Leveraging Computer Science Expertise to Enhance Tumor Marker Design

Neuralink and BCIs: Addressing Policy Gaps and Societal Concerns

A Thesis Prospectus
In STS 4500
Presented to
The Faculty of the
School of Engineering and Applied Science
University of Virginia
In Partial Fulfillment of the Requirements for the Degree
Bachelor of Science in Computer Science

By
Simon Anderson

November 21, 2024

On my honor as a University student, I have neither given nor received unauthorized aid on this assignment as defined by the Honor Guidelines for Thesis-Related Assignments.

ADVISORS

Joshua Earle, Department of Engineering and Society

Introduction

A month after Neuralink, Elon Musk's pioneering brain-computer interface (BCI) company, successfully implanted its first device, known as the "Link," into a human subject, the wires came loose. As a result, the device malfunctioned, leading to a decrease in data transmission (Gibson, 2024). Reports soon revealed that Neuralink had been aware of these potential issues for years but proceeded with human trials regardless (Levy, 2024). This incident highlights the complex relationship between corporate decision-making, trust in private companies, and the policies governing life-altering technology.

While BCIs mark a major advancement in both medical and cognitive fields, the question remains whether the policies currently regulating their development adequately reflect societal values. These values may include promoting health equity, ensuring privacy, maintaining ethical boundaries between therapy and human enhancement, and much more. My STS project seeks to explore whether the current policy landscape around BCIs aligns with these societal values or exposes gaps and tensions. Specifically, I will investigate how current policies balance the therapeutic promise of BCIs with the risks of cognitive enhancement and corporate control. The key research question driving my project is: Does the current policy landscape surrounding BCIs reflect societal values, and if not, where are the gaps that need addressing?

Technical Project

As a computer science student who does not have a formal technical project as part of my degree, I will be writing about how my skills may be useful for a specific biomedical engineering technical project. The goal of the project is to design a new type of tumor marker that localizes breast cancer tumors effectively without migrating through tissue. Current markers inserted during biopsies often move from their original position. Those that do not migrate must be

inserted during pre-operative procedures, which increases patient discomfort, costs, and infection risks. The team seeks to develop a new marker that can be inserted during a biopsy and stay in place safely until surgery. This project also pertains to STS, as studies have shown disparities in breast cancer diagnosis and treatment outcomes across different racial and socioeconomic groups (Wheeler et al., 2013). Understanding these trends is critical when designing diagnostic tools to ensure that they are effective and equitable for diverse populations.

My technical expertise in computer science could come in handy for this project. Namely, I could contribute by developing image recognition algorithms using skills from my machine learning coursework. These algorithms would analyze medical images to automatically detect and locate tumor markers. This would streamline the process of tracking the marker's position, reduce the need for manual inspection, and allow the team to learn migration patterns of existing markers. This would also be beneficial when testing their new marker to ensure that it works as intended. Additionally, machine learning models could be adjusted to account for variations in imaging quality across different demographic groups, which may be influenced by factors like tissue density, among others (Smilg, 2018). By intentionally using data from a diverse group of people or creating separate models which specialize in certain demographic groups, it would ensure that the technology is equally effective across diverse patient groups.

One of the key ways in which I could help the team avoid potential issues is by integrating the tool I just mentioned early on in the development process. With machine learning models, the team could quickly identify migration patterns or malfunctions in existing markers to refine their approach, as well as in their new marker during testing, allowing them to save time and money before potentially moving onto clinical trials. This also reduces the risk of human error, such as incorrectly identifying the marker's location. On top of this, by using medical data

that includes a representative sample of patients from diverse racial and socioeconomic backgrounds, we could train algorithms that are not only accurate but also fair and unbiased, which may not always be the case with humans.

If I were to be a part of this team, I would definitely learn a lot in the process. I would be exposed to the intersection of healthcare, clinical operations, and software development, which aligns closely with what my job after graduating will entail. I would obviously gain technical knowledge as well by working with machine learning algorithms and applying them to real-world medical images such as ultrasound, MRIs, and X-rays. This would also come with learning more about regulatory considerations and clinical processes, which is relevant to my STS project revolving around Neuralink. In all, working in a team with engineers from other disciplines would expose me to new subjects while also allowing me to refine my skills in my area of expertise, and together, it would enable us to create a practical, patient-centered product which keeps all patients in mind, not just the majority.

STS Project

Background

As mentioned, for my STS project, I will be exploring the alignment of policy surrounding Neuralink/BCIs and societal values as a whole. To answer this question, I will first identify the stakeholders, examine their values, and analyze how these values intersect with the policy frameworks that govern BCIs.

First, we have Neuralink and its owner, Elon Musk, the richest man in the world according to Forbes (Forbes, 2024). As a private company, Neuralink has lots of power over the direction of BCIs, allowing them to prioritize either therapeutic or enhancement applications. In the enhancement case, it could eventually take forms such as Link apps which enable gaming,

strengthened cognitive abilities, or sensory enrichment (Maynard & Scragg, 2019). Musk's leadership style, marked by a willingness to push regulatory limits, raises important questions about how corporate values, such as profit and market expansion, might be enabled by lack of applicable policy.

Next, we have patients and people with disabilities. In this area, BCIs hold the most promise for therapeutic uses. Neuralink's current focus is on enabling people with quadriplegia to control computers using thoughts, with ambitions to restore vision, motor function, and speech (Neuralink, n.d.). These medical uses align with societal values around health equity and improving quality of life. However, if enhancement becomes appealing to a wealthier demographic or more profitable, therapeutic applications could lose priority.

Wealthy individuals seeking enhancement represent another stakeholder group. BCIs used for enhancement could deepen inequalities by creating advantages for those who can afford them. As the Royal Society states, "If cognitive enhancement confers a long-term advantage to users who can afford it, this increases inequity within generations; if those users are then better able to afford enhancement for their children, disadvantage is multiplied across generations" (The Royal Society, 2019). This raises critical questions about fairness, equity, and whether current policies adequately address these risks.

Finally, policymakers and regulatory bodies are crucial stakeholders. Privacy laws like HIPAA in the U.S. and GDPR in the EU regulate neural data, while agencies like the FDA oversee BCI approval. However, these policies may not fully address the ethical challenges BCIs present, especially when it comes to enhancement. Some experts argue that the FDA may not be equipped to handle the subjective, value-laden assessments required for BCI regulation (Binkley et al., 2021). Also, privately held companies such as Neuralink are not obligated to provide

details about these regulatory interactions, leaving a lot of details out of view from the public (Dickey, 2023). With this in mind, it could be useful to explore whether current policies reflect societal concerns around privacy, equity, and health or leave critical issues unaddressed.

Research Methods

To address my research question, I will combine policy analysis and philosophy to explore both the practical and moral implications of BCI development at Neuralink.

On the policy side, I will examine existing frameworks governing BCI production, use, and corporate power. This includes FDA medical device regulations and the EU's GDPR, allowing me to identify risks and policy gaps using insights from the *Journal of Law and the Biosciences*, where they have begun to explore this topic (Rainey et al., 2020). I will assess whether these policies provide sufficient regulatory coverage to address the technological and social risks posed by BCIs.

From the philosophical side, I am looking to explore the ethical questions surrounding two main areas: human enhancement and corporate power. To look into the ethics of human enhancement, I plan to read texts from key thinkers in this area such as Julian Savulescu or Allen Buchanan. I can also read articles which summarize and compare different perspectives on human enhancement, such as *The Ethics of Human Enhancement* written by Alberto Giubilini and Sagar Sanyal (Giubilini & Sanyal, 2015). To ensure a balanced perspective, I will incorporate viewpoints from scholars like M. A. Torres and L. Gebru, who critically examine the ethical implications of advanced AI and human enhancement. In their research, Torres and Gebru explore how motivations for creating such technologies may be influenced by historical and societal biases, including those rooted in discriminatory traditions, leading to potentially harmful outcomes for marginalized communities (Gebru & Torres, 2024). For corporate power, I will use

Zuboff's *The Age of Surveillance Capitalism*, which discusses how technological companies accumulate power by monetizing human behavior, raising ethical concerns about concentrated control over neural data (Zuboff, 2019).

Using these two approaches, policy and philosophy, I hope to provide a well-rounded examination of Neuralink's influence over the future of BCIs. This analysis will not only identify existing gaps but could also guide suggestions on implementing ethical and equitable policies that better align with societal desires..

STS Frameworks

Going hand in hand with my research methods, I will be leveraging two main STS frameworks in my writing: ethics/philosophy and policy/risk analysis. These frameworks provide a way to assess both the practical and social dimensions of Neuralink's BCI development.

As mentioned, the ethics and philosophy framework will allow me to analyze the moral implications of both human enhancement and corporate power. It focuses on whether the use of BCIs aligns with societal values and how it could affect fairness and equity. Also, by examining ethical concerns related to concentrated control over neural data, I can explore whether a single company like Neuralink can be trusted with such power. To address these questions, I plan to use care ethics as my primary philosophical framework. Care ethics emphasizes the importance of relational and contextual considerations in moral decision-making, focusing on how care and responsibility should guide societal choices (Sander-Staudt, n.d.). This perspective will be useful in analyzing how policies and practices related to neurotech and BCIs could prioritize the well-being and needs of individuals and communities.

The policy and risk analysis framework will help me evaluate regulatory gaps and risks associated with BCIs. As discussed, I will analyze existing regulations and laws on medical

devices and data privacy to see if they adequately address both the technological risks and the societal implications of BCIs. This framework is also useful in identifying relationships between the different groups of stakeholders and exploring who might benefit from certain groups being able to interpret and manage these risks. As noted by Sheila Jasanoff, regulatory decisions often reveal conflicts between scientific, corporate, and political interests, highlighting the struggle to balance innovation with accountability (Jasanoff, 1987). This framework will help me address both the limits of regulations and opportunities to align policies with societal needs.

Timeline

In early December, I will focus on primary sources, including policy documents and Neurlink's publications, and begin with initial analysis. During the break from mid-December to mid-January, I will focus on organizing what I have found so far, outlining key sections, and beginning to review secondary sources about ethics, corporate power, and values. Once back from break in January, I will continue this review and see if there are any gaps in my research that still need to be addressed. In early February, I will complete the first full draft of my research paper and use the rest of the semester for review, revisions, and any additional research that becomes necessary.

Key Texts

To guide my research on how policies surrounding brain-computer interfaces (BCIs) align with societal values, I have identified four main resources– some of which I have already briefly touched on.

The first is *The Age of Surveillance Capitalism* by Shoshana Zuboff, which explores how technology companies profit from personal data, leading to new forms of corporate power that threaten individual privacy and autonomy (Zuboff, 2019). This text will help me analyze how

companies like Neuralink may exercise power over neural data, raising ethical concerns about privacy, autonomy, and corporate control.

Another key text is Sheila Jasanoff's *Designs on Nature: Science and Democracy in Europe and the United States*. Jasanoff argues that science and technology are shaped by social values and that policies reflect societal priorities (Jasanoff, 2005). This book will provide insight into how regulatory frameworks are developed, helping me assess whether current policies on BCIs align with societal values like privacy and health equity, or if gaps and risks remain.

Two more sources which will be considered key texts for my research are two primary sources: HIPAA and GDPR. HIPAA, the Health Insurance Portability and Accountability Act, is a U.S. law focused on protecting individuals' medical privacy and security (HIPAA Administrative Simplification, 2013). GDPR is the General Data Protection Regulation, and it is a European Union regulation offering broader protections by giving individuals control over their personal data and setting strict guidelines for companies handling their data (General Data Protection Regulation, n.d.). Since Neuralink's technology collects highly sensitive neural data, these regulations will be helpful in determining whether current policies adequately protect privacy and autonomy or if additional safeguards are needed.

Finally, I will use Neuralink's website, which features blogs and press releases, to understand the company's priorities and vision. While Neuralink emphasizes therapeutic goals, such as restoring motor function and vision, they also reference cognitive enhancement, raising questions about future directions (Neuralink, n.d.). This analysis will help me evaluate how the company frames its technology and whether this aligns with societal concerns around privacy, equity, and human enhancement.

References

Binkley, C. E., Politz, M. S., & Green, B. P. (2021). Who, If Not the FDA, Should Regulate Implantable Brain-Computer Interface Devices? AMA Journal of Ethics, 23(9), 745–749. https://doi.org/10.1001/amajethics.2021.745

Dickey, J. (2023). From Lab to Legislation: Brain-Machine Interfaces and Privacy Considerations | Bipartisan Policy Center. Retrieved October 11, 2024, from https://bipartisanpolicy.org/blog/from-lab-to-legislation-brain-machine-interfaces-and-privacy-considerations/

Forbes Real Time Billionaires List—The World's Richest People. (2024). Forbes. Retrieved October 11, 2024, from https://www.forbes.com/real-time-billionaires/

General Data Protection Regulation (GDPR) – Legal Text. (n.d.). General Data Protection Regulation (GDPR). Retrieved October 14, 2024, from https://gdpr-info.eu/

Gebru, T., & Torres, É. P. (2024). The TESCREAL bundle: Eugenics and the promise of utopia through artificial general intelligence. First Monday. https://doi.org/10.5210/fm.v29i4.13636

Gibson, K. (2024, May 9). Neuralink brain-chip implant encounters issues in first human patient—CBS News. CBS News. https://www.cbsnews.com/news/elon-musk-neuralink-first-implant-problem/

Giubilini, A., & Sanyal, S. (2015). The Ethics of Human Enhancement. Philosophy Compass, 10(4), 233–243. https://doi.org/10.1111/phc3.12208

HIPAA Administrative Simplification. (2013, March 26). U.S. Department of Health and Human Services.

https://www.hhs.gov/sites/default/files/ocr/privacy/hipaa/administrative/combined/hipaa-simplification-201303.pdf

Jasanoff, S. (2005). Designs on Nature: Science and Democracy in Europe and the United States. Princeton University Press. http://ebookcentral.proquest.com/lib/uva/detail.action?docID=713600

Jasanoff, S. S. (1987). Contested Boundaries in Policy-Relevant Science. Social Studies of Science, 17(2), 195–230. https://doi.org/10.1177/030631287017002001

Levy, R. (2024, May 15). Exclusive: Musk's Neuralink has faced issues with its tiny wires for years, sources say. Reuters. https://www.reuters.com/technology/musks-neuralink-has-faced-issues-with-its-tiny-wires-years-sources-say-2024-05-15/

Maynard, A. D., & Scragg, M. (2019). The Ethical and Responsible Development and Application of Advanced Brain Machine Interfaces. Journal of Medical Internet Research, 21(10), e16321. https://doi.org/10.2196/16321

Neuralink. (n.d.). Neuralink—Pioneering Brain Computer Interfaces. Neuralink. Retrieved October 11, 2024, from https://neuralink.com/

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

Sander-Staudt, M. (n.d.). Care Ethics | Internet Encyclopedia of Philosophy. Retrieved November 20, 2024, from https://iep.utm.edu/care-ethics/

Smilg, J. S. (2018). Are you dense? The implications and imaging of the dense breast. SA Journal of Radiology, 22(2), 1356. https://doi.org/10.4102/sajr.v22i2.1356

The Royal Society. (2019). Blurring lines between mind and machine. The Royal Society. Retrieved October 11, 2024, from https://royalsociety.org/-/media/policy/projects/ihuman/report-neural-interfaces.pd

Wheeler, S. B., Reeder-Hayes, K. E., & Carey, L. A. (2013). Disparities in Breast Cancer Treatment and Outcomes: Biological, Social, and Health System Determinants and Opportunities for Research. The Oncologist, 18(9), 986–993. https://doi.org/10.1634/theoncologist.2013-0243

Zuboff, S. (2019). The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power—Book—Faculty & Research—Harvard Business School. Retrieved October 12, 2024, from https://www.hbs.edu/faculty/Pages/item.aspx?num=56791