously recorded. Then, the only thing that malicious entities would need is information on what and when the person is looking at; this kind of technology already exists in Google Glass.

In between voluntary focus group studies and permanent neural recordings are situations where the technology isn't permanent, but may be required. For example, the Army is investigating using neural signals to determine if soldiers are falling asleep while driving convoys [69]. Eventually, it may be feasible to determine more than just sleep state from these kinds of devices. Driverless cars may eventually supercede this kind of technology, but it does raise the issue of using neural signals as a form of surveillance.

### *3.3 Prior Literature on Elicited Information*

The subfield of elictation of information is newer, but relies on the same techniques as those used in the guilty knowledge test. This method, which relies on information that is produced through the process, was identified in 1996 as another way of obtaining knowledge about a system without directly attacking the system itself: a side-channel attack [32]. Using this kind of technique it is possible to obtain information by observing characteristics of the system, like the acoustics developed by a computer's CPU as it works [25].

So just like listening to the sounds a computer makes, one can interpret the signals generated in the brain in response to stimuli to infer information [33]. There are distinctive findings from the existing literature that can inform this research, and these are addressed next.

*Martinovic et al: Demonstration of Concept*

Martinovic *et al* published the first results on the viability of a side-channel attack on a BCI in 2012 [37]. They used a combination of analysis techniques and stimuli in an experiment that was modeled on potential real-life use cases. For example, some subjects were compliant in classifier training (counting digits), like a user might do to calibrate their BCI for optimal performance. Others participated in passive training, where they were shown stimuli with one target that would evoke a response with high probability compared to the rest (for this experiment, they used a picture of former president Barack Obama as the known target). Neural signals were recorded with a commercially-available EPOC Emotiv EEG system; the authors note that electrode placement with this device is not ideal for P300 detection since there is inadequate electrode coverage in the parietal area of interest.

The experiment had five categories of information: the first digit of a fake pin number (thought of in the experiment by the subject), which bank the subject used, month of birth, identifying a face, and where they lived (using images of parts of a map). For each category of stimuli, subjects were shown a question on a computer screen and then shown a series of stimuli (each stimuli was on screen for 250ms). There were no instructions about what to think of during the trials, and subject interaction was kept to a minimum. The chosen metric for demonstrating information leakage was comparing Shannon entropy[11] to random guessing (chance). This method relates to how well the classifier works: if a classifier can perfectly identify a target on the first guess there is 100% information leakage, while lower percentage leakage means the classifier is worse at extracting desired information.

On average, they demonstrated overall information leakage of 10-20%. Information leakage went as high as 43%, depending on the type of calibration (if they utilized the participatory training classifier or the passive facial response classifier). Of note is the maps condition, where subjects were asked to count the number of times they saw their location of residence

[11]Shannon's entropy, from information theory, is a way to measure the amount of information that can potentially be gained when one learns of the outcome of a random process. More entropy means more information can be gained once the experiment outcome is known.[35]

on the map on the screen. This active participation did lead to a higher percent reduction in entropy over random guessing (on this scale, 0% chance and 100% means always finding the correct answer).

An interesting finding of this research is the need for specificity in stimuli. In the category of banks, they originally showed bank logos but were unable to obtain sufficient results. Then, they changed to ATM logos and were able to calculate over a 15% reduction of entropy. The reason for this difference could be that we are accustomed to seeing all sorts of bank logos: as we drive, on television, on the internet, etc. However, we preferentially will notice the ATM logo for our bank, because we are used to seeing it when we get cash out (and will avoid other banks and convenience fees).

*Frank et al: Subliminal Stimuli*

Following the Martinovic experiment, many of the same researchers did a follow-on study where that involved recognition of subliminal faces. Each subject participated in training, with digit stimuli shown for 250ms. But for the experiment, stimuli were shown for 13.3ms.[12] There were two types of stimuli, a black and white picture of former president Barack Obama, or a blurred out shape of equivalent area. These were shown every 5 seconds, overlaid in a Charlie Chaplin movie. The authors discuss several analysis methods, but only in their conclusion section do they report that they could achieve 66% accuracy with a naive attack strategy in determining if a subject recognzied the president, and this could be increased to 90% accuracy.[13]

In addition to the difference between stimuli presentation time, this experiment used an ActiveTwo BioSemi 64-wet electrode system. This drastic increase in electrode coverage, as well as the improved signal to noise ratio from the conductive gel, definitel helped in increasing accuracy.

---

[12]Due to using a CRT monitor that had a maximum 75Hz refresh rate.

[13]It is less than clear where these numbers come from since they are not reported in the results or discussion section first.

(a) Original video

(b) Video with hidden target stimulus

Figure 3.3: Examples of experimental stimuli, from author manuscript [24]

*Bonaci: Elicitation During Gameplay*

Both the Martinovic et al and Frank et al papers discuss at length the kinds of scenarios where one of their attacks could occur. In her dissertation research, Dr Tamara Bonaci developed and tested subject responses to stimuli in a computer game [14]. Subjects were instructed how to play Flappy Whale, in which they control the vertical displacement of a whale by contracting and relaxing their forearm; muscle electrical signals were recorded using an electrode. While subjects played this game, their brain signals were also monitored. This was to record their neural response to subliminal stimui (7ms display time) displayed on the screen. Subjects were shown images in a range of categories, including fast food chains, car logos, coffee shop logos, and logos of sports teams. In subsequent analysis, the result was 44% success in determining a subject's affinities, aggregated across all categories[14]

*Lange et al: Elicitation Without Training Data*

The final paper to be discussed involves follow-up research. Here, Lange et al replicated the experiment performed by Martinovic et al, but used only 4-digit PIN numbers as stimuli [34]. Subjects were asked to provide a 4-digit PIN number at the start of the test, and this was shown in addition to five randomly generated 4-digit numbers. Brain signals were recorded with 32 electrodes on a Cephalon Waveguard cap. In the first analysis, they used

[14]Results calculated by Matthew Ehlert during a 2016 CNT Research Experience for Veterans.


Figure 3.4: A subject plays the computer game Flappy Whale. The electrode recording muscle electrical activity (that controls the whale) can be seen on the subject's left arm.

the concept of Perceived Information (PI, from the field of cryptography) to develop their classifier. This worked for seven fo the eight subjects, and interestingly for one subject they could not calculate a PI because one of the randomly generated non-target stimuli was also a combination of digits relevant to the subject.

Two of the subjects were also analyzed using an unsupervised analysis; this meant there was no existing training data. This was done with the assumption that the recordings for the targe digit (1/6) could be differentiated from the others (5/6). Compared to the results from the trained PI data, it took 5-10 times the number of additional observations to caluculate a result. So this kind of attack is possible, but takes more information.

### *3.4 Limitations of Prior Work*

The biggest fault of the Lange paper is that they only performed the unsupervised analysis on two of their seven subjects (discounting the one who had ended up with two target stimuli). While it's true that for a BCI to work efficiently, a subject will have training data, but a malicious entity may not have access to it. The Martinovic paper did have a novel solution to this problem, by guessing that there is one stimuli (a famous face) that could be used in

passive training [37]. Depending on the scale of the attack, there may not be enough time to show passive training stimuli and then the probe+non-target stimuli.

Even though many of the experiments described here used subliminal stimuli [24][14], subjects still reported that they could see and identify the stimuli or could at least see something. Unless an individual has a computer with a fast enough refesh rate (at least 144Hz), it may not be possible to show many stimuli before the attack is noticed. Or, the stimuli would have to be part of the environment.[15]

Another point is that the Martinovic paper is the only one that uses a commercially-available EEG headset; the rest were research-grade. It is important that those results were possible without the need for conductive gel or active current electrodes, but the relatively low redution in entropy (10-20% average) shows that there is the possibility of progress. The other results should not be discounted yet either. The ongoing work of companies from Sect 3.2 means electrode technology and signal processing will increase as they aim for a commercially viable product. Thus, it is not unreasonable to use research-grade EEG systems, which is why one was chosen for the experiment described next.

### *3.5 Experimental Design*

The experiment for this dissertation was designed to determine how accurately a digit could be calculated based on an "odd one out" methodology alone. To remove counfounding factors, it was a standalone experiment without a game component. This research was approved by the University of Washington Human Subjects Division under protocol 52192/2056.[16] All subjects were volunteers who were given a $20 Amazon gift card for their participation.

---

[15]One interesting concept identified by Tim Brown is that easter eggs in video games could be used as stimuli, and even if they are completely irrelevant in the context of the game (cheeseburger in a medieval chest), it won't raise suspicion because these kinds of non-sequiteurs are a part of gaming.

[16]The study number changed because of an update to the tracking system in 2017.

### *3.5.1 Protocol*

Ten subjects (8 male, 2 female, average age 31.3 minus one subject who did not provide age, average Edinburgh Handedness 69) completed the experiment, and all self-reported lack of dyslexia or neurological disorder that could be impacted by flashing stimuli. Only seven were used in analysis: one subject kept falling asleep (resulting in a protcol change described later in this section) while two other subjects failed to follow experimental protocols (e.g. drinking from a cup in the middle of an experimental recording).

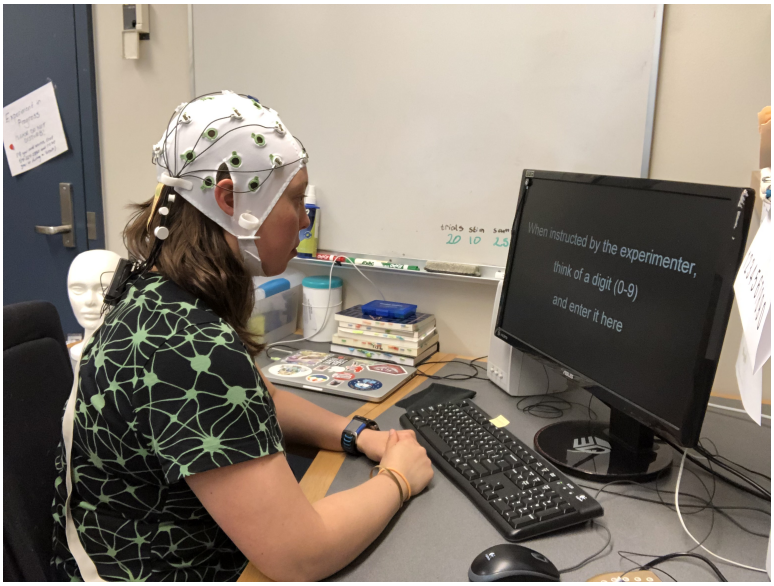


Figure 3.5: Experimental Setup

Neural electrical activity was recorded with a BrainProducts system (actiCAP slim electrodes, V-Amp amplifier), with 16 electrodes placed using the 10-20 system [27], as shown in Fig 3.6. Maximum electrode impedance was 50KΩ at the start of the experiment and fell to below 20KΩ by the end (approximately 60 minutes). There were two computers involved in the experiment: one dedicated to showing the stimuli, and the other for saving data.

After placing the cap, subjects were given instructions about the experiment. They were told not to move too much because their movements would impact the recording (demon-

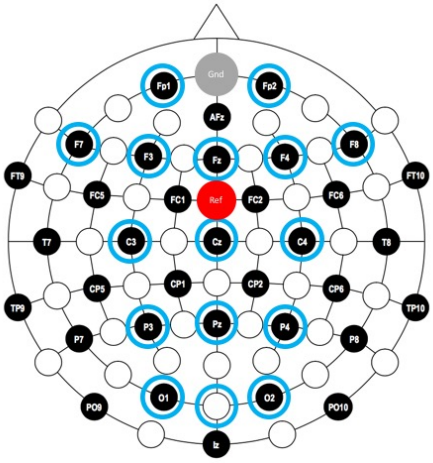


Figure 3.6: Modification of a BrainProducts EasyCap image to show 16-electrode placement

strated to the subjects by having them blink and clench their teeth while watching a live datastream). Subjects were informed that they would be doing 10 rounds of the experiment. At the beginning of each round of the experiment, they would be asked to pick a number and enter it on the computer. Then they should stare at a fixation point at the center of the monitor. The stimuli were placed in cardinal (top, bottom, left and right) directions 1.14" away from the fixation point. Each trial consists of 1500ms: 200ms baseline, 10ms stimuli shown, and 790ms to record for P300. There was also a 500ms gap introduced to ensure that data streaming was successful and there weren't any buffer issues. See Fig 3.7 for a graphical depiction of a trial. The stimuli were shown in a pseudorandom order. At the end of the round, the subject would be asked to enter their number again, and allowed to take a break if necessary.

To overcome fatigue, subject responses were implemented. Every other round, the subjects were told to press the keyboard spacebar when they saw their selected number on the screen; the ordering was spacebar/no spacebar was counterbalanced between subjects. The spacebar entries were not recorded.

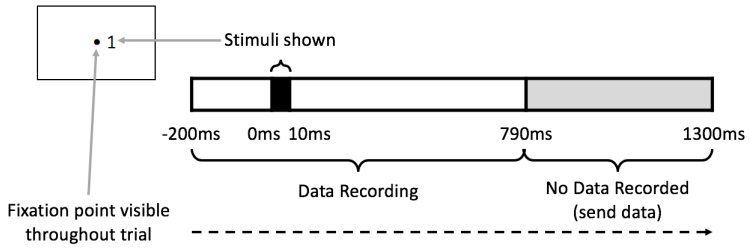


Figure 3.7: Graphical demonstration of trial timeline

### *3.5.2 Analysis: Regression in Real-Time*

The data was recorded and processed in Matlab on a secondary computer. The end of each trial was noted via light sensor on the stimuli display monitor. The data was processed using a zero-phase Buttereworth digital filter (3rd order, 30Hz lowpass), then an absolute value Hilbert function. It was then parsed into trial epochs by displayed stimuli for analysis.

At the end of the round, a least-squares linear regression was performed on the average of all epochs for each stimuli; the maximum regression value was output as the potential target digit.[17]

Overall, linear regression is accurate at calculating the target digit 2-3 times per 10 rounds per subject (minus one subject where no correct answers were caluclated). A tabulation of the correct and incorrect calculations is presented in Table 3.1.

### *3.5.3 Passive vs Active Involvement*

Five of the 7 subjects analyzed performed the alternating space/no space task for each round of the experiment. For space-bar rounds, the regression algorithm correctly calculated the target digit 8/25 times, compared to the 6/25 correct calculations for non-space rounds. This

[17]Ideally the linear regression would have been done 20 times, after each 10 demonstrations of stimuli. This would have enabled calculating the potential target digit and stopping the experiment in real time when a stimuli was determined to be the target with strong certainty. However, processing constraints on the analysis computer prevented this without multithreading or extending the interval between stimuli.

| Digit | Number Correct | Incorrect digits |
|---|---|---|
| 0 | 0/6 | 4, 4, 5, 7, 8, 9 |
| 1 | 3/6 | 4, 5, 8 |
| 2 | 1/8 | 0, 1, 4, 6, 7, 8, 9 |
| 3 | 1/8 | 0, 1, 2, 2, 5, 5, 6 |
| 4 | 1/7 | 0, 0, 1, 2, 2, 7 |
| 5 | 2/8 | 1, 2, 6, 6, 7, 8 |
| 6 | 4/8 | 1, 3, 8, 9 |
| 7 | 0/5 | 0, 1, 5, 8, 9 |
| 8 | 1/8 | 1, 5, 5, 6, 6, 9, 9 |
| 9 | 1/6 | 2, 4, 4, 6, 8 |

Table 3.1: List of possible target digits, number of times guessed correctly, and incorrect guesses calculated by the linear regression code.

matches with a finding from Martinovic, where subject involvement in counting the number of times their location of residence showed up on a map fragment increased the reduction of entropy compared to other stimuli categories [37]. This could indicate that entirely passive subjects are less desirable, and that some method of maintaining attention is required for more accurate results.

### *3.5.4 Predicting Current vs Future Number*

In addition to calculating the current digit, using linear regression was also successful at calculating the *next* digit the subject was going to choose. For example, a subject selects the number 5, and the computer caluclates 8. Then on the next round, the subject selects the number 8. This did not happen with the same consistency as calculating the current digit, but overall happend 13 times out of 70 trials (compared to 14 correct guesses of the

current digit over the same number of trials). The odds ratio[18] of calculating the digit for the following round over caculating the current digit is 1.04, which means the algorithm is slightly more likely to calculate the future digit.

As subjects were not instructed on how to think about or mentally maintain their chosen digit during the experiment, this could indicate how subjects were thinking about the experiment, and possibly preparing for further rounds (subjects did know that they were performing 10 rounds of the experiment).

### *3.6 Discussion and Future Work*

While this experiment demonstrates that elicitation of information without training information is feasible, there are numerous experiments that can be done to more concretely describe the effects. In the Lange et al experiment [34], they reported it took 5-10 times the number of trials to calculate a pre-generated PIN number without any classifier. Given that with their PI classifier it took between 5-10 and 30-40 trials to calculate the correct PIN, it might be unreasonable to expect 150-400 trials of stimuli from a subject. Even though this experiment did not calculate the correct digit every time, it was still able to do so 20-30% of the time with only 200 trials (20 demonstrations of each digit). It would be interesting to study the tradeoff of accuracy versus number of trials, and would might actually be considered acceptable for a particular use case. For example, if a company was only interested in whether or not a certain demographic mostly liked its product, 100% accuracy isn't necessary. However, if you're trying to discover someone's bank PIN number, you'd want a much higher accuracy rate.

Without statistical significance, it is possible to look at the incorrect digits in Table 3.1 and ask if there is any recognition by similarity. The digit 8 is an excellent example of this; 5, 6, and 9 all have similar curve patterns. However, without more information this could just be coincidence. Further experimentation with duration of stimuli presentation

[18] Association strength of two events, or the ratio of the odds of one outcome in the presence or absence of a second outcome.

and context would be helpful in understanding if this is a real phenomenon. Context may also be important: subjects were only choosing random digits that were displayed one at a time. The presence of other digits (like with a PIN or phone number) would provide more information to the subject about the overall relevance of the number combinations.

The calculation of future digits is also worthy of additional study. Since subjects knew that they would be performing 10 rounds of the experiment, it's possible that some of them were thinking ahead with their digit planning. A related experiment would be to have subjects preselect their digit series, to see if future digits can be calculated (given that the subject is trying to keep the digit order in mind). An n-back or forward test would also demonstrate the effect of P300 reaction to a prior/future digit compared to the nth digit.

Finally, it would be timely to revisit the experiments of Dr Bonaci [14] with the linear regression algorithm, to see if accuracy is altered by the presence of othe stimuli. The results of active involvement (pressing the space bar) would indicate that accuracy might be lower, but this could be tested by making the stimuli part of game play (rather than just passive subliminal images on the screen). The use of an eye tracker could also show effectiveness of calculating target stimuli; for example, does the subject have to be fixated on the stimuli?

# Chapter 4

# NEUROETHICS AND PRIVACY

It has been demonstrated in Ch 3 that information is obtainable through analysis of elicited neural signals. The accuracy of these techniques and precision of stimuli used will only increase as electrodes and processing techniques advance. Rather than wait to see how people react to violations of their privacy via BCI, it would be educational to examine perceptions about neural signals now. This will help educate the strength and direction of policy decisions that will be discussed later in Ch 5.

The purpose of this chapter is to create a definition of neural privacy, drawing on concepts (rights, interests, property, intimacy, and trust) from existing literature. While this definition can be used to frame what future legislation is protecting, it also frames the expectations of future BCI users. The definition will then be contrasted with the results from a neuroethics survey that asked subjects about their current perceptions of the privacy of neural versus non-neural information. This analysis includes analyzing the contrast between how mobility-impaired and non-mobility-impaired individuals perceive neural privacy. This particular contrast is examined because of the difference in life experiences that may occur for an individual with a mobility impairment: increased interaction with the medical community for continued care, physical characteristics (like a cane or wheelchair) that distinguish them from others, and the potential reliance on a family member or caregiver to accomplish mundane tasks (e.g. brushing teeth, grocery shopping, etc.) These kinds of relationships may lead to different thresholds and conceptions of neural privacy, given the bodily origin of the information in question.

But first, an important question is whether there is a difference between the raw, original neural electrical signals as the information that is derived from them. While the signals in

isolation are much less informative without contextual clues such as stimuli and movement, there may be some neurological conditions that can be insinuated using signal processing [6][55] and this information could be used against an individual (e.g. an indicator of dementia could be used to deny life insurance). However for the purposes of this analysis and discussion, privacy of neural signals will refer to quantifiable information that is determined from a combination of EEG and relevant environmental stimuli.

### *4.1 Context from the Origins of the Modern Definition of Privacy*

While privacy in general denotes the ability to hide or prevent access to a person/place/idea, it is very much defined by the social contexts and expectations of a society.[1] Particular to this dissertation, technology has proved to be a major catalyst that pushes the notions of what privacy means and what it should cover.[2]

One of the most famously cited and foundational definitions of privacy comes from a time when photography was becoming more commonplace. Samuel Warren and Louis Brandeis wrote in 1890 that "[i]nstantaneous photographs and newspaper enterprise have invaded the sacred precincts of private and domestic life" as they described necessary legal changes to keep up with the new modern invention. The centering of the individual in making their own choices is a theme that remains constant:

> In every such case the individual is entitled to decide whether that which is his shall be given to the public. No other has the right to publish his productions in any form, without his consent. This right is wholly independent of the material on which, or the means by which, the thought, sentiment, or emotion is expressed. [68]

---

[1]This dissertation is examining a more Western-centered theory of privacy, but it is worthwhile to point out that this interpretation could be significantly different from another cultural perspective. For example, the 2019 CPDP conference featured a panel on Islamic concepts of privacy: `https://www.cpdpconferences.org/cpdp-panels/islamic-legal-conceptions-of-privacy`.

[2]The legal construct of privacy will be discussed in Ch 5.

The sentiment still exists today, but it is much more nuanced and difficult to apply. Warren and Brandeis reference the case of unauthorized photography of a woman acting in a play, and now in the 21st century we have smartphone cameras and social media to help make any photo go viral, even when it may not have been appropriate to take the picture in the first place. But the line between public and private is increasingly fuzzy, making it incredibly context-dependent: who is taking the picture, of whom, and for what purpose? Take the following example: the European Court of Human Rights has ruled photos that purport to show the health of Monaco's Prince Rainier III are appropriate for the public to see, but that photos of the royal family on vacation are not [23]. Assuming that both photos were taken in accordance with the Independent Press Standards Association Editor's Code of Practice [41], perhaps the health of a soverign is newsworthy in the interest of national security, while the family vacation is off-limits because we all assume our vacations are our own business.

These kinds of privacy discussions are complex because they involve societal norms and multiple parties (the photographer, the photographee, the publisher, etc). It may be easier to analyze neural privacy because it is fundamentally a relationship between a customer and the company that designed the hardware and software being used.[3] The following sections address ways to qualify the relationship between the two.

### *4.2 Privacy as a Right versus Interest*

Declarations of rights are a succinct way for a governing body (international to local) to set a threshold standard of behavoir for its members. But oftentimes their scope is limited (e.g. US constitutional bill of rights only applies to acts related to the government); and, there is no consistent enforcement mechanism (see the US failing to observe the right to seek asylum as proclaimed by the United Nations Universal Declaration of Human Rights [8]). If a right to privacy is claimed (particular to the United States in this analysis) does it mean anything,

[3]It could very well be that the hardware and software are from different companies, but for the sake of simplicity this discussion will consider them the same, as one entity receiving neural information.

or is there a better way to understand what privacy baselines we are owed? An alternative to a right to privacy is that privacy is an interest. The merits and drawbacks of each will be discussed in the following sections.

### *4.2.1 The Right to Privacy*

The simplest way to articulate a right to privacy is to point to the 1st, 3rd, 4th, 5th, and 9th amendments of the US Constitution [19]. But the word privacy isn't mentioned of any of them; the precedent has come through interpretations and case rulings. So what then grounds this right to privacy? Judith Jarvis Thomson examines this question and determines that the right to privacy is derived from other rights in her 1975 essay, "The Right to Privacy" [62].

Thomson's feelings on the subject of privacy are clear: "Perhaps the most striking thing about the right to privacy is that nobody seems to have any very clear idea what it is." She also asks, " Is *every* violation of a right a violation of the right to privacy?" To answer this, Thomson asks us to consider the concepts of positive and negative rights. It is the ability to do things or prevent others from doing things to an entity; it can also include one's ability to violate another person's positive and negative rights. Then, by examining the progression of these concepts through several vignettes, Thomson concludes with the statement that instead of a right to privacy, the act of violating other rights is inevitably what we call a violation of a right to privacy: it is not unique on its own.

Thomson offers the following illustrative example. A man who owns a piece of pornography he doesn't want others to see. The ways it can be protected — closing the blinds, locking it in a safe — can be countered in any number of ways if someone truly does want to view the picture. In this case, the positive right is the man's ability to do things with respect to the picture (like look at it), while the negative right is its prevention from being sold by others or torn apart. It is his right of things that *shall* happen or not happen to something. This is in opposition to something like a a subway map posted on the wall of the station. Even if you cover it up to prevent others from seeing it, the act of someone tearing

the cover off is not a violation of your rights. It is the property and ownership of the picture that differentiates it from the subway map, and Thomson uses this distinction to posit that it is a property right that is violated first, instead of a privacy right.

In another scenario that involves listening to a husband and wife arguing, the using of an amplifying device to eavesdrop on them violates their right to not be listened to (a right of a person). This goes on for cases like torturing a man for information (violating his right to not be hurt or harmed or to be forced to divulge information), or for sharing a fact that was given to you in confidence (violating a right to confidentiality). Thus there are clusters of rights that can be violated in a given scenario, and while privacy can be one right in a cluster, it is never in its own unique cluster. She closes the article by confirming that privacy is thus derivative because privacy does not need to be invoked to describe the other rights in the cluster.

Unfortunately, this might be too much of a simplification that discounts some of the nuance. On this topic, Julie Inness provides an extended critique of Thomson's work[29].[4] Inness also disagrees with how the concept of positive and negative rights are applied because they limits rights to physical objects that require attribution of ownership. This point about ownership is worth noting for this discussion about neural privacy. From Thomson's perspective an individual can violate another's right to not be seen or touched, which is also a privacy violation, while Inness responds that we can limit access to our bodies and posessions. So for an arguing couple, it is not the content of the conversation overheard that warrants the claim of a violation of privacy, but the mere fact that in this scenario a property of the body was accessed to begin with. The lack of a need for ownership is also demonstrated with the example of love letters. By sending them, you no longer have ownership and yet your privacy can still be violated if the information is revealed. The necessity of deriving rights from other rights, as Thomson seems to do, may deprecate the gravity of the original right (e.g reading someone's text messages over their shoulder on the

[4]She is making these statements as a counterpoint to Thomson's work, instead claiming that privacy is directly related to intimacy, as will be discussed in Sect 4.4.

bus is a violation of their right to text in peace instead of a privacy violation), while possibly allowing for dissimilar rights to suddenly be grouped together.

This contrast in absolute rights versus derivative rights is intriguing for the discussion about neural privacy in two ways: is neural privacy its own right distinct from ordinary privacy, and do the concepts of ownership and bodily autonomy matter more or less to neural information? Thomson might assert that neural privacy is simply derivative of whatever was used to obtain the information (right to only have pertinent BCI-related information elicited). The argument from Inness' perspective is that the person may not have original ownership of the content of the neural information (you don't personally own the logo of your favorite coffee brand). Plus, your bodily autonomy is being infringed upon when you are no longer needed to physically express a thought (such as stating your favorite coffee brand).

If neural privacy isn't intuitively a right from either viewpoint, perhaps there is a different way of considering it. Next, privacy will be discussed as an interest, with the contrast of right versus interest at the end of the section.

### *4.2.2 Privacy as an Interest*

Judith Wagner DeCew chooses instead to talk about privacy as an interest, which is something that can be "invaded" or protected as deemed necessary [20]. She chooses this distinction over calling privacy a right, because a right is a justified claim, and a claim is only an assertion that has merit. This allows her to discuss points of legality and morality that are significant on their own and do not require the additional framework that would be required of a right (e.g. why it must be followed). And while she would prefer to refer to an interest in privacy, she still utilizes the terminology of rights because they are so pervasive in the broader privacy discussion.

The fact that she is unable to completely dissociate her arguments from others to which she doesn't entirely agree, indicates how deeply the concepts of privacy and rights are linked, and how culturally/regionally these norms change. On the one hand, it would be convenient to completely dispense with the right to privacy and allow individuals to set their own

standards. On the other, it would be chaotic to not know if and when you are violating someone's privacy interest. As an analogy, some people think that it is their right to be able to walk on the left side of an escalator while others stand on the right. But entering or leaving the UW lightrail station during rush hour demonstrates that not everyone believes this. You may have an interest in walking, but your ability to is physically invaded by a person standing in front of you.

But seeing as there is no universally stated and enforced right to privacy, perhaps what we all experience is an interest in privacy that is normalized to a particular environment: a right to privacy is different between countries, or even different times of day (leaving the curtains of your bedroom open during the day but closing them when you're changing at night or sleeping). Then it is up to legal interpretation to decide when an interest becomes a right.

### *4.2.3 Privacy as a Right and Interest*

Thomson, Innes, and DeCew all have different perspectives on the right (or lack thereof) to privacy. I propose that each can contribute to the understanding of neural privacy. Innes would declare accessing neural information without consent to be a violation of privacy, because it is an act of obtaining a property of the body. Thomson, not ascribing to the validity of a distinctive right to privacy, might say the violation not to have our thoughts taken by others (right of the person). And DeCew could consider how we have an interest in preventing others from learing our secrets, but invading it is acceptable.

It is perhaps beneficial for the remainder of this chapter to state that until there is a law or statute that is violated, neural privacy is merely an interest. Then when the terms of service are breached or a company does something illegal according to state or federal law, it is a right to nerual privacy that is violated.

### *4.3 Are Our Thoughts Our Property?*

Innes and Thomson disagree on whether or not property is necessary for a right to be violated. Inness used the example of a love letter for why ownership isn't necessary to have one's privacy violated. But, the letter was written or typed by hand, using words articulating thoughts and mental imagery. So if a malicious entity steals the letter and reads it, they have access to the thoughts you were having when you wrote the letter. Are these thoughts now stolen, or owned by the person who read them? Meanwhile, Thomson relies heavily on the concept of ownership in her discussion of privacy not being a right: it is not a privacy violation, but a violation of one's property rights that is then a violation of privacy.

There may not be a good theoretical answer to this question, but it may be better answered in legal terms. There is a vast field of law dedicated to protection of ownership. Copyright, trademarks, and patents can all be conceived of as protections of expression for someone's thoughts. But before these designations, there is an uncertain period of ambiguous ownership; yes, an individual conceived of them, but what is preventing someone else from using them without consent? Much of what we do does not warrant some sort of legal protection (e.g. you don't need to copyright the imaginary song that you sing in the shower), and it would be unreasonable to assume that every elicited thought can be immediately provided some sort of legal status to assert a property claim. The problem is that we don't know how valuable our neural thoughts are until something is done with them — it is difficult to predict what will become important.

One option is to assign ownership to everything that is elicited neurally; either to the individual from which it was elicited (it's their neural information), or the company that elicited it (their algorithms calculated it). This eliminates any legal gaps, but in the latter case means that whatever is elicited is owned by the company, regardless of if it was the intended information. Particular to cases of malicious use, this might even be some kind of theft. It also incentivizes the monetization of everything about us, as is already happening on the internet. The more consumer-centric approach, and that of Inness' perspective, is

that physical ownership of the information is not important to the content itself. So even if a company somehow elicited the lyrics of your imaginary shower song, they should still belong to you, particularly if the lyrics of the song had nothing to do with the purpose of you using a BCI to begin with. Any elicited information should be considered your own property that you are providing for a purpose (e.g. controlling a videogame). And, information that is elicited without a purpose to which you have agreed, should be considered a form of property theft.

### *4.4 Using Intimacy to Interpret the Sharing of Neural Information*

Another way to interpret the relationship of our neural information with the outside world is through the lens of intimacy. Julie Inness [29] says that, "The intimacy of information stems from the role the information plays for the agent, specifically the role the information plays for the agent when she conveys it to others." Her definition of intimacy includes the specification that choice is invovled on the part of the agent providing information, and that the value comes from the original relationship one has with the information.[5] So, a person can assert their privacy if they do not believe that the receiving party has an intimate understanding of the information; there is no respect for the care embodied in the fact that information was shared.

The choice described by Inness on the part of the agent (person using a BCI) may be what's important in using intimacy to understand neural privacy. If one chooses to allow a company to elicit information that can be used for future targeted ads, the user is allowing for an intimate relationship because the receiver (the company) is able to create value (monetary) from the information.

It may be difficult to convince a user that they have an intimate relationship with a

[5]The full definition of intimacy is:
Let Y = intimate decision about x.
Let x range over instances of access, instances of information dissemination, and the agent's activities.
To call Y "intimate" is to claim that it involves a choice on the agent's part about how to embody her love, care, or liking. Y involves such a choice because x draws its meaning and value from the agent's care, love, or liking.

gaming company or online retailer, because the word intimacy tends to have different connotations (this is not a sexual or physical relationship). However, it may be more effective to ask the user, "Do you believe that by providing information at this time it will be respected by the entity with care equivalent to the love or liking you have for it." And if the answer is yes, then it is truly the user's choice to share.

When information is elicited not by choice, the violation of privacy can be linked to a lack of acknowledging that the sharing of information is in fact a relationship. And if a company is being truthful up front about why information is being requested, they may be initiating a intimate relationship that reassures the user. Having transparency and choice changes the relationship between user and company from one of taking to one of mutual understanding.

### *4.5 The Importance of Trust*

The transparency and choice from the previous section are only possible if the user trusts that the company is actually going to comply. Without this trust, there is no ability to foster a relationship of transparency and choice. The fact that consumers continue to use websites even if they know their information is being collected and sold (but not what and to whom), indicates that the lack of trust may no longer be an acceptable deterrent. From the perspective of industry, there is not sufficient oversight or punishment to deter obtaining information without consent or proper notice, so they are free to utilize any tactics they see fit.

Some studies have asked consumers about their trust in companies. A 2018 Pew Research Study [46] aggregated recent findings of consumer trust: 91% of Americans in 2016 agreed or strongly agreed that users had lost control over the collection and use of of their information, while 2/3rd of those surveyed in 2017 said that existing legislation is inadequate in protecting private information. And yet, people still engage on social media, surf the web, and use credit cards to pay for transactions. Would they continue this behavior with neurally-based information. This question of trust, and willingness, was asked on a neuroethics survey and will be discussed later in Sect 4.7.2.

So in addition to intimacy, our privacy can be defined by the trust we have in those who take our information. The concept of "privacy as trust" is discussed in depth by Ari Ezra Waldman, who defines this relationship through regulation (restricting for some, open to others) [67]. So if we trust an entity, we are open to having an intimate relationship of sharing information.

### *4.6 Defining Neural Privacy*

To sum up the arguments from the previous sections, neural privacy encompasses the following statements:

- We all should have an interest in protecting our neural privacy, but require additional legal frameworks to make it a right

- Defining and ascribing ownership is necessary to provide value to what is being elicited (controlling a video game vs monitoring emotions)

- Users should be able to trust that the information taken from elicited neural signals by a company will be used and interpreted properly, making the relationship between user and company an intimate one

### *4.7 Neuroethics Survey*

If consumers percieve a difference between information that is taken using existing, conventional means, and that which is elicited using neural signals, then there is a different kind of privacy and intimacy associated with the latter over the former. To investigate this, a neuroethics survey was distributed online. Subjects were asked to anonymously provide basic demographic data and fill out the survey. Over a perid of 24 days, 77 respondants completed the survey. Of the 77 respondants, 18 indicated that they had a mobility impairment and

used some sort of assistive device (e.g. wheelchair or cane).[6] A full printout of the Google Forms survey can be found in Appendix A.

At the beginning of questions that involved BCIs, the following definition was provided:

> Brain Computer Interfaces (or BCIs) can record brain activity while an individual is performing different actions (for example, blinking their eyes, playing a video game, or texting on a phone). BCIs are often used to give a user control of a computer using their brain activity.

### *4.7.1 Physical Versus BCI Privacy*

The first two series of questions were designed to see if subjects perceive information taken by two different methods. The intended goal was to understand if there is a difference in perceived violation of privacy between humans and computers. The two scenarios were designed to be as equivalent as possible, but in order to suspend disbelief (how would a human sitting behind you on the bus have access to text messages you haven't written yet). So each case involves a person using a BCI to control a cellphone; in the first question a person is intercepting the communication between the neural recording device and the phone, while in the second question it is an app not being controlled by a BCI on the phone that is taking the information. It was intentionally not stated whether the app was installed by the user or if it was supposed to have access to the neural information.

Answers were requested on 5-point Likert scale from Disagree Completely to Agree Completely, and the prompts were given as:

1. Imagine using a BCI-enabled mobile device while sitting on the bus. In this scenario, it would be a serious violation of my privacy if a person sitting behind me on the bus...

[6] Many of the demographic categories did not have enough responses to perform statistically significant comparisons, e.g. only nine people indicated that they had spent any significant amount of time outside of the US/were not currently in the US. Other categories such as race and gender are also available for future analysis.

(a) ... read the content of a text message I sent to my family member

(b) ... recorded a video of me typing a text message to my family member which they watched later to find out the content of the text message.

(c) ... obtained a recording of my brain activity while I typed a text message to my family member from which the content of my text message could be figured out.

(d) ... obtained a recording of my brain activity *as I was planning* to type out a text message to my family member such that the content of the text message could be predicted before I typed it.

(e) ... obtained a recording of my brain activity that included my emotional state while typing a message to my family member (e.g., to reveal feelings of inattention, boredom, excitement, anxiety, other states?)

2. Imagine a scenario where there is a recording app on your phone that can store everything you do on it. You can control this phone with a BCI device. It would be a serious violation of my privacy if an app on my phone...

(a) ... downloaded and stored the content of a text message I sent to my family member

(b) ... used a keylogger to record me typing out the message in real time, including what was typed and deleted.

(c) ... obtained a recording of my brain activity while I typed a text message to my family member from which the content of my text message could be figured out.

(d) ... obtained a recording of my brain activity *as I was planning* to type out a text message to my family member such that the content of the text message could be predicted before I typed it.

(e) ... obtained a recording of my brain activity that included my emotional state while typing a message to my family member (e.g., to reveal feelings of inattention, boredom, excitement, anxiety, other states?)

A visual depiction of the results are shown with two different analyses: a comparison of the person versus app privac violations in Fig 4.1, and a comparison of perceptions of privacy violations between indiviuals who report no mobility impairment versus those who did in Fig 4.2.

*Person vs App*

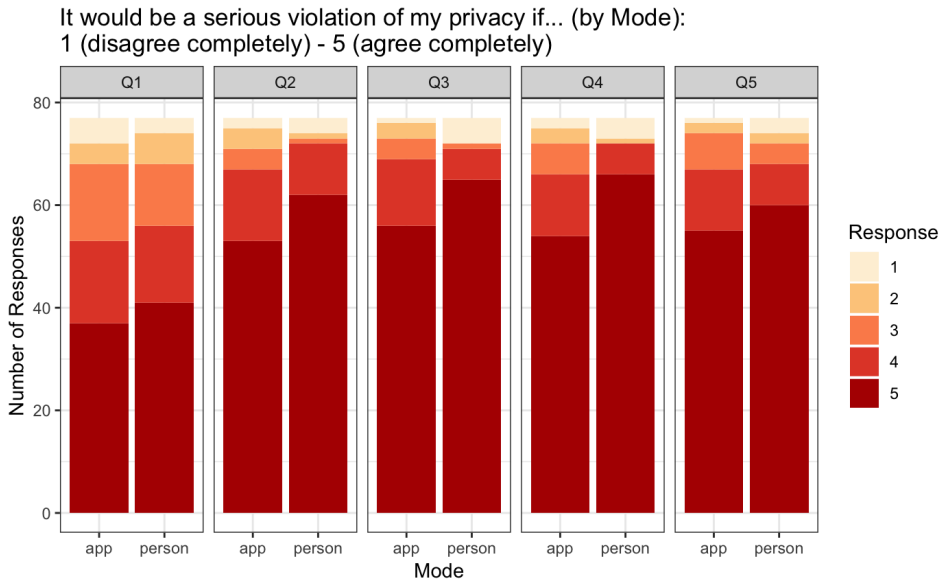


Figure 4.1: Comparison of number of responses between privacy violations that involve a human versus those that involve an app, where 1 is Disagree Completely and 5 is Agree Completely

To analyze the statistical significance of these results, the two modes of privacy violation were compared using an odds ratio. This method calculates the odds that exposure to a factor affects the outcome [59]. When comparing a person vs an app violating privacy for each question, an odds ratio of $< 1$ means subjects perceive the app-collected information as less of a privacy violation compared to if equivalent information was collected by a human. These results, along with the associated 95% confidence intervals are shown in Table 4.1.

| Question | Odds Ratio | 95% CI |
|---|---|---|
| 1 (Content) | 0.826 | 0.484 - 1.41 |
| 2 (Video) | 0.554 | 0.288 - 1.066 |
| 3 (Brain Activity) | 0.774 | 0.421 - 1.421 |
| 4 (Neural Planning) | 0.559 | 0.315 - 0.991 |
| 5 (Emotions) | 0.973 | 0.617 - 1.537 |

Table 4.1: Odds ratios for comparison of person vs app violations of privacy

For all five questions, subject responses indicated that an app posed less of a privacy violation compared to a person. However, only the scenario of planning neural activity is significant within a 95% confidence interval.

*Mobility vs Non-Mobility Impaired*

This question compares mobility vs non-mobility impairment, aggregating across the person vs app questions. The responses as shown in Fig 4.2 are plotted as proportion of responses, since this analysis is done across questions.

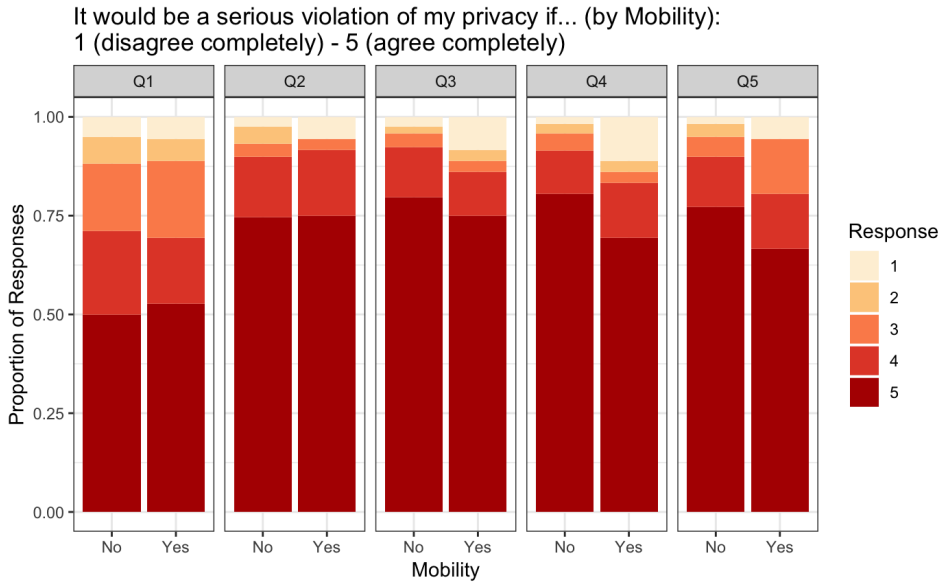


Figure 4.2: Proportional response comparison between indiviuals who reported a mobility vs non-mobility impairment (combines person and app scenarios), where 1 is Disagree Completely and 5 is Agree Completely

As with the person versus app comparison, the scenario of mobility vs non-mobility impaired individuals and their perceptions of privacy violations was compared via odds ratios. In this case, an odds ratio <1 means individuals who did not indicate a mobility impairment perceive the data collected to be less of a privacy violation, compared to those who did indicate a mobility impairment. These results, along with the associated 95% confidence intervals are shown in Table 4.2.

| Question | Odds Ratio | 95% CI |
|---|---|---|
| 1 (Content) | 0.946 | 0.479 - 1.869 |
| 2 (Video) | 0.960 | 0.403 - 2.287 |
| 3 (Brain Activity) | 1.238 | 0.458 - 3.344 |
| 4 (Neural Planning) | 1.760 | 0.655 - 4.726 |
| 5 (Emotions) | 1.704 | 0.613 - 4.739 |

Table 4.2: Odds ratios for comparison of mobility vs non-mobility impaired perceptions of privacy volations

With the 95% confidence interval, none of the scenarios had statistical significance to indicate that subjects with and without a mobility impairment had different perceptions of privacy violations.

### *4.7.2 Willingness Versus Trust*

The next two questions are related to feelings about who has access to information. This relates to the argument from Sect 4.1 and the relationship between privacy and trust from Sect 4.5. In the current internet ecosystem, there may be a willingness to have a social media website, but you don't trust the social media company with your data. These questions are directed at this sentiment: how willing a consumer is to share their neural information, and how much the consumer trusts an entity. The prompts were given as:

3. Imagine that a BCI device could detect your thoughts or feelings, such as desires for particular foods, attraction to particular movie stars, discomfort with particular political views, or how you're feeling physically and mentally. Please indicate how *willing* you are to share information with the following entities, on a scale from extremely unwilling (1) to extremely willing (5).

4. Imagine that a BCI device could detect your thoughts or feelings, such as desires

for particular foods, attraction to particular movie stars, discomfort with particular political views, or how you're feeling physically and mentally. Please indicate how *trustworthy* you believe each of the following entities are obtaining this information, on a scale from extremely untrustworthy (1) to extremely trustworthy (5).

The list of entities includes: a trusted individual/family member[7], medical professional, university researcher with Institutional Review Board (IRB) approval, governmental organization, non-profit company, and for-profit company. The results to these questions is shown in Fig 4.3.

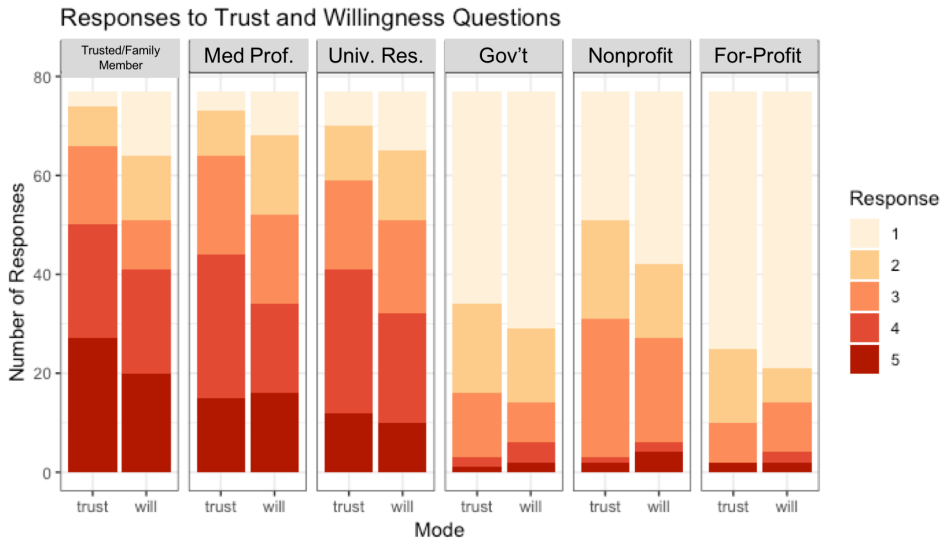


Figure 4.3: Aggregate responses to Trust and Willingness Questions. For trust, 1=extremely untrustworthy, 5=extremely trustworthy; for willingness, 1=extremely unwilling, 5=extremely willing

[7]Through a formatting error, a different entity was listed between willingness and trustworthy questions. It should have been "Emergency Contact" in both scenarios.

### *4.7.3 Comparing Value of Types of Data*

The last questions on the survey for this section compared neural data to other forms of data that could also be perceived as "private."[8] The categories were chosen because of their usage in prior literature, media prominence, or existing regulatory coverage. They were (in order): FitBit or equivalent fitness tracker, record of medical history (e.g. at your doctor's office), genetic information from a company like 23andMe, online shopping history, monthly credit card statement, and a journal/diary. The results for the following prompt are shown in Fig 4.4.

8. For each of the following, please indicate if you think recordings of your neural information should be more, equally, or less private than the comparison.

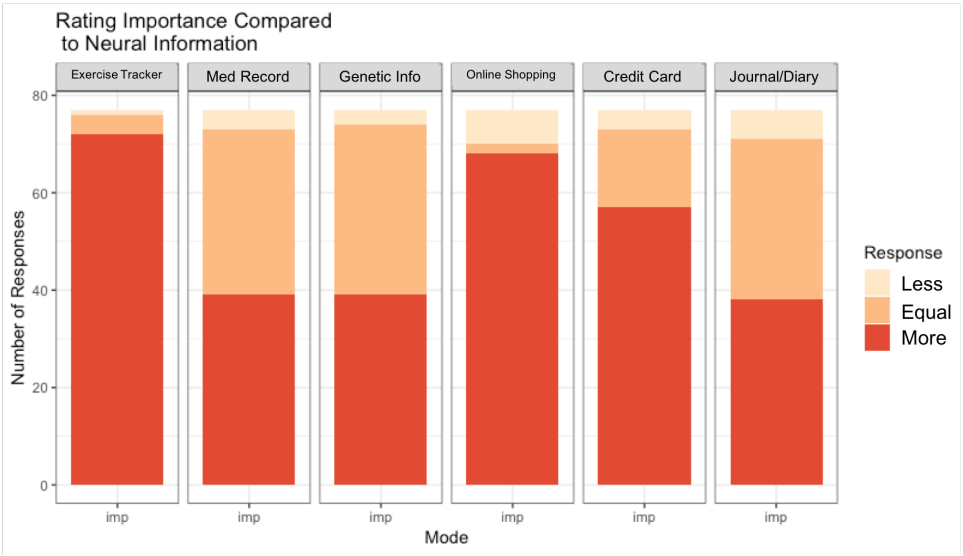


Figure 4.4: Results for Question 8

[8]Questions 5-7 are discused in Ch 5.

### *4.8 Discussion and Future Work*

Some of the results from this chapter corroborate findings from other chapters, and others have intereseting implications for future research. Most surprisingly, subjects indicated with statistical signficance that an app posed less of a privacy violation compared to a person only if neural signals indicating future actions were involved. And, the results from Ch 3 show that linear regression is almost as good at calculating future digits as it is the current digit. So, the scenario from Sect 4.7.1 is accurate for a future use case. The premise of the human subject experiments and most of the neuroethics survey was what is obtainable about one's current or prior state (already developed preferences, etc.), when in fact the future state (will I like this brand) could be a more frightening scenario. A future survey could address this, possibly with more detailed experimental results that outline varying levels of risk.

Next, one of the hypotheses for the neuroethics survey was that mobility impairment would affect a subject's notions about neural privacy. In fact, it seems that it only is a factor in a small subset of scenarios. This could mean that feelings on neural privacy are uniform across populations, or that there is more nuance than has been illuminated by this survey and analysis. For example, there might be significance between the mobility and non-mobility impaired subjects if analysis is done within the respective questions about person versus app monitoring of neural signals (instead of across). The large confidence intervals speaks to a large variability in the data.[9]

In the section on trust and willingness to share neural information, it is evident that individuals neither trust nor are willing to provide neural information with either the government, non-profit, or for-profit entities. This would track with the larger supposed societal narrative that individuals don't want to use things like social media because they don't trust them. What was unclear in the questionnaire was if the thoughts and feelings were the only thing being shared with the entities, or they were in addition to other information being

[9]Much of the statistical analysis in this chapter was done with the assistance of the UW Stats Consult program, and a graduate student who worked on this data for a class project. The full statistics report can be found in Appendix B.

collected (e.g. controlling a video game, typng text on a government website). Whether or not thoughts and feelings are the primary or secondary data may be of relevance for a question like this.

Finally, in the comparison of neural information to other kinds of information, it seems that the categories that are more closely tied to bodily information as well as journal/diary entry, are the ones that subjects are more likely to think need equal protection to neural information. A future survey could include a short answer space so that respondants can indicate why they responded the way they did, particularly in cases where someone indicated that they thought neural information should be less private than fitness tracker information.

# Chapter 5

# REGULATING EMERGING TECHNOLOGIES

The breadth of responses subjects had to ethical questions of BCIs in Ch 4 demonstrates the range of opinions and thoughts on the subject. Contrary to the myth that no one cares about privacy anymore, individuals do have quantitative opinions about who should have access to personal information. This leads to the question: how should bodily-derived information be protected, and does its use warrant increased (or at least explicitely stated) rules? There is an obvious lacunae in the legal landscape, as BCI-derived information falls between the protections afforded by HIPAA to medical data, and the mostly unregulated realm of personal data.

This chapter will describe initial questions of enacting regulation (what is it protecting, how do you define harm to an indiviual) and discuss them using analogous contemporary examples. Finally, I will present the responses to survey questions about what entities should be involved at various steps in the development of BCI technology.

## *5.1 What is Regulation Protecting?*

There are many ways regulation can be applied to neurally-derived information. The question is, what should it be protecting: the neural signals themselves, or the data that is derived from them? While it is the information obtained that is of interest to advertisers and companies, its source is a record of neural signals collected over a period of time (minutes to years). We have already seen how aggreated and inferred data have essentially become proprietary and protected by the companies that collect them, like clothing stores [12], or data brokers [47]. It is also essential to note that consumers may not know what is being collected and inferred about them: a BCI user could assume that their neural signals are

only being used to control a computer, when a plethora of other details about their lives is simultaneously being collected.

This broader scope of general data privacy is mostly outside the scope of this dissertation (see Sect 5.2 for more), but questionnaire data in Sect 5.3 describes when and who should be involved in oversight and regulation. To the question of what can be done for BCIs, there is one industry in particular that is already embracing bodily-derived information: insurance.

### *5.1.1 Genetic Information*

The Genetic Information Non-Discrimination Act (GINA) was passed in 2008 for the purpose of "establishing a national and uniform basic standard is necessary to fully protect the public from discrimination and allay their concerns about the potential for discrimination, thereby allowing individuals to take advantage of genetic testing, technologies, research, and new therapies"[2]. The increasing use of genetic testing helps with identifying and diagnosing medical conditions. But because many of these are correlated or occur only in unique populations, it may deter an individual to seek out treatment (sickle cell anemia in African-Americans is explicitely called out as previously addressed example in the legislative text), due to subsequent discrimination [2]. From an implementation standpoint, it invokes the 1964 Civil Rights Act as well as the Equal Employment Opportunity Commission to amend various existing federal regulations to prohibit employment and health insurance discrimination based on one's genetic profile.[1]

These kinds of protection are valuable to maintaining the privacy and trust of individuals who seek out genetic screening or counseling. However, it does run afoul of a new employment trend: insurance discounts for sharing information about how healthy you are (see Sect 5.1.2 for more information). A bill in the previous Congress was introduced to address this: HR 1313, the Preserving Employee Wellness Programs Act [22]. Contrary to GINA, the bill would mean companies could acquire the genetic information of a worker, or family member,

[1]Note that this does not apply to other kinds of insurance [38], which are discussed in Sect 5.1.2.

if requested for a wellness program and it would not violate GINA (Sect 3b in the bill language). And while the Affordable Care Act codified insurance discounts for participation in wellness programs, consumer rights and physicians organizations criticized that the bill increases the penalty for not participating in such a program as well as skirting other non-discrimination statutes because the information is being volunteered [4].

This kind of subvertive punishment for not sharing personal information is precisely what Peppet describes in his 2011 paper on the unraveling of privacy [43]. Many proposed privacy policies in the US require companies to put consumer in control of their own information, by requiring opt-in to collect information. Peppet states that this very act of giving control to consumer, a signaling economy, does in fact the opposite. By providing this control, consumers can be both benefitted (they know where their information is going) but also harmed. If everyone else is choosing to share information, and you don't, it could inhibit your ability to perform or obtain a good (say insurance). Or, you may be forced to disclose information that you would rather not, because it's required for a particular task. This is a novel form of what Peppet calls economic distress, as you begin to be punished for the privacy choices that you make. What is originally advertised as a fee for good behavior (putting a tracker in your car to get an insurance discount) can become a fee for perceived bad behavior. This is already happening with tools like Snapshot, which allows an insurance company to monitor how you drive and provide a discount. However, certain activities are deemed to be higher risk and increase premiums, such as driving at night (you might get hit by a drunk driver). For individuals who work swing or night shifts, disproportionately people of color and minorities, and who have to drive to get to work, they may be penalized even though they are technically doing nothing wrong [64]. One possible mitigation for unraveling is policies like "don't ask don't tell" or "do not use" which would prevent even asking for the informtion to begin with, of which GINA is one example.

### *5.1.2 Life Insurance*

Social media and personal health have proven to be a boon for collecting granular data on individual habits. The insurance industry in particular is able to benefit because they can exactly track how active or truthful someone is about their lifestyle. It incentivises "good" (i.e. healthy, non-smoker, not a risk taker) people to apply for discounts or monitoring, while others who cannot participate (e.g. preexisting condition) shoulder a greater financial burden: again, Peppet's privacy unraveling [43].

The ultimate endgame is insurance that is only available through persistent monitoring, and it has arrived. The John Hancock company has announced that it will only sell "interactive" life insurance starting in 2019 [11]. It's billed as a win-win, where customers are encouraged to become more healthy through discounts and perks, while the company saves money on payouts. This also means that a company has incredibly detailed and invasive information about your daily life, which can indicate larger patterns in the population. Aggregate GPS data has already shown the location of secret military installations [28], and what would prevent a polic department from seizing the location information of everyone at a protest?

If you don't have a smartwatch or equivalent (e.g. it wasn't subsidized by your employer or the insurer), then you have to input your activities manually. Should you be penalized if you are inconsistent with reporting your habit tracking, or for refusing pervasive monitoring? Rather than wait for you to upload your workout or grocery receipt, insurers in New York received approval to set premiums based in part on information that can be inferred from social media. The Wall Street Journal subsequently posted the following tips on how to avoid more invasive methods like a blood test for setting premiums:

- Don't post photos of yourself smoking on social-media sites.
- Do post photos of yourself running. Riskier sports, like skydiving, could complicate the situation.
- Use fitness-tracking devices that indicate an interest in fitness.

- Purchase food from online meal-preparation services that specialize in healthy choices.
- Visit the gym with a phone linked to a location-tracking service. If you visit the bar, leave your phone at home [54].

This kind of information mining is incredibly invasive, particularly since neither you or the social media platform may have consented to this kind of use of your posts. However, some see these troves of data as sources of valuable information. For example, researchers found that the kind of Instagram filter used could indicate depression [48]. With advanced computing and algorithms it's conceivable that a myriad of conditions could be discovered from BCI-recorded neural signals; for example, it's already possible to identify symptoms of Parkinson's Disease based on typing on a smartphone [7]. While these rates may be individualized for the person who is able to participate, the behaviors of others can help or hurt your own status if you choose not to. Just like the higher insurance premium for driving at night (Sect 5.1.1), you will be assigned the attributes of others in the absence of information. So even if you're in reasonable shape, if you're older you could still be subjected to higher rates because older people tend to be sicker or get injured more easily.

It is evident that companies are already utilizing every means available to discover more about us, and BCIs will provide them with another data stream. The next section will describe what obtaining this information means from a legal perspective.

### *5.2 What is the Legal Harm of Elicited Information?*

Say information were elicited from a BCI user, for example an unlisted phone number. Does the malicious entity have to call or publish it, or is the mere fact that they have obtained it sufficient for legal recourse? Is it possible to define a legal harm? The increasing adoption of biometric identification for personal and commercial use means private companies have access to all sorts of bodily information.

Few states have laws specifically limiting collection and use of biometric information: Texas (Business and Commerce Code, Title 11.A Ch 503), Washington (Chapter 19.375 RCW), and Illinois (740 ILCS 14/). The latter recently survived a challenge at the state's

supreme court that is relevant to the situation of neurally elicited information.

In Rosenbach v Six Flags [30], a mother sued on behalf of her son, who was required to provide a fingerprint to obtain a season pass to an amusement park. They were not notified that a fingerprint was required to obtain said pass, and even though the boy never returned the park could not identify if they still had the fingerprint or how long they would maintain it. In their suit, the mother (on behalf of her son) claimed that Six Flags' violation of two principles in Illinois law (consent and rentention policies) was enough of a harm to enable a private right of action. The amusement park claimed that there was no actual injury (e.g. misuse) of the fingerprint, and therefore there was no harm. In its analysis, the opinion of Chief Justice Lloyd Karmeier overturned the ruling of a lower court, noting that failure to follow the statute was in fact a harm. To quote from the ruling: "The injury is real and significant." In this case, it wasn't harm as if the company had gone out and used the boy's fingerprint for a nefarious purpose. The particular harm was in not being able to articulate up front that obtaining the fingerprint was a necessary condition to obtain the season pass, and not having (legally mandated) retention and destruction policies. Overall the decision was celebrated by civil liberties organizations such as EFF [53] (which had previously filed a friend of the court brief along with a consortium that included the American Civil Liberties Union and Center for Democracy and Technology).

The underlying requirement for consent is part of a larger conversation about how we blindly click or sign terms of service and other policies. This has been studied by Joseph Turow and others, who have looked at the longitudinal attiudes toward privacy policies [63]. Many consumers accept these policies without reading them because they inherently believe that the text is about how the company is going to protect their information, when in fact it is the other way around. The very act of calling it a privacy policy is a deception.

This all has implications for neural privacy. From the definition in Ch 4.6, there is the condition that neural privacy is an interest until there is a legal enforcement mechanism, and property rights over neural signals are retained until they have further legal standing. The implication is that privacy policies must either be clearer (so that users don't automatically

check the box without reading) or renamed so that there is no misinterpretation of their purpose. It also means that laws like the one in Illinois are needed in other states and/or federally, to ensure that neural information is not hoarded and misused.

### *5.3 Survey Responses*

To better gauge perceptions on the topic of regulation, three questions were included on the questionnaire from Ch 4 asking subjects to indicate whether specific entities should play more, less, or the same role in the development, regulation, and response to misuse of neural information.

Respondants were asked a series of questions related to how different entities should be involved in the development and deployment of BCI technologies; responses were recorded as less than current involvement (1), same as current involvement (2), and more than current involvement (3). The current level of involvement was intentionally excluded for each entity so that answers could be provided relative to the subject's perceived knowledge of existing regulatory structures. Not every combination realistically exists in real life, particularly cases of less involvement. However, they were included for uniformity and so that subjects could respond truthful to their opinions.[2] The entities asked about for each question were (in order): User, University Researcher with ethics board approval, Independent Regulatory Organization, Legislators, and Device Manufacturers. The questions asked are as follows (with the numbering indicating the order from the original survey):

5. A *blank* should play more/same/less of a role in oversight of the development of devices that can record and use your private neural information for the purposes of using a BCI, compared to current involvement.

6. A *blank* should play more/same/less of a role in regulating the use of devices that can record and use your private neural information for the purposes of using a BCI when they come to market, compared to current involvement.

[2] A future survey should ask subjects to provide examples to expand on their responses.

7. A *blank* should hold more/same/less responsibility in seeking reparations and compensation if neural information is stolen, or malicious entities are able to elicit private information from you while you are using a BCI, compared to current involvement in similar cases like data breaches.

### *5.4 Results*

The responses for all three questions are shown in Fig 5.1, and the results for this question were again analyzed using an odds ratio (same method as from Ch 4.7.1).

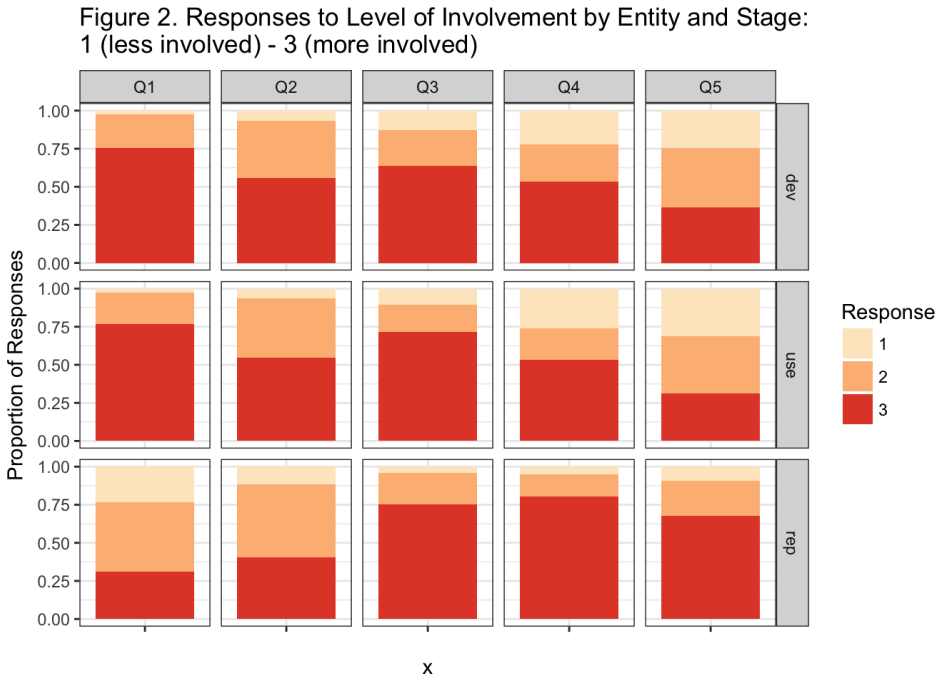


Figure 5.1: Aggregate responses to Questions 5-7; Columns represent the entity involved, and rows represent period of involvement.

*Entity Involvement*

First, the involvement of an entity was evaluated through the development cycle (i.e. how much should a user be involved in development versus usage versus reparations). An odds

ratio >1 means less involvement through each stage, while an odds ratio <1 indicates the entity should be more involved through each stage. The results are shown in Table 5.1. The entities that have odds ratios with confidence intervals <1 are significant.

| **Entity** | **Odds Ratio** | **95% CI** |
|---|---|---|
| User | 2.481 | 1.785 - 3.45 |
| University Researcher | 1.248 | 0.971 - 1.605 |
| Indep Reg Org | 0.709 | 0.563 - 0.891 |
| Legislators | 0.685 | 0.555 - 0.847 |
| Device Manufacturers | 0.741 | 0.597 - 0.921 |

Table 5.1: Odds ratios for entity involvement across stages (going down the column in Fig 5.1)

Subjects indicated that the larger the entity, essentially, the more it should be involved through the development lifecycle; this is statistically significant at the 0.05 level for independent regulatory organizations, legislators, and device manufacturers. Conversely, users and university researchers should be less involved through the development lifecycle, but this is only statistically significant for the user. Ideally, this is how it would be: users show preference for features through their purchasing patterns, or request features to be added to a device, and manufacturers respond. It order of operations may still be true in some sense, but it is not such a linear process and oftentimes device manufacturers are guessing or creating the features that they want consumers to purchase.

In between users and device manufacturers are some of the entities that can also have a say. University researchers have the ability to research and invent next generation tools and features that will eventually end up in a device. But, with revenues at stake these kinds of advances can be pushed to industry where there is more money and less oversight (no institutional review boards).

Independent regulatory organizations like IEEE (Institute of Electrical and Electronics Engineers) or ACM (Association for Computing Machinery) already have published standards that guide technology development. But without the close involvement of those developing the technology, they cannot proactively develop standards in any efficient manner unless there is a concerted effort on behalf of those developing it. To be more proactive, as the subjects said these independent regulatory organizations should be, they could be creating systems-level guidance for security, privacy, and interface requirements that can apply to any emerging technology.

Finally, legislators had the lowest odds ratio, indicating the strongest desire to be more involved through development to reparations and responses to malicious elicitation. The current state of the latter is perhaps the weakest and in need of the most reform to adequately respond to neural elicitation and other technologies. Since no baseline was provided to the subjects for this question, they may or may not know about the FTC's Section 5 authority in pursuing unfair and deceptive practices [18]. But a lack of resources (personnel and money) severely limits what the agency can do; this is something that legislators can fix, by mandating an increase in personnel and funding to accomodate investigations. Additionally, the lack of a private right of action (indiviual suing the company) in almost all cases means that consumers have little means of recourse. The private right of action in Illinois allowed Rosenbach to sue Six Flags, but this right would have to be applied in other states or federally.

*Involvement Through Stages*

Next, the responses were analyzed by stage. For Table 5.2, the odds ratios indicate how involvement should change based on who the entity is (as ranked by the order in which questions were asked). For odds ratios >1, involvement should decrease going from user to device manufacturer; for odds ratios <1, involvement should increase going from user to device manufacturer. The entities that have odds ratios with confidence intervals <1 are significant.

| Stage | Odds Ratio | 95% CI |
|---|---|---|
| Development | 1.461 | 1.276 - 1.673 |
| Use | 1.536 | 1.355 - 1.742 |
| Reparations | 0.620 | 0.513 - 0.748 |

Table 5.2: Odds ratios for entity involvement at each stage

This data confirms that the burden of responsibility shifts from consumer to device manufacturer as a technology moves from development to usage and responses to malicious uses. Subjects also indicated that they did not trust device manufacturers with their neural information (Sect 4.7.2), which is why they may desire increased reparations for any violations of privacy.

### *5.5 Discussion and Future Work*

Much of the above discussion about state laws and when/how entities should be involved in the development lifecycle would be simplified if there was comprehensive data privacy legislation passed in the US. This baseline legislation would go a long way to providing generic consumer data protection, but would not preclude consideration of stricter protections for biometric and neural information.

A future survey could include questions that baseline a subject's knowledge of a particular entity (e.g. "Are you aware of what empowers the Federal Trade Commission to investigate unfair and deceptive practices?"), and asking specifically how the burden of responsibility passes through the development cycle. Another interesting question to ask is what subjects think is appropriate compensation, and whether this is based on the kind of information that is obtained. This alludes to the larger question of how much our information is worth to companies, and whether consumers believe that biometric information is more valuable than other types of personal information. We already know from Sect 4.7.3 that neural information is more private than existing kinds of information: does this mean it is more

expensive?

The existing survey data presented here and in Ch 4 is enough to begin conversations with device manufacturers and other regulators. The subjects for this survey had statistically significant opinions about who and when should be involved in developing neural technologies.

# BIBLIOGRAPHY

[1] Regulation 2016/679 of the European Parliament and of the Council. https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:31995L0046, 1995.

[2] H.R. 493 (ENR) -Genetic Information Nondiscrimination Act of 2008. https://www.govinfo.gov/app/details/BILLS-110hr493enr/summary, 2008.

[3] Directive 95/46/EC of the European Parliament and of the Council. https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32016R0679, 2016.

[4] Reed Abelson. How Healthy Are You? G.O.P. Bill Would Help Employers Find Out. *The New York Times*, Dec 2017.

[5] Alessandro Acquisti. Privacy in electronic commerce and the economics of immediate gratification. In *Proceedings of the 5th ACM conference on Electronic commerce*, pages 21–29. ACM, 2004.

[6] Noor Kamal Al-Qazzaz, Sawal Hamid Bin Ali, Siti Anom Ahmad, Kalaivani Chellappan, Md Islam, Javier Escudero, et al. Role of eeg as biomarker in the early detection and classification of dementia. *The Scientific World Journal*, 2014, 2014.

[7] Teresa Arroyo-Gallego, María Jesus Ledesma-Carbayo, Álvaro Sánchez-Ferro, Ian Butterworth, Carlos S Mendoza, Michele Matarazzo, Paloma Montero, Roberto López-Blanco, Veronica Puertas-Martin, Rocio Trincado, et al. Detection of motor impairment in parkinson's disease via mobile touchscreen typing. *IEEE Transactions on Biomedical Engineering*, 64(9):1994–2002, 2017.

[8] United Nations General Assembly. Universal Declaration of Human Rights. https://www.un.org/en/universal-declaration-human-rights/, 1948.

[9] Kenneth A Bamberger and Deirdre K Mulligan. Privacy on the books and on the ground. *Stan. L. Rev.*, 63:247, 2010.

[10] Kenneth A Bamberger and Deirdre K Mulligan. *Privacy on the ground: driving corporate behavior in the United States and Europe*. MIT Press, 2015.

[11] Suzanne Barlyn. Strap on the fitbit: John hancock to sell only interactive life insurance. `https://www.reuters.com/article/us-manulife-financi-john-hancock-lifeins/strap-on-the-fitbit-john-hancock-to-sell-only-interactive-life-insurance-idUSKCN1LZ1WL`.

[12] Erika Beras. Gap and Old Navy wrangle over customer data. `http://www.marketplace.org/2019/03/04/business/big-question-gap-and-old-navy-post-breakup-what-do-customer-data`, Mar 2019.

[13] Emilio Bizzi, Steven E Hyman, Marcus E Raichle, Nancy Kanwisher, Elizabeth A Phelps, Stephen J Morse, Walter Sinnott-Armstrong, Jed S Rakoff, and Henry T Greely. Using imaging to identify deceit: Scientific and ethical questions. *American Academy of Arts and Sciences, Cambridge, MA*, 2009.

[14] Tamara Bonaci. *Security and Privacy of Biomedical Cyber-Physical Systems*. PhD thesis, 2015.

[15] Timothy Brown, Margaret C. Thompson, Jeffrey Herron, Andrew Ko, Howard Chizeck, and Sara Goering. Controlling our brains a case study on the implications of brain-computer interface-triggered deep brain stimulation for essential tremor. *Brain-Computer Interfaces*, 3(4):165–170, October 2016.

[16] Robert M. Chapman and Henry R. Bragdon. Evoked Responses to Numerical and Non-Numerical Visual Stimuli while Problem Solving. *Nature*, 203(4950):1155–1157, September 1964.

[17] Marco D Comerchero and John Polich. P3a and p3b from typical auditory and visual stimuli. *Clinical neurophysiology*, 110(1):24–30, 1999.

[18] Federal Trade Commission. A Brief Overview of the Federal Trade Commission's Investigative and Law Enforcement Authority. `https://www.ftc.gov/about-ftc/what-we-do/enforcement-authority`, 2008.

[19] Judith DeCew. Privacy, Jan 2018.

[20] Judith Wagner DeCew. *In pursuit of privacy: Law, ethics, and the rise of technology*. Cornell University Press, 1997.

[21] L. A. Farwell and E. Donchin. Talking off the top of your head: toward a mental prosthesis utilizing event-related brain potentials. *Electroencephalography and Clinical Neurophysiology*, 70(6):510–523, December 1988.

[22] Virginia Foxx. Preserving employee wellness programs act. `https://www.congress.gov/bill/115th-congress/house-bill/1313`.

[23] Leslie Francis and John G Francis. *Privacy: what everyone needs to know*. Oxford University Press, 2017.

[24] Mario Frank, Tiffany Hwu, Sakshi Jain, Robert Knight, Ivan Martinovic, Prateek Mittal, Daniele Perito, and Dawn Song. Subliminal Probing for Private Information via EEG-Based BCI Devices. *arXiv:1312.6052 [cs]*, December 2013. arXiv: 1312.6052.

[25] Daniel Genkin, Adi Shamir, and Eran Tromer. Rsa key extraction via low-bandwidth acoustic cryptanalysis. In *Annual Cryptology Conference*, pages 444–461. Springer, 2014.

[26] Kevin Granville. Facebook and Cambridge Analytica: What You Need to Know as Fallout Widens. *The New York Times*, May 2018.

[27] Richard W Homan, John Herman, and Phillip Purdy. Cerebral location of international 1020 system electrode placement. *Electroencephalography and Clinical Neurophysiology*, 66(4):376–382, April 1987.

[28] Jeremy Hsu. The strava heat map and the end of secrets. `https://www.wired.com/story/strava-heat-map-military-bases-fitness-trackers-privacy/`, Jan 2018.

[29] Julie C Inness. *Privacy, intimacy, and isolation*. Oxford University Press on Demand, 1996.

[30] Lloyd Karmeier. Rosenback v. Six Flags. `http://www.illinoiscourts.gov/Opinions/SupremeCourt/2019/123186.pdf`, 2019.

[31] Peter H Klopfer and Daniel I Rubenstein. The concept privacy and its biological basis. *Journal of social Issues*, 33(3):52–65, 1977.

[32] Paul C Kocher. Timing attacks on implementations of diffie-hellman, rsa, dss, and other systems. In *Annual International Cryptology Conference*, pages 104–113. Springer, 1996.

[33] Calvin J Kraft and James Giordano. Integrating brain science and law: neuroscientific evidence and legal perspectives on protecting individual liberties. *Frontiers in neuroscience*, 11:621, 2017.

[34] Joseph Lange, Clment Massart, Andr Mouraux, and Franois-Xavier Standaert. Side-Channel Attacks Against the Human Brain: The PIN Code Case Study. pages 171–189, July 2017.

[35] Jianhua Lin. Divergence measures based on the shannon entropy. *IEEE Transactions on Information theory*, 37(1):145–151, 1991.

[36] David T Lykken. The gsr in the detection of guilt. *Journal of Applied Psychology*, 43(6):385, 1959.

[37] Ivan Martinovic, Doug Davies, Mario Frank, Daniele Perito, Tomas Ros, and Dawn Song. On the feasibility of side-channel attacks with brain-computer interfaces. In *Presented as part of the 21st USENIX Security Symposium (USENIX Security 12)*, pages 143–158, Bellevue, WA, 2012. USENIX.

[38] NIH. What is genetic discrimination? `https://ghr.nlm.nih.gov/primer/testing/discrimination`.

[39] US Department of Labor. Wage and hour division (WHD): Employee polygraph protection act (EPPA). `https://www.dol.gov/whd/polygraph/`.

[40] The MacArthur Foundation Research Network on Law and Neuroscience. fMRI and Lie Detection. `http://lawneuro.org/LieDetect.pdf`, 2016.

[41] Independent Press Standards Organization. Editor's Code pf Practice. `https://www.ipso.co.uk/editors-code-of-practice/`, 2018.

[42] Holly Pauzer. 71% of Consumers Prefer Personalized Ads. `https://www.adlucent.com/blog/2016/71-of-consumers-prefer-personalized-ads`.

[43] Scott R Peppet. Unraveling privacy: The personal prospectus and the threat of a full-disclosure future. *Nw. UL Rev.*, 105:1153, 2011.

[44] Andrea Peterson. Holocaust Museum to visitors: Please stop catching Pokémon here. `https://www.washingtonpost.com/news/the-switch/wp/2016/07/12/holocaust-museum-to-visitors-please-stop-catching-pokemon-here/`, Jul 2016.

[45] John Polich. Updating P300: An integrative theory of P3a and P3b. *Clinical Neurophysiology*, 118(10):2128–2148, October 2007.

[46] Lee Rainie. Americans complicated feelings about social media in an era of privacy concerns. `https://www.pewresearch.org/fact-tank/2018/03/27/americans-`

complicated-feelings-about-social-media-in-an-era-of-privacy-concerns/, Mar 2019.

[47] Edith Ramirez, Julie Brill, Maureen K Ohlhausen, Joshua D Wright, and Terrell McSweeny. Data brokers: A call for transparency and accountability. *Federal Trade Commission*, pages 97–100, 2014.

[48] Andrew G Reece and Christopher M Danforth. Instagram photos reveal predictive markers of depression. *EPJ Data Science*, 6(1):15, 2017.

[49] J. Peter Rosenfeld, Bradley Cantwell, Victoria Tepe Nasman, Valerie Wojdac, Suzana Ivanov, and Lisa Mazzeri. A Modified, Event-Related Potential-Based Guilty Knowledge Test. *International Journal of Neuroscience*, 42(1-2):157–161, January 1988.

[50] J. Peter Rosenfeld, Xiaoqing Hu, and Kristine Pederson. Deception awareness improves P300-based deception detection in concealed information tests. *International Journal of Psychophysiology*, 86(1):114–121, October 2012.

[51] J.Peter Rosenfeld, Matthew Soskins, Gregory Bosh, and Andrew Ryan. Simple, effective countermeasures to P300-based tests of detection of concealed information. *Psychophysiology*, 41(2):205–219, March 2004.

[52] J.Peter Rosenfeld, Matthew Soskins, Gregory Bosh, and Andrew Ryan. Simple, effective countermeasures to P300-based tests of detection of concealed information. *Psychophysiology*, 41(2):205–219, March 2004.

[53] Jennifer Lynch and Adam Schwartz. Victory! Illinois Supreme Court Protects Biometric Privacy. https://www.eff.org/deeplinks/2019/01/victory-illinois-supreme-court-protects-biometric-privacy, Jan 2019.

[54] Leslie Scism. New York Insurers Can Evaluate Your Social Media Use–if They Can Prove Why It's Needed. *Wall Street Journal (Online)*, Jan 2019.

[55] SJM Smith. Eeg in neurological conditions other than epilepsy: when does it help, what does it add? *Journal of Neurology, Neurosurgery & Psychiatry*, 76(suppl 2):ii8–ii12, 2005.

[56] Sarah Spiekermann, Jens Grossklags, and Bettina Berendt. E-privacy in 2nd generation e-commerce: privacy preferences versus actual behavior. In *Proceedings of the 3rd ACM conference on Electronic Commerce*, pages 38–47. ACM, 2001.

[57] Potter Stewart and Supreme Court Of The United States. U.S. Reports: Katz v. United States, 389 U.S. 347. `https://www.loc.gov/item/usrep389347/`, 1967.

[58] Samuel Sutton, Margery Braren, Joseph Zubin, and E. R. John. Evoked-Potential Correlates of Stimulus Uncertainty. *Science*, 150(3700):1187–1188, 1965.

[59] Magdalena Szumilas. Explaining odds ratios. *Journal of the Canadian academy of child and adolescent psychiatry*, 19(3):227, 2010.

[60] Dean Takahashi. Valve psychologist explores controlling games directly with your brain. `https://venturebeat.com/2019/03/24/valve-psychologist-explores-controlling-games-directly-with-your-brain/`, Mar 2019.

[61] Margaret Claire Thompson. *Side Effect Mitigation Methods for Closed-Loop Deep Brain Stimulation*. PhD thesis, 2018.

[62] Judith Jarvis Thomson. The right to privacy. *Philosophy & Public Affairs*, pages 295–314, 1975.

[63] Joseph Turow, Michael Hennessy, and Nora Draper. Persistent misperceptions: Americans misplaced confidence in privacy policies, 2003–2015. *Journal of Broadcasting & Electronic Media*, 62(3):461–478, 2018.

[64] Upturn. Civil rights, big data, and our algorithmic future: A september 2014 report on social justice and technology. `https://bigdata.fairness.io/`, Sep 2014.

[65] Jacques J Vidal. Toward direct brain-computer communication. *Annual review of Biophysics and Bioengineering*, 2(1):157–180, 1973.

[66] Ari Ezra Waldman. Designing without privacy. *Hous. L. Rev.*, 55:659, 2017.

[67] Ari Ezra Waldman. *Privacy as Trust: Information Privacy for an Information Age*. Cambridge University Press, 2018.

[68] Samuel D Warren and Louis D Brandeis. Right to privacy. *Harv. L. Rev.*, 4:193, 1890.

[69] D. Wu, V. J. Lawhern, S. Gordon, B. J. Lance, and C. T. Lin. Driver Drowsiness Estimation From EEG Signals Using Online Weighted Adaptation Regularization for Regression (OwARR). *IEEE Transactions on Fuzzy Systems*, 25(6):1522–1535, December 2017.

[70] Rafael Yuste, Sara Goering, Guoqiang Bi, Jose M Carmena, Adrian Carter, Joseph J Fins, Phoebe Friesen, Jack Gallant, Jane E Huggins, Judy Illes, et al. Four ethical priorities for neurotechnologies and ai. *Nature News*, 551(7679):159, 2017.

[71] Lawrence A. Zeidman, James Stone, and Daniel Kondziella. New revelations about hans berger, father of the electroencephalogram (eeg), and his ties to the third reich. *Journal of Child Neurology*, 29(7):1002–1010, 2014.

# Appendix A

# NEUROETHICS QUESTIONNAIRE

## Research Study: Brain-Computer Interface (BCI) Security and Privacy

UNIVERSITY OF WASHINGTON CONSENT FORM
Please read the following consent form and indicate at the bottom that you understand and are willing to complete this research questionnaire. This questionnaire is completely anonymous. If you decide not to participate, please close the form.

* Required

### Investigator

Howard Jay Chizeck
Professor
Electrical Engineering
Bioengineering
chizeck@uw.edu
(206) 221-3591

### Co-Investigator

Katherine Pratt
Graduate Student
Electrical Engineering
inglek@uw.edu
(541) 224-2850

### Co-Investigator

Tim Brown
Graduate Student
Philosophy
timbr@uw.edu
(206) 221-3591

### Researcher's Statement

We are asking you to participate in a research study. Your participation in this research study is entirely voluntary and you may choose to stop participating at any time. There are no adverse effects or penalties if you refuse to participate, or if you discontinue participation after you have started. Please read this disclosure carefully.

### Purpose of the Study

Brain-computer interfaces (BCIs) are a communication channel between the brain and the external environment. Although primarily used for medical research, recently BCIs have gained popularity in gaming, entertainment and marketing industry. Recently it has been shown that some non-medical BCI applications may allow an unauthorized access to the users' private information. For example, nosy or

malicious application developers may use BCIs to infer users' personal information, such as financial information, PIN numbers or passwords, or product preferences. The purpose of this study is to investigate possible privacy implications of BCIs, in particular those used for non-medical applications.

## Study Procedures

At the beginning of the study, we will ask some optional demographic questions. Then we'll ask you a series of questions related to how you interact with technology, and imaginary scenarios. This will allow us to investigate perceptions of this kind of technology from an ethical perspective. There are no physical risks associated with filling out an online survey. The questionnaire data we will obtain from the described experiments will be completely anonymous. It should take you between 10-15 minutes to complete the survey.

## Benefits of the Study

You will not immediately benefit from taking part in this study. We hope, however, that the results of the study will help us assess potential privacy threats of using BCI devices. That will further help us develop more secure and privacy-preserving BCI devices.

## Sources of Funding

The study team and/or the University of Washington are receiving financial support from the Tech Policy Lab, the National Science Foundation, and the Center for Sensorimotor Neural Engineering (CSNE).

## Reporting of Harms

There should be no harms related to filling out this questionnaire. However, if you feel you have experienced a research-related harm please contact one of the investigators listed at the top of this consent form. If you would like to contact the University of Washington Human Subjects Division for questions, concerns, or complaints about the research, questions about your right to obtain information, or to offer input, they can be reached at 206-543-0098 or hsdinfo@uw.edu.

1. **Subject's Statement** *
   Check all that apply.

   ☐ I have read the above information and I volunteer to take part in this research. If I have questions later about the research, I can ask one of the researchers listed above.

## Demographic Questions

Please answer as many questions as you feel comfortable.

2. **Year of Birth**

3. **Gender**

4. **Race/Ethnicity**

5. **Current Occupation**

6. **What is your highest level of education obtained?**
   Mark only one oval.
   - ◯ Did not complete high school
   - ◯ High school degree
   - ◯ Associates degree
   - ◯ Bachelors
   - ◯ Masters
   - ◯ PhD or other professional degree

7. **Do you have significant coursework (major or minor) or work experience in a STEM (Science, Technology, Engineering, Mathematics) related field?**
   Mark only one oval.
   - ◯ Yes
   - ◯ No

8. **If you answered yes to the above question, please list the experience here.**

9. **Years lived in US (if born in the US, just provide your current age)**

10. **If you haven't lived in the US your whole life, indicate the country where you consider you have lived the most outside the US**

## Use of Technology Questions

Please answer as many questions as you feel comfortable about your current technology usage

11. **How many digital devices do you use on a daily basis? (For example, smartphone, desktop computer, laptop computer, personal gaming device, etc.)**
Mark only one oval.

- ◯ 0
- ◯ 1-3
- ◯ 4-6
- ◯ 7-9
- ◯ 10+

12. **Do you use one or more digital assistants? (For example: Siri, Cortana, Alexa, etc.)**
Mark only one oval.

- ◯ No
- ◯ Yes, only one
- ◯ Yes, more than one

13. **On average, how many hours of media (television shows, streaming video, movies, books, comics/graphic novels, etc) do you consume in a week?**
Mark only one oval.

- ◯ None
- ◯ 1-5
- ◯ 6-10
- ◯ 11-15
- ◯ 16-20
- ◯ 20-24
- ◯ 25-30
- ◯ 30-34
- ◯ 35 and above

14. **Do you self-identify as a science-fiction fan?**
Mark only one oval.

- ◯ Yes
- ◯ No

15. **Are you mobility-impaired, or do you use a prosthesis?**
Mark only one oval.

- ◯ Yes
- ◯ No

16. **If you answered yes to the previous question, please list any technologies you use as a result (e.g. motorized wheelchair, prosthetic limb, exoskeleton, etc.)**

## Research Section 1

Brain Computer Interfaces (or BCIs) can record brain activity while an individual is performing different actions (for example, blinking their eyes, playing a video game, or texting on a phone). BCIs are often used to give a user control of a computer using their brain activity.

## Imagine using a BCI-enabled mobile device while sitting on the bus. In this scenario, it would be a serious violation of my privacy if a person sitting behind me on the bus...

Please answer the following questions that complete the above phrase, on a scale from disagree completely to agree completely

17. **... read the content of a text message I sent to my family member** *
Mark only one oval.

| | 1 | 2 | 3 | 4 | 5 | |
|---|---|---|---|---|---|---|
| Disagree Completely | ◯ | ◯ | ◯ | ◯ | ◯ | Agree Completely |

18. **... recorded a video of me typing a text message to my family member which they watched later to find out the content of the text message.** *
Mark only one oval.

| | 1 | 2 | 3 | 4 | 5 | |
|---|---|---|---|---|---|---|
| Disagree Completely | ◯ | ◯ | ◯ | ◯ | ◯ | Agree Completely |

19. **... obtained a recording of my brain activity while I typed a text message to my family member from which the content of my text message could be figured out.** *
Mark only one oval.

| | 1 | 2 | 3 | 4 | 5 | |
|---|---|---|---|---|---|---|
| Disagree Completely | ◯ | ◯ | ◯ | ◯ | ◯ | Agree Completely |

20. **... obtained a recording of my brain activity *as I was planning* to type out a text message to my family member such that the content of the text message could be predicted before I typed it.** *

Mark only one oval.

| | 1 | 2 | 3 | 4 | 5 | |
|---|---|---|---|---|---|---|
| Disagree Completely | ◯ | ◯ | ◯ | ◯ | ◯ | Agree Completely |

21. **... obtained a recording of my brain activity that included my emotional state while typing a message to my family member (e.g., to reveal feelings of inattention, boredom, excitement, anxiety, other states?)** *

Mark only one oval.

| | 1 | 2 | 3 | 4 | 5 | |
|---|---|---|---|---|---|---|
| Disagree Completely | ◯ | ◯ | ◯ | ◯ | ◯ | Agree Completely |

## Research Section 2

Brain Computer Interfaces (or BCIs) can record brain activity while an individual is performing different actions (for example, blinking their eyes, playing a video game, or texting on a phone). BCIs are often used to give a user control of a computer using their brain activity.

## Imagine a scenario where there is a recording app on your phone that can store everything you do on it. You can control this phone with a BCI device. It would be a serious violation of my privacy if an app on my phone...

Please answer the following questions that complete the above phrase, on a scale from disagree completely to agree completely

22. **... downloaded and stored the content of a text message I sent to my family member** *

Mark only one oval.

| | 1 | 2 | 3 | 4 | 5 | |
|---|---|---|---|---|---|---|
| Disagree Completely | ◯ | ◯ | ◯ | ◯ | ◯ | Agree Completely |

23. **... used a keylogger to record me typing out the message in real time, including what was typed and deleted.** *

Mark only one oval.

| | 1 | 2 | 3 | 4 | 5 | |
|---|---|---|---|---|---|---|
| Disagree Completely | ◯ | ◯ | ◯ | ◯ | ◯ | Agree Completely |

24. **... obtained a recording of my brain activity while I typed a text message to my family member from which the content of my text message could be figured out.** *

Mark only one oval.

| | 1 | 2 | 3 | 4 | 5 | |
|---|---|---|---|---|---|---|
| Disagree Completely | ◯ | ◯ | ◯ | ◯ | ◯ | Agree Completely |

25. **... obtained a recording of my brain activity *as I was planning* to type out a text message to my family member such that the content of the text message could be predicted before I typed it.** *

Mark only one oval.

| | 1 | 2 | 3 | 4 | 5 | |
|---|---|---|---|---|---|---|
| Disagree Completely | ◯ | ◯ | ◯ | ◯ | ◯ | Agree Completely |

26. **... obtained a recording of my brain activity that included my emotional state while typing a message to my family member (e.g., to reveal feelings of inattention, boredom, excitement, anxiety, other states?)** *

Mark only one oval.

| | 1 | 2 | 3 | 4 | 5 | |
|---|---|---|---|---|---|---|
| Disagree Completely | ◯ | ◯ | ◯ | ◯ | ◯ | Agree Completely |

## Research Section 3

Brain Computer Interfaces (or BCIs) can record brain activity while an individual is performing different actions (for example, blinking their eyes, playing a video game, or texting on a phone). BCIs are often used to give a user control of a computer using their brain activity.

**Imagine that a BCI device could detect your thoughts or feelings, such as desires for particular foods, attraction to particular movie stars, discomfort with particular political views, or how you're feeling physically and mentally. Please indicate how *willing* you are to share information with the following entities, on a scale from extremely unwilling to extremely willing.**

27. **The person you list as an emergency contact on waivers** *

Mark only one oval.

| | 1 | 2 | 3 | 4 | 5 | |
|---|---|---|---|---|---|---|
| Extremely unwilling | ◯ | ◯ | ◯ | ◯ | ◯ | Extremely willing |

28. **Medical Professional (doctor, nurse, etc.)** *
Mark only one oval.

| | 1 | 2 | 3 | 4 | 5 | |
|---|---|---|---|---|---|---|
| Extremely unwilling | ◯ | ◯ | ◯ | ◯ | ◯ | Extremely willing |

29. **University Researcher (with ethics board approval)** *
Mark only one oval.

| | 1 | 2 | 3 | 4 | 5 | |
|---|---|---|---|---|---|---|
| Extremely unwilling | ◯ | ◯ | ◯ | ◯ | ◯ | Extremely willing |

30. **Governmental Organization** *
Mark only one oval.

| | 1 | 2 | 3 | 4 | 5 | |
|---|---|---|---|---|---|---|
| Extremely unwilling | ◯ | ◯ | ◯ | ◯ | ◯ | Extremely willing |

31. **Non-profit company** *
Mark only one oval.

| | 1 | 2 | 3 | 4 | 5 | |
|---|---|---|---|---|---|---|
| Extremely unwilling | ◯ | ◯ | ◯ | ◯ | ◯ | Extremely willing |

32. **For-profit company** *
Mark only one oval.

| | 1 | 2 | 3 | 4 | 5 | |
|---|---|---|---|---|---|---|
| Extremely unwilling | ◯ | ◯ | ◯ | ◯ | ◯ | Extremely willing |

## Research Section 4

Brain Computer Interfaces (or BCIs) can record brain activity while an individual is performing different actions (for example, blinking their eyes, playing a video game, or texting on a phone). BCIs are often used to give a user control of a computer using their brain activity.

**Imagine that a BCI device could detect your thoughts or feelings, such as desires for particular foods, attraction to particular movie stars, discomfort with particular political views, or how you're feeling physically and mentally. Please indicate how *trustworthy* you believe each of the following entities are obtaining this information, on a scale from extremely untrustworthy to extremely trustworthy**

33. **Family Member** *
Mark only one oval.

| | 1 | 2 | 3 | 4 | 5 | |
|---|---|---|---|---|---|---|
| Extremely untrustworthy | ◯ | ◯ | ◯ | ◯ | ◯ | Extremely trustworthy |

34. **Medical professional** *
Mark only one oval.

| | 1 | 2 | 3 | 4 | 5 | |
|---|---|---|---|---|---|---|
| Extremely untrustworthy | ◯ | ◯ | ◯ | ◯ | ◯ | Extremely trustworthy |

35. **University Researcher (with ethics board approval)** *
Mark only one oval.

| | 1 | 2 | 3 | 4 | 5 | |
|---|---|---|---|---|---|---|
| Extremely untrustworthy | ◯ | ◯ | ◯ | ◯ | ◯ | Extremely trustworthy |

36. **Governmental Organization** *
Mark only one oval.

| | 1 | 2 | 3 | 4 | 5 | |
|---|---|---|---|---|---|---|
| Extremely untrustworthy | ◯ | ◯ | ◯ | ◯ | ◯ | Extremely trustworthy |

37. **Nonprofit company** *
Mark only one oval.

| | 1 | 2 | 3 | 4 | 5 | |
|---|---|---|---|---|---|---|
| Extremely untrustworthy | ◯ | ◯ | ◯ | ◯ | ◯ | Extremely trustworthy |

38. **For-profit company** *
Mark only one oval.

| | 1 | 2 | 3 | 4 | 5 | |
|---|---|---|---|---|---|---|
| Extremely untrustworthy | ◯ | ◯ | ◯ | ◯ | ◯ | Extremely trustworthy |

## Regulation of Neural Technology, part 1

Each of the following individuals/groups plays some role in privacy. Please indicate if they should be doing more or less, or if they should continue to do what they are doing, for each of the scenarios.

## A _____ should play more/same/less of a role in oversight of the development of devices that can record and use your private

**neural information for the purposes of using a BCI, compared to current involvement.**

39. **User** *
Mark only one oval.
- More
- Same
- Less

40. **University researcher with ethics board approval** *
Mark only one oval.
- More
- Same
- Less

41. **Independent Regulatory Organization** *
Mark only one oval.
- More
- Same
- Less

42. **Legislators** *
Mark only one oval.
- More
- Same
- Less

43. **Device Manufacturer** *
Mark only one oval.
- More
- Same
- Less

## Regulation of Neural Technology, part 2

Each of the following individuals/groups plays some role in privacy. Please indicate if they should be doing more or less, or if they should continue to do what they are doing, for each of the scenarios.

**A _____ should play more/same/less of a role in regulating the use of devices that can record and use your private neural information for the purposes of using a BCI when they come to market, compared to current involvement.**

44. **User** *
Mark only one oval.

- More
- Same
- Less

45. **University researcher with ethics board approval** *
Mark only one oval.

- More
- Same
- Less

46. **Independent Regulatory Organization** *
Mark only one oval.

- More
- Same
- Less

47. **Legislators** *
Mark only one oval.

- More
- Same
- Less

48. **Device Manufacturer** *
Mark only one oval.

- More
- Same
- Less

## Regulation of Neural Technology, part 3

Each of the following individuals/groups plays some role in privacy. Please indicate if they should be doing more or less, or if they should continue to do what they are doing, for each of the scenarios.

**A _____ should hold more/same/less responsibility in seeking reparations and compensation if neural information is stolen, or malicious entities are able to elicit private information from you while you are using a BCI, compared to current involvement in similar cases like data breaches.**

49. **User** *
Mark only one oval.

- More
- Same
- Less

50. **University researcher with ethics board approval** *
Mark only one oval.

- More
- Same
- Less

51. **Independent Regulatory Organization** *
Mark only one oval.

- More
- Same
- Less

52. **Legislators** *
Mark only one oval.

- More
- Same
- Less

53. **Device Manufacturer** *
Mark only one oval.

- More
- Same
- Less

## Final Questions

## For each of the following, please indicate if you think recordings of your neural information should be more, equally, or less private than the comparison

54. **Neural information is _____ private than data from a Fitbit or similar exercise tracker.** *
Mark only one oval.

- More
- Equally
- Less

55. **Neural information is _____ private than a record of my personal medical history (e.g. at your doctor's office).** *
Mark only one oval.

- More
- Equally
- Less

56. **Neural information is _____ private than genetic information from a company like 23andMe.** *
Mark only one oval.

- More
- Equally
- Less

57. **Neural information is _____ private than my online shopping history.** *
Mark only one oval.

- More
- Equally
- Less

58. **Neural information is _____ private than my monthly credit card statement.** *
Mark only one oval.

- More
- Equally
- Less

59. **Neural information is _____ private than a journal/diary.** *
Mark only one oval.

- More
- Equally
- Less

## Thank you!

We appreciate your time in filling out this questionnaire. If you have any questions, please feel free to contact the researchers (Howard Chizeck - chizeck@uw.edu, Katherine Pratt - inglek@uw.edu, Tim Brown - timbr@uw.edu) or the University of Washington Human Subjects Division at 206-543-0098 or hsdinfo@uw.edu.

# Appendix B

# STATISTICAL CONSULT ANALYSIS

## Neuro-ethics Survey in Brain-Computer Interface Security and Privacy

*Marlena Bannick*

*06/05/2019*

As Brain-Computer Interface (BCI) devices become more integrated into our lives, there is interest in the ethics around their use. In particular, around how people perceive privacy in relation to these technologies, and how much of a role different entities should play in regulating these technologies. In order to assess attitudes in the general public, a survey was sent to groups at the University of Washington and Oregon Health Sciences University asking questions about key demographics and their perceptions of privacy and regulation in BCI devices.

Seventy-seven individuals responded to the survey. First, we analyze whether or not beliefs about privacy violations change depending on if information is collected by a person or by a BCI-enabled device, and whether or not this belief changes between people with and without mobility impairments. Second, we analyze beliefs about the involvement of various entities (users, device manufacturers, etc.) over stages in the BCI implementation process.

The results indicate that people believe that person-obtained information is more of a privacy violation than BCI-device obtained information, and that people with mobility impairments believe BCI-device obtained information is less of a privacy violation than those without mobility impairments. The results also indicate that people believe users should be more involved in the early stages of the process, such as development, than other entities, and that legislators and device manufacturers should be more involved in the later stages of the process, such as regulation and compensation.

### Data Sources

The data collected in the surveys included demographic information: year of birth, gender, race/ethnicity, occupation, level of education obtained, experience in STEM-related fields, years lived in the U.S., and country of birth. Respondents were also asked about experience with technology such as digital assistants, smartphones, television, and use of prosthesis.

The survey had two main aims that are addressed in this report:

1. Are there differences in perceptions of privacy violations if neural-behavioral information is obtained by a person, or through a BCI-enabled device? Do different groups perceive neural-behavioral data being obtained as a more or less severe privacy violation?
2. Characterize the involvement that people think is necessary for different entities (e.g. user, government, corporation) involved in BCI applications at different stages of the development and roll-out process.

**Aim 1**

For the first aim, the respondents were given a survey with five sets of two questions each. Each question in a set describes a person- or BCI-application (app) related scenario:

- "Imagine using a BCI-enabled mobile device while sitting on the bus. In this scenario, it would be a serious violation of my privacy if a **person** sitting behind me on the bus..."
- "Imagine a scenario where there is a recording app on your **BCI-enabled phone** that can store everything you do. It would be a serious violation of my privacy if the app on my phone...".

The scenarios differ in their ending, with increasing invasiveness. The first of the five types of scenario focuses on reading or downloading content of text messages sent to family. This contrasts with the last of the five scenarios where the person or the BCI-enabled phone obtains recordings of brain activity when the user is in an emotional state.

The responses to each of these questions (for 10 total questions) were on a Likert-scale from 1 (disagree completely that this is a serious violation of privacy) to 5 (agree completely that this is a serious violation of privacy).

The questions for aim 1 are labeled in figures as 'Q1', 'Q2', etc. The key for these questions is given in Table 1.

Table 1: Question Definitions for Privacy Violations in Aim 1

| Question Number | Person Scenario | BCI-enabled Device Scenario |
|---|---|---|
| Preamble | Imagine using a BCI-enabled mobile device while sitting on the bus. In this scenario, it would be a serious violation of my privacy if a person sitting behind me on the bus... | Imagine a scenario where there is a recording app on your phone that can store everything you do on it. You can control this phone with a BCI device. It would be a serious violation of my privacy if an app on my phone... |
| Q1 | ... read the content of a text message I sent to my family member. | ... downloaded and stored the content of a text message I sent to my family member. |
| Q2 | ... recorded a video of me typing a text message to my family member which they watched later to find out the content of the text message. | ... used a key-logger to record me typing out the message in real time, including what was typed and deleted. |
| Q3 | ... obtained a recording of my brain activity while I typed a text message to my family member from which the content of my text message could be figured out. | ... obtained a recording of my brain activity while I typed a text message to my family member from which the content of my text message could be figured out. |
| Q4 | ... obtained a recording of my brain activity "as I was planning" to type out a text message to my family member such that the content of the text message could be predicted before I typed it. | ... obtained a recording of my brain activity "as I was planning" to type out a text message to my family member such that the content of the text message could be predicted before I typed it. |
| Q5 | ... obtained a recording of my brain activity that included my emotional state while typing a message to my family member (e.g., to reveal feelings of inattention, boredom, excitement, anxiety, other states). | ... obtained a recording of my brain activity that included my emotional state while typing a message to my family member (e.g., to reveal feelings of inattention, boredom, excitement, anxiety, other states). |

**Aim 2**

For the second aim, respondents were asked questions about specific entities involved with the BCI roll-out process and whether those entities in a given stage of the process should be "more" or "less" involved than they already are. The entities include: (1) device user, (2) university researcher with IRB approval, (3) independent regulatory organization, (4) legislators, and (5) BCI-device manufacturers. The stages of the process included (1) oversight in development of product, (2) regulating use of product, and (3) reparations or compensation to the user in the event that neural information is stolen.

The responses to each of these questions (for a total of 15 questions) were on a Likert-scale from 1 (less involvement) to 3 (more involvement).

## Statistical Methods and Considerations

The survey had a respectable number of respondents considering the limited number of groups, but is still considered a small sample size for the types of analyses that we conduct here. Most of the groups in the demographic groupings were not large enough for us to include in the analysis, with the exception of mobility impairments.

It is important to account for repeated measures in the analysis, since the same individual answered multiple questions that we will use within the same regression model. Additionally, the dependent variable is ordered on the Likert-scale. To account for these features of the data set, we use ordinal logistic regression with clustering by individual, for both aims 1 and 2. We assume that observations from the same individual are independent, which is probably an incorrect assumption; we account for this misspecification by using robust standard errors in constructing the confidence intervals. We model the cumulative probability of responding less of a privacy violation:

$$\text{logit}(P[Y \leq i|X]) = \alpha + \beta X, \quad i \in 1, ..., n$$

where $i = 1$ represents the cutoff from 1 to 2 (level 0 and 1), etc., $n$ is the number of cutoffs, $Y$ is the ordinal response, and $X$ is the independent variable of interest. To fit this model, we will use the `ordLORgee` function from the R package `multgee`.

For aim 1, the response variable $Y$ has 5 levels (4 thresholds), and represents the level of agreement with a scenario being a privacy violation. The key independent variable $X$ is whether or not the scenario was about a person obtaining the information or a BCI-enabled device obtaining the information; we will call this variable `mode`. We run this analysis separately for each of the 5 question sets/scenarios. We also run this regression including the independent binary variable of mobility impairment (`mobility`) with an interaction term between `mobility` and `mode`. For these two regressions, we are interested in the coefficient that represents the magnitude and direction of association for the independent variable, `mode`, and the interaction term between `mode` and `mobility` on the perceived privacy violation.

For aim 2, the response variable $Y$ has 3 levels (2 thresholds), and represents the level of involvement (more/less) that the entity should have at a given stage. The key independent variables are entity and stage. We will first characterize the level of involvement for each entity across stages separately. Then we will

characterize the level of involvement for each stage separately, moving across the entity levels (1: user to 5: device manufacturer). Then we will include both entity and stage in the same model. Finally, we will include an interaction between entity and stage.

#### Coefficient Interpretation

We are modeling the odds of responding at or below an arbitrary threshold of privacy violation. The $\beta$ coefficients are log odds ratios, so the exponentiated coefficients ($e^{\beta}$) are odds ratios comparing two groups. Thus, the interpretation of the exponentiated coefficients are odds ratios of responding at or below an arbitrary threshold of privacy violation, comparing two groups. The group that is coded as `0` is the reference group, and `1` is the alternative group. An odds ratio greater than 1 means that group `1` has higher odds of responding below a threshold than group `0`. For example, if lower response values mean less of a privacy violation, then an odds ratio greater than 1 indicates that group `1` thinks it is less of a privacy violation than group `0`. Conversely, an odds ratio less than 1 indicates that group `1` thinks it is more of a privacy violation than group `0`. When we include an interaction term, the exponentiated coefficient is a ratio of odds ratios.

## Results and Discussion

#### Aim 1

Before running the analyses, we visualize the responses to the survey questions. Figure 1A shows the number of respondents who answer 1-5 in terms of disagreeing/agreeing with each question being a major privacy violation, broken down by whether it was a person or a BCI-enabled device obtaining the information. We can see that, overall, more respondents agree that it would be a major privacy violation for a person to obtain their information, across all scenarios proposed, as compared to a BCI-enabled device obtaining their information. A regression analysis will help us understand analytically what we see visually by comparing people across all of the categories of responses.

Figure 1A. Responses to Privacy Violation Questions by Mode
1 (low) − 5 (high)

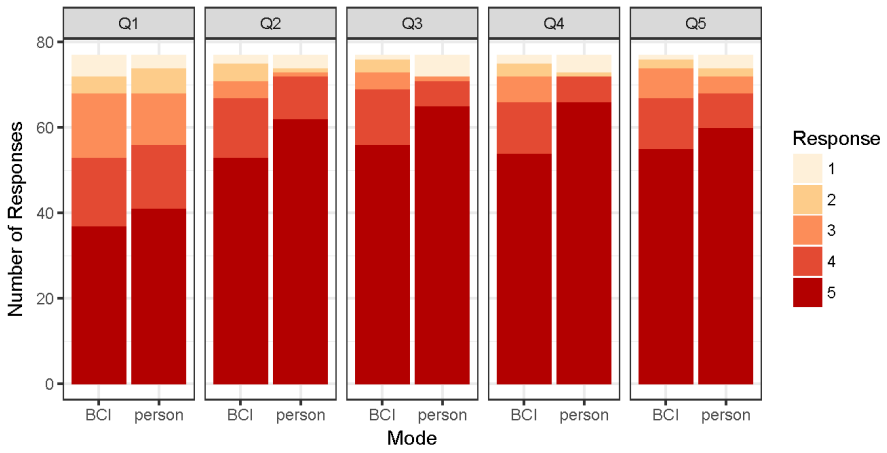


Figure 1B shows the same responses, but this time broken down by people with and without mobility impairments, and broken down by person-obtained information and BCI-obtained information. With the exception of Questions 1 and 2, regardless of whether or not the information was obtained by a person or a BCI-enabled device, more respondents without mobility impairments thought that it was a serious privacy violation compared to people with mobility impairments.

Figure 1B. Responses to Privacy Violation Questions by Mobility and Mode:
1 (low) − 5 (high)

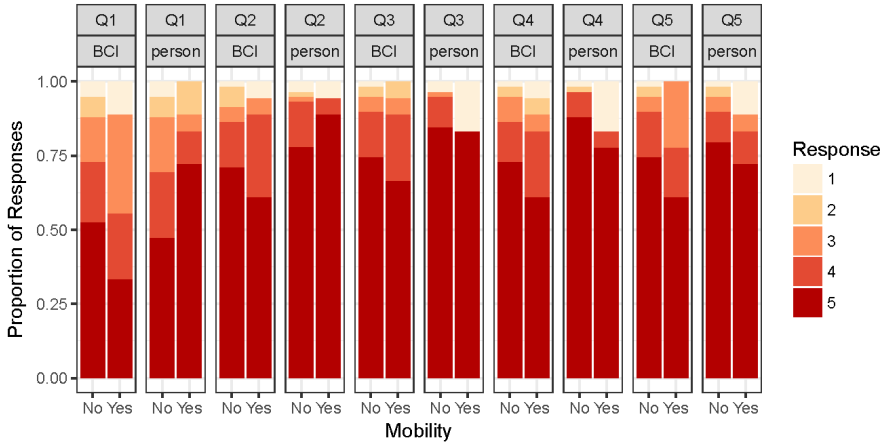


Table 2 shows the resulting odds ratios for responding "less of a privacy violation" by person (alternative)

compared to BCI-enabled device (reference). For all questions, the odds of responding "less of a privacy violation" when information is collected by people are less than when information is collected by a BCI-enabled device. In other words, the odds of responding more of a privacy violation are greater when the information is collected by people. This is consistent with the trend that we see in Figure 1A.

The confidence intervals for Questions 1, 2, 3, and 5 overlap with 1, so they are "not statistically significant". However, the analysis of the survey data was primarily exploratory, i.e. we were not looking to confirm prior hypotheses with rigorous statistical tests, but rather to explore the trends and relationships in the data. As such, we focus on the magnitude and direction of the odds ratios of interest and their confidence intervals, rather than on p-values.

Table 3 shows three odds ratios per question. The first is the odds ratio for responding "less of a privacy violation" for BCI-device obtained information comparing people with mobility impairments to people without mobility impairments. The second is the odds ratio for only people without mobility impairments responding "less of a privacy violation" comparing information collected by person to BCI-device. The third is the ratio of odds ratios of responding "less of a privacy violation" for person versus BCI-devices, comparing people with mobility impairments to people without mobility impairments.

For all but the third question, people with mobility impairments perceive BCI-device obtained information as less of a privacy violation than those without mobility impairments, though none of these odds ratios are significant. In the first question, people without mobility impairments think that person-obtained information is less of a violation than BCI-device obtained information, whereas people with mobility impairments think person-obtained information is more of a violation. In the second question, both people with and without mobility impairments think person-obtained information is more of a violation than BCI-device obtained information. In questions three through five, people without mobility impairments think that person-obtained information is more of a privacy violation than BCI-device obtained information, whereas people with mobility impairments think that person-obtained information is less of a privacy violation than BCI-device obtained information. The confidence intervals for the interaction terms are very wide.

Table 2: Odds Ratio for Mode by Question/Scenario (person v. BCI-application reference)

| **Question** | **Odds Ratio** | **95% CI** |
|---|---|---|
| Q1 | 0.83 | 0.48 - 1.41 |
| Q2 | 0.55 | 0.29 - 1.07 |
| Q3 | 0.77 | 0.42 - 1.42 |
| Q4 | 0.56 | 0.32 - 0.99 |
| Q5 | 0.97 | 0.62 - 1.54 |

**Aim 2**

Figure 2 shows the proportion of respondents answering that entities in a particular stage should be less involved (1), no more or less involved (2), or more involved (3) than they are currently. Every participant responded to every question. We can look across rows, which means looking across entities within the same stage. Or we can look down columns, which means looking across stages within the same entity.

Table 3: Odds Ratios for Mode and Mobility, Interaction

| | Odds Ratio | 95% CI |
|---|---|---|
| **Question 1** | | |
| Mobility Impairment | 1.94 | 0.77 - 4.89 |
| Person v. BCI | 1.18 | 0.67 - 2.08 |
| Mobility : Person Interaction | 0.20 | 0.04 - 0.88 |
| **Question 2** | | |
| Mobility Impairment | 1.34 | 0.44 - 4.12 |
| Person v. BCI | 0.68 | 0.35 - 1.33 |
| Mobility : Person Interaction | 0.37 | 0.05 - 2.52 |
| **Question 3** | | |
| Mobility Impairment | 0.84 | 0.25 - 2.89 |
| Person v. BCI | 0.63 | 0.37 - 1.07 |
| Mobility : Person Interaction | 2.23 | 0.47 - 10.63 |
| **Question 4** | | |
| Mobility Impairment | 1.16 | 0.37 - 3.64 |
| Person v. BCI | 0.43 | 0.26 - 0.73 |
| Mobility : Person Interaction | 2.46 | 0.75 - 8.02 |
| **Question 5** | | |
| Mobility Impairment | 1.41 | 0.43 - 4.70 |
| Person v. BCI | 0.88 | 0.61 - 1.27 |
| Mobility : Person Interaction | 1.40 | 0.39 - 5.03 |

Figure 2. Responses to Level of Involvement by Entity and Stage: 1 (less involved) − 3 (more involved)

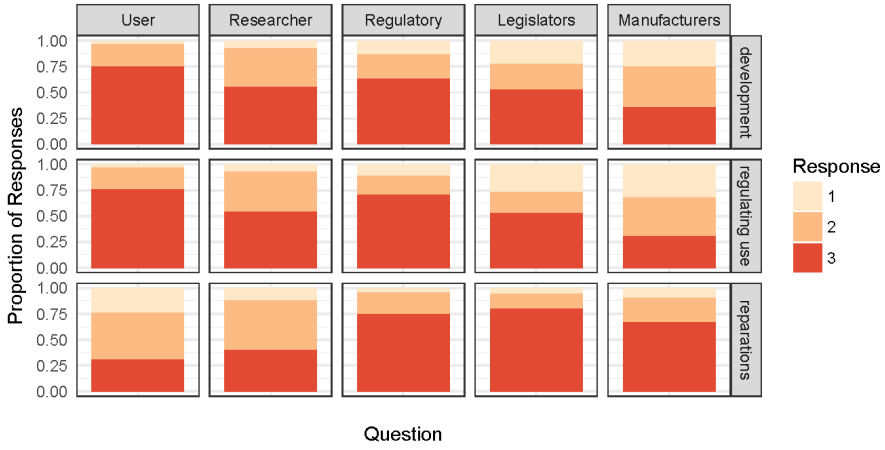


Table 4 shows the results of running ordinal logistic regression models within each entity across stages. The odds ratios represent the ratio of odds of responding lower involvement ($1 < 2 < 3$) across stages. People think that users and university researchers should be less involved over the stages, i.e., relatively more involvement at the development stage than at the regulatory stage than at the reparations stage. The odds ratio for users is statistically significant. Conversely, people think that independent regulatory organizations

Table 4: Odds ratios for relatively more involvement at earlier stages and less at later stages, by entity

| **Entity** | **Odds Ratio** | **95% CI** |
|---|---|---|
| User | 2.48 | 1.78 - 3.45 |
| University Researcher | 1.25 | 0.97 - 1.61 |
| Independent Regulatory Organization | 0.71 | 0.56 - 0.89 |
| Legislators | 0.69 | 0.55 - 0.85 |
| Device Manufacturers | 0.74 | 0.60 - 0.92 |

(IROs), legislators, and device manufacturers should be more involved over the stages, i.e., relatively more involvement at the reparations stage than at the regulatory stage than at the development stage.

Table 5: Odds ratios for relatively more involvement for users and less for device manufacturers, by stage

| **Stage** | **Odds Ratio** | **95% CI** |
|---|---|---|
| Development | 1.46 | 1.28 - 1.67 |
| Regulating Use | 1.54 | 1.35 - 1.74 |
| Reparations | 0.62 | 0.51 - 0.75 |

Table 5 shows the odds ratios for responding lower involvement comparing entities ordered by user to device manufacturer. People think that in both the development stage and the regulatory stage, users should be the most involved entities, and that the device manufacturers should be the least involved entities. Conversely, people think that in the reparations and compensations stage, users should be less involved than researchers, IROs, legislators, and device manufacturers, where device manufacturers are the most involved. All of these results are statistically significant.

As a secondary analysis for aim 2, we can look at the overall feelings of participants for involvement over stages and entities. Table 6 shows the odds ratios in two separate models: Model 1 includes stage and entity, and Model 2 includes stage and entity plus an interaction term between the two. The results of Model 1 suggest that, overall, people think that entities should be more involved in earlier stages than later stages (the `Stage` odds ratio), and that device manufacturers should be less involved than legislators, IROs, researchers, and users, where users are the most involved (the `Entity` odds ratio). Adding an interaction term to Model 2, we see people think that entities like device manufacturers should be relatively more involved as the stages progress compared to legislators, IROs, researchers, and users.

Table 6: Odds ratios for overall less involvement across entity, and overall less involvement across stage

| | **Odds Ratio** | **95% CI** |
|---|---|---|
| **Model 1** | | |
| Entity | 1.14 | 1.04 - 1.24 |
| Stage | 1.18 | 1.05 - 1.32 |
| **Model 2** | | |
| Entity | 2.19 | 1.79 - 2.68 |
| Stage | 2.75 | 2.07 - 3.66 |
| Entity : Stage | 0.71 | 0.65 - 0.78 |

## Conclusions

It is clear from the above analyses that individuals think that levels of involvement should change from where they currently are for various entities and across stages. It is most helpful to look at the tables that stratify by entity and by stage, as they are the most detailed and make the fewest assumptions. It is less clear if people perceive privacy violations as more or less severe by mode, as the odds ratios were closer to 1 and had very large confidence intervals.

As these analyses were exploratory, we performed many hypothesis tests. Therefore, it is likely that we have "false discoveries" for the coefficients that have a significant confidence interval. In a confirmatory analysis, we would need to adjust for many hypothesis tests to prevent this inflated Type I error rate. A limitation of the analysis for aim 2 is that the results can only be interpreted as relative to the respondents' current beliefs about the amount of involvement of the entities at various stages. If similar surveys are conducted in the future, we would recommend adjusting the wording of the questions to represent absolute levels of involvement.