

5-15-2020

An Inventory of Existing Neuroprivacy Controls

Dustin Steinhagen
Dakota State University

Houssain Kettani
Dakota State University

Follow this and additional works at: <https://scholar.dsu.edu/ccspapers>



Part of the [Databases and Information Systems Commons](#), and the [Other Computer Sciences Commons](#)

Recommended Citation

Steinhagen, Dustin and Kettani, Houssain, "An Inventory of Existing Neuroprivacy Controls" (2020).
Research & Publications. 10.
<https://scholar.dsu.edu/ccspapers/10>

This Article is brought to you for free and open access by the Beacom College of Computer and Cyber Sciences at Beadle Scholar. It has been accepted for inclusion in Research & Publications by an authorized administrator of Beadle Scholar. For more information, please contact repository@dsu.edu.

An Inventory of Existing Neuroprivacy Controls

Dustin Steinhagen, Houssain Kettani

The Beacom College of Computer and Cyber Sciences
Dakota State University, Madison, South Dakota, USA
dustin.steinhagen@trojans.dsu.edu; houssain.kettani@dsu.edu

ABSTRACT

Brain-Computer Interfaces (BCIs) facilitate communication between brains and computers. As these devices become increasingly popular outside of the medical context, research interest in brain privacy risks and countermeasures has bloomed. Several neuroprivacy threats have been identified in the literature, including brain malware, personal data being contained in collected brainwaves and the inadequacy of legal regimes with regards to neural data protection. Dozens of controls have been proposed or implemented for protecting neuroprivacy, although it has not been immediately apparent what the landscape of neuroprivacy controls consists of. This paper inventories the implemented and proposed neuroprivacy risk mitigation techniques from open source repositories, BCI providers and the academic literature. These controls are mapped to the Hoepman privacy strategies and their implementation status is described. Several research directions for ensuring the protection of neuroprivacy are identified.

CCS Concepts

• Security and privacy → Human and societal aspects of security and privacy → Privacy protections

Keywords

Brain-Computer Interfaces; Brain Hacking; Neural Data Protection; Neurosecurity; Neuroprivacy; Privacy Controls

1. INTRODUCTION

Privacy concerns have come to the fore of public consciousness in recent years, with strict privacy regulations such as the California Consumer Privacy Act and the European Union's General Data Protection Regulation making headlines and privacy controversies such as the Snowden revelations and Facebook's Cambridge Analytica scandal roiling popular trust in government and commercial institutions. In recent decades, researchers have shown heightened interest in anticipating and documenting privacy threats to the human brain and mind involving the use of Brain-Computer Interfaces (BCIs). There is much discussion of neurosecurity and neuroprivacy risks and controls, although it is unclear what is being done to develop and make widely available privacy controls for protecting the privacy of people who utilize BCIs.

This paper seeks to summarize the current state of neuroprivacy controls implementation, utilizing the Hoepman privacy strategies

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.
ICISDM 2020, May 15-17, 2020, Hawaii, HI, USA
© 2020 Association for Computing Machinery (ACM).
ISBN 978-1-4503-7765-2/20/05...\$15.00

<https://doi.org/10.1145/3404663.3404664>

as its control taxonomy [21]. Although brain security and privacy scholars have identified many privacy controls for the brain, there has yet to be an inventory of theoretical and implemented neuroprivacy controls that are classified into privacy control families. This study found that few privacy controls have moved beyond their theoretical formulations and that no open source solutions exist for individuals wanting to protect themselves from brain hackers. A novel classification and inventory of BCI privacy controls is offered, along with discussion of research gaps in developing compensating controls for neuroprivacy risk.

Brain-Computer Interfaces (BCIs) are devices that enable bidirectional or unidirectional communication between brains and computers [33]. These devices may or may not be additionally classifiable as neuroprosthetics depending on whether they are integrated into the neural circuitry of the organism [19]. Applications of BCIs range from diagnosing diseases and restoring bodily functions to remotely controlling robots, lie detection, authentication and gaming, among others [33]. Electroencephalography (EEG) is a low-cost and non-invasive method for obtaining neural data, in use by various consumer BCI headset providers such as Muse, NeuroSky and Emotiv. Various other methods for obtaining neural data as well as methods for stimulating the brain with BCIs were discussed in [4].

Brain-computer interfaces can be classified according to the direction(s) communication flows – BCIs can record from the brain, stimulate the brain or both record from and stimulate the brain [33]. These interfaces can also be ranked according to their invasiveness in relation to the brain – invasive BCIs are situated within the brain, semi-invasive BCIs on the brain surface or nerves and non-invasive BCIs outside of the skin of the head and skull [33]. A design classification for BCIs – between passive, active, reactive and hybrid – also exists, which makes distinctions according to whether the initiator of the neural data acquisition is the device or the person and the nature of their interaction [4]. A novel five-phase BCI cycle was proposed that explains the high-level signal processing steps involved with neural data acquisition and neural stimulation, which is useful for neurosecurity threat modeling [4].

Privacy is difficult to define, but at a high level can be thought of as “the relief from a range of kinds of social friction” [35]. Adapting this definition to BCIs, neuroprivacy can be thought of as relief from social friction stemming from the processing of neural data. The original neuroprivacy article defined neuroprivacy as privacy concerns of neurodiagnostic and neuroimaging techniques [31]. “Privacy of thoughts and feelings” has been identified as one of the seven types of privacy, inspired by the privacy issues of neurotechnology [17]. Another author has identified issues related to BCI-enabled blackmail and decisional interference (loss of autonomy), which are privacy harms within Solove's taxonomy [23]. Since insecurity is a privacy harm, neurosecurity issues that result in social friction for the data subject can also be thought of as neuroprivacy issues.

Neuroprivacy concerns represent a unique and pressing challenge for privacy professionals as the mind becomes ever more connected and discernible with advances in neurotechnology. There is uncertainty as to whether data protection regimes can adequately address neuroprivacy concerns [20], and other scholars have identified a lack of basic privacy protections such as applications having excessive access to personal information in popular BCI headsets [39]. There was one application that sent users' raw EEG data to cloud storage, potentially allowing unknown parties to extract sensitive personal information from users' brainwaves at some point in the future. Obtaining informed consent from the users of BCIs is particularly difficult due to previously collected EEG data having the potential to later be processed in novel ways that allows extraction of information that was not possible before due to advancements in the methods used to process EEG signals [37]. The neural data extracted with BCIs is inherently personal – each person's brainwaves are unique [37] and change little over their lifetime [14] and may contain sensitive information such as religiosity [25], drowsiness levels [40] and “guilty knowledge” [33], among other aspects of cognition.

The original neurosecurity paper defined neurosecurity as “the protection of confidentiality, integrity, and availability of neural devices from malicious parties with the goal of preserving the safety of a person's neural mechanisms, neural computation, and free will” [13]. Another researcher has proposed the term “neurocrime” to refer to criminal activities that involve neural information [23]. “Brain spyware” has been demonstrated in the lab, illustrating that brain hackers can deduce private information such as banking information, recognized faces, PIN numbers, location of residency and month of birth from the brains of BCI users by showing them visual stimuli and using a machine learning model to detect familiar information based on the brain's response [28]. This type of attack can also be performed subliminally, without the conscious awareness of the target [18]. Brain hacking scholars are anticipating the growing significance and complexity of these privacy and security threats [4], as novel architectures for neurotechnology enable technology-assisted telepathy [7], using thoughts to remote-control animals [41] and the emergence of a global “Internet of Neurons” [34].

Hoepman distilled the privacy design literature into eight distinct kinds of strategies made up of 26 tactics, which address various high-level requirements of data protection regulation [10, 11, 21, 22]. His privacy design strategies and tactics have been proposed as a potential control taxonomy for mitigating privacy risks [11]. Although intended to categorize privacy design patterns, Hoepman envisioned that his taxonomy could be used to analyze the privacy impact of information systems more generally [21], suggesting that it is an appropriate framework for understanding neuroprivacy risks and controls.

The first four strategies – Minimize, Separate, Abstract and Hide – are data-oriented strategies, dealing with the architectural considerations of privacy described in [36]. Architectures that support less identifiability and less centrality promote stronger privacy. Identifiability is the ease with which personal data can be linked to a natural person, while centrality is the degree to which an organization's network systems enable control over the data subject's information [36]. The tactics within these four strategies can be thought of as specific techniques for reducing centrality and identifiability in an organization's architecture. The remaining four strategies – Inform, Control, Enforce and Demonstrate – are process-oriented, or privacy-by-policy, strategies [21]. They are organizational processes and procedures for promoting strong privacy [22]. Inform and Control are strategies that focus on

empowering the data subject, whereas Demonstrate and Enforce focus on the role of the data controller in maintaining privacy protections.

The existing brain hacking literature identifies numerous privacy and security controls that could be deployed to mitigate risks related to the use of BCI technology. However, these controls have not been organized under a privacy control taxonomy such as the Hoepman strategies. The rest of this paper is as follows: Section 2 discusses the methodology for inventorying the neuroprivacy controls; Section 3 contains the results; Section 4 discusses the results, while concluding remarks are in Section 5.

2. METHODOLOGY

For a piece of research to be included in this inventory, it had to contain at least one concrete, specific action an individual or organization could take to protect neuroprivacy; be from an academic database; be available for reading, either through library subscriptions or be openly available for download; and be readily traceable from the neurosecurity and neuroprivacy literature or from searches for relevant keywords in the titles of the articles. Only controls mentioned within the context of BCIs were included in the inventory – the neuroprosthetics literature was not included, such as the neurosecurity controls mentioned in [19].

Controls were classified into four categories: academic, commercial, open source and theorized. Academic controls have been demonstrated in research labs; commercial controls were countermeasures that were known to be in use by BCI providers; open source controls were tools that were hosted on openly available repositories; and theorized controls were those mentioned in the literature, but had not been implemented in academia, in industry or by individuals maintaining open source solutions. For theorized controls, only the earliest citation was used, and all citations were presented if multiple independent researchers proposed or implemented the control during the same year. The Hoepman strategies were loosely interpreted when possible, which included counting security tools that could be used to indirectly protect personal data as privacy tools. Security controls that were not classifiable as Hoepman controls were given a “Non-Hoepman” security control designation. The Smartphone Brain Scanner openPDS system described in [37] was accounted for as a collection of separate, constituent privacy controls rather than one large privacy control that spanned multiple Hoepman tactics and families.

Mentions of security or privacy controls for BCI devices on certain code repositories and commercial BCI provider websites were also included in this study. The following BCI provider domains were searched for security, privacy and data protection in the context of BCIs or neural data:

- BioSemi (<https://www.biosemi.com>)
- Emotiv (<https://www.emotiv.com>)
- Halo Neuro (<https://www.haloneuro.com>)
- Kokoon (<https://kokoon.io>)
- Muse (<https://choosemuse.com>)
- MyndPlay (<https://myndplay.com>)
- NeuroOptimal (<https://neurooptimal.com>)
- NeuroSky (<http://neurosky.com>)
- OpenBCI (<https://openbci.com>)

These providers' privacy policies and other web pages from the domains listed above were analyzed for any sort of privacy protections that specifically dealt with neural data or BCI device privacy protections. Only four of the nine providers listed above dealt specifically with neuroprivacy issues in their privacy

notices: Emotiv, Kokoon, Muse and NeuroSky. For NeuroSky, neural data processing was discussed in the privacy notice of the Effective Learner App, but not in the main privacy notice. With regards to all privacy notices, protections that explicitly generalized to covering all personal data of which EEG data was a specified type were included in the inventory by default, even if it was unlikely that those protections applied to neural data. For instance, a BCI provider that stated that EEG data is personal data and that they will allow the data subject to correct inaccuracies in their personal data were classified in the inventory as allowing the correction of EEG data unless there was an explicit exception for EEG data. Thus, for the purposes of this paper, all privacy notices were interpreted literally. In addition to BCI provider websites, the following code repositories were searched for mention of open source privacy or security tools that are explicitly for protecting EEG data or BCI devices:

- GitHub (<https://github.com>)
- GitLab (<https://gitlab.com>)
- BitBucket (<https://bitbucket.org>)
- SourceForge (<https://sourceforge.net>)
- LaunchPad (<https://launchpad.net>)

3. RESULTS

This section begins with a general overview of the findings, followed by six subsections describing the controls, sorted by their Hoepman classifications. In total, 94 neuroprivacy controls were identified from the literature, BCI provider websites and in code repositories. Figures 1 through 3 were generated from this collected data and were designed to answer the following research questions:

- What proportion of neuroprivacy controls have implementations and what is the relative distribution of implementation types?
- What is the relative control distribution among the eight Hoepman privacy strategies?
- What is the proportion of each control group that has been implemented?

Figure 1 illustrates the relative distribution of controls based on whether they were theorized, academic, commercial or open source. This inventory revealed that most neuroprivacy controls are theoretical, with 30% in use by BCI providers. There were two controls, blockchain and Secure Multiparty Computation, that have been implemented in the lab [1-3]. There was a single open source solution, Open Brain Consent, which is a template for obtaining consent from research participants for research that involves neural data collection [32]. Figure 2 shows the relative distribution of the 94 controls according to their Hoepman classifications. There were seven controls that could not be classified under the Hoepman strategies, so were given the “Security” designation. The most represented techniques were those in the Hide and Demonstrate strategies, making up 18% and 22% of the controls respectively. Enforce, Separate and Abstract were the least represented controls, with five or fewer controls each. Figure 3 shows the relative distribution of neuroprivacy controls that have been implemented. Implemented controls were any controls that were not categorized as theorized. In total, just under 30% (28) of the 94 controls identified were implemented. Of the privacy strategies, only the Abstract strategy lacked any implemented controls, and the Enforce strategy had the highest proportion of implemented controls (80%).

There were seven controls identified that were not classifiable as Hoepman controls, none of which have been implemented in the

context of BCIs. Adversarial training, architecture modifications, defense distillation and ensemble method are techniques for hardening machine learning models against exploitation [4]. These techniques would involve protecting the algorithm(s) responsible for decoding intentions in neural data acquisition and those responsible for encoding neural firing patterns in neurostimulators. Other proposed neurosecurity controls included utilizing robust programming languages, malware visualization and compilation techniques and options [4].

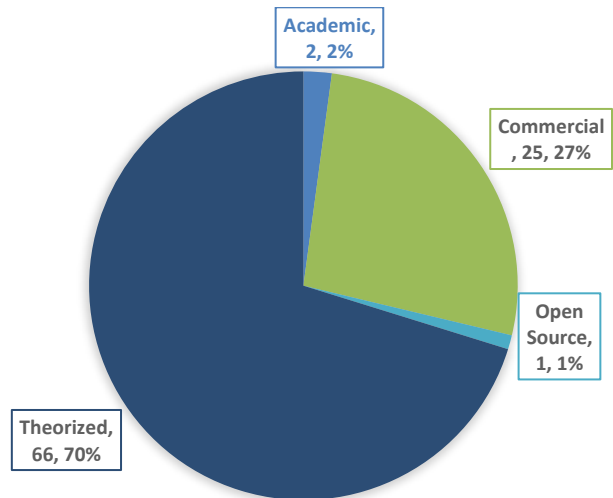


Figure 1. Distribution of Existing Neuroprivacy Controls by Type




Figure 2. Distribution of Controls by Hoepman Strategy




Figure 3. Distribution of Controls by Implementation Status

The Minimize strategy limits as much as possible the processing of personal data [11]. Of the twelve Minimize controls identified, only one was implemented: BCI providers stating in their privacy notices that personal data is retained only as long as necessary, which presumably applies to neural data as well [26, 29, 30]. A similar and cheaper control would be deleting neural data regularly as it becomes unnecessary [6]. Limiting retention of neural data was not universal among the BCI providers, as Emotiv explicitly stated in their privacy notice that they retain EEG data indefinitely [15]. Minimize can also be realized by selectively choosing what personal data to process or by refusing to process personal data. The use of directional antennas and low transmission power could limit the processing of personal data in physical space [4]. Utilizing whitelists, running apps with lowest privileges and designing platforms to only have necessary features are methods for being selective about which neural data will be processed and limiting the processing to relevant parties and applications [4]. Methods for not processing certain neural data included blacklists [4], traffic filtering [34] and input sanitization and validation [34]. Safe APIs and languages were proposed as a way to automate input sanitization and validation [4]. Most of the proposed Minimize controls are familiar from the cybersecurity field, with the exception of the BCI Anonymizer [5, 9]. This tool has been proposed as a way to remove private information from EEG data before they are stored or transmitted [5], although it has not been invented as suggested by the abandoned status of the BCI Anonymizer patent application [9]. Several issues with the BCI Anonymizer idea have been identified, including resource constraints in BCI devices, lack of access to proprietary algorithms, lack of a clear method for separating private information from intentions and a general lack of any implementation details [39].

Separate involves isolating or distributing the processing of personal data as much as possible to prevent correlation [11]. Only five Separate controls were identified in this inventory. Keeping the data subject’s pseudonymization ID separate from their neural data was the only Separate control with a real-world implementation [15]. Other proposed techniques for promoting isolation in BCIs included application sandboxing [34] and segmented application architectures [4]. Suggestions for

distributing neural data processing across physical locations included utilizing an external device for authorizing neural stimulation [4] and keeping EEG data under the control of the end user through the use of a locally hosted Personal Data Store (PDS) [37]. Abstract techniques limit the detail in which personal data is processed [11]. This strategy had the fewest proposed controls of any strategy and was the only one that did not have any real-world implementations. Adding noise to EEG data before applications can process it could decrease the risk of personal data leaks [28]. Differential privacy is a specific application of noise addition that could be deployed in BCIs to deidentify brainwaves [39]. Personal Data Stores could be equipped with capabilities for aggregate computation of neural data across multiple PDS instances and summarizing EEG data into high-level attributes by reducing the dimensionality of the data [37].

The Hide strategy reduces the identifiability of personal data and prevents it from becoming public or known [11]. This strategy had a total of six implementations, five commercial and one academic. Emotiv encrypts neural data at rest and in transit [15]. There were two BCI providers that anonymize neural data by removing any associated identifying information [29, 30]. Emotiv utilizes pseudonymization, which involves processing neural data with a unique identifier rather than directly attributing the neural data to the data subject [15]. NeuroSky, Muse and Kokoon limit access to neural data based on a “need to know” [26, 29, 30]. Secure Multiparty Computation was demonstrated in the lab to allow estimation of driver drowsiness without exposing any individual user’s EEG data [1, 2]. Proposed Hide techniques included fine-grained context-based access control [27], maintaining different permission granularity [4] and restricted APIs [28], all of which could be used to limit access to neural data. Functional encryption and homomorphic encryption have been suggested as potential methods for obfuscating raw EEG data or the extracted features [39]. Application hardening could be used to make BCI software more difficult to reverse engineer, increasing the difficulty of brain exploitation [4]. Mix techniques reduce the risk of unwanted correlation and could include randomization [39], utilizing mix networks such as The Onion Router [34] and using spread spectrum for enhanced wireless security [4]. Mental firewalls and mental encryption are theoretical controls unique to neuroprivacy. As BCIs are equipped with advanced brain-to-brain communication capabilities, there will be a need to protect against “malicious brainwaves,” which could compromise the integrity of the mind [34]. Restricting access to the mind or specific parts of the mind could be achieved with a firewall that is designed to filter brainwaves. Mental encryption would involve using a person’s cognition as part of an algorithm for encrypting their brainwaves, although it is unclear how this would work in practice [34].

The Inform strategy empowers the data subject by informing them about the processing of their personal data [11]. Inform was the third most represented Hoepman strategy with 14 controls, of which eight were implemented. Having a privacy notice that discusses neuroprivacy concerns can be thought of as the foundational control in this strategy. The four BCI providers that mentioned neuroprivacy issues in their privacy notices – Emotiv, Kokoon, Muse and NeuroSky – each disclose neural data processing details, inform data subjects of changes to the notice, and direct data subjects with privacy questions to the proper contact [15, 26, 29, 30]. Other Inform controls that were present in BCI provider privacy notices included explaining why neural data is processed [15, 26, 30] as well as stating that regulators [26, 30] and the data subjects [15, 26, 30] may be notified in the case

of a breach. The Open Brain Consent form is an open source template for a privacy notice and consent form that can be furnished to research participants whose neural data will be collected [32]. Researchers have identified a need for using accessible language in BCI privacy disclosures [6] and explaining all neural data processing in clear and plain language [37], although only dedicated research, such as analyzing the reading difficulty of BCI provider privacy notices, could discern whether this has been carried out in practice. The data subject could also receive regular reminders of the status of the neural data processing that affects them [37]. End user BCI security training [24], security demos and serious games [24] and subliminal stimuli awareness [18] have been proposed as controls for informing the data subject about brain hacking risks and defense. These measures can be considered Inform controls insofar as they provide information on how the data subject's brains could be manipulated by adversaries, as this manipulation is a form of neural data processing.

Control techniques allow data subjects to control aspects of the processing of their personal data [11]. The majority of Control techniques, six of nine, were implemented. The BCI providers had statements regarding revocation of consent and the rights of the data subject to access, port and update or correct personal data, which presumably apply to neural data as well [15, 26, 29, 30]. Kokoon stated that all neural data processing was explicitly opt-in [26], and all providers besides Emotiv implied that neural data could be retracted [26, 29, 30]. Proposed controls in this strategy include easy opt-out mechanisms for data subjects [6], obtaining informed consent for neural data processing [6] and allowing the data subject to retain control of their neural data through the use of a Personal Data Store [37]. Obtaining informed consent was marked as unimplemented due to the concerns that have been raised over whether this is theoretically possible for raw EEG data [37]. The Enforce strategy is concerned with the privacy policy of the organization [11]. Enforce had the highest percentage of implemented controls at four of five of those identified. Privacy notices and privacy policies are two different pieces of documentation – the former informs the data subject about data processing [8] and the latter is an internal piece of governance documentation [12]. Every BCI provider that discussed EEG data processing in its privacy notice was assumed to have also implemented an internal privacy policy that covers EEG data processing to some degree [15, 26, 29, 30]. Other implemented Enforce controls included ensuring employees comply with the privacy notice [15], ensuring service providers comply with the privacy notice [15, 26, 29, 30] and training all employees on requirements under legal regimes [16]. The privacy addendum was proposed as a mechanism for managing neuroprivacy risk associated with third parties such as BCI application developers [6]. The addendum requires adherence to a set of BCI Privacy Principles and would be included as a contractual requirement in various circumstances.

Demonstrate controls involve verifying and showing that personal data is being processed in compliance with privacy expectations [11]. It was the most represented Hoepman strategy in this inventory with 21 unique controls, only two of which have been implemented: Emotiv conducts internal audits for compliance with their privacy notice [15] and a blockchain platform that is capable of detecting violations of EEG data integrity was demonstrated in the lab [3]. Similar methods for ensuring neural data integrity could include those familiar from the cybersecurity field, namely checksums, digital signatures, hash functions and message authentication codes [34]. Brain antivirus would prevent,

detect and respond to brain malware [4]. An external stimuli monitor would allow the data subject to review stimuli for malicious content [4] and “detecting rapid screen changes” through the use of a subliminal stimuli detector would allow for the detection and review of malicious subliminal content [18]. Another control that could screen stimuli for malicious content is code analysis techniques that can determine the legitimacy of any stimuli presented [38]. Analyzing the physical medium for exploitation attempts could help thwart jamming attacks [4]. Various other types of detection have been proposed, such as intrusion detection systems [4], malicious noise detection [24], inconsistent classification detection [24], machine learning-based anomaly detection of malware [4] and machine learning-based inconsistency detection of processing [4]. Feedback mechanisms for undesired and uninitiated output would allow BCI users to alert BCI providers, law enforcement or other outside parties to potential neural tampering or device malfunction, which may be desirable for sensitive uses of neurotechnology such as rehabilitation where the risk of physical or psychological harm is high [24]. Auditing techniques could include periodic configuration reviews and updates [4], verifying the legitimacy of BCI software [6] and ensuring a Personal Data Store system allows for the auditing of neural data access [37]. Tracking neural data access, usage and flow throughout systems could promote accountability [27].

4. DISCUSSION

It was not apparent prior to this study that nearly one hundred neuroprivacy controls have been discussed or implemented, and that neuroprivacy controls spanning all eight Hoepman strategies have been proposed, with at least one real-world implementation in each control family except for Abstract. There were unimplemented controls in each of the eight strategies and there may be many more controls that have not been theorized yet within each strategy. Implemented controls could be improved by deploying tools that are purely academic such as Secure Multiparty Computation [1, 2] and the blockchain platform [3]. As malware and privacy attacks become more severe and commonplace in the context of BCIs, it is imperative that robust security and privacy solutions such as intrusion detection systems and brain antivirus are developed, widely deployed and made available to all data subjects and BCI providers.

The most difficult controls to actualize may be those in the Minimize and Abstract strategies, as both control families face research challenges. Minimize controls such as the BCI Anonymizer, blacklists and whitelists require being able to make a distinction between necessary and unnecessary EEG data, with no proposed methods existing in the literature for making this distinction. Even if algorithms are developed for classifying neural data in this way, BCI devices may not be able to support the control due to BCI resource constraints or BCI providers not wanting to reveal their proprietary feature extraction algorithms [39]. Abstract controls that involve adding noise to EEG data, such as differential privacy, face theoretical challenges. The accuracy of EEG data may be important to the BCI application in question, meaning these controls would directly conflict with the primary functionality of the system [28].

Perhaps one of the hard problems of neuroprivacy is that of obtaining informed consent for neural data processing. Adequately informing the data subject about neuroprivacy risks related to raw EEG data processing may be impossible due to the complexity of neural data and unforeseen advancements in data extraction [37]. If obtaining informed consent for raw neural data

processing is impossible, then perhaps informed consent can only practically be achieved by alerting the data subject to this impossibility. A related problem is deciding whether and how to inform the data subject about brain hacking risks, such as the possibilities of brain malware and subliminal probing. Providing such information to users of BCI headsets, although a potentially disturbing or frightening subject for some, may be important information for data subjects to consider when making an informed decision to use a BCI headset. Privacy notices could limit such discussion to the attacks that have been demonstrated by researchers, such as those in [18, 28], and include information about what the BCI provider is doing to prevent, detect and respond to such brain hacking threats.

Secure file deletion for neural data, secure neural data deletion, could be deployed for decreasing the likelihood of forensic recovery of neural data from storage devices. Expanding the current suite of Control techniques could involve allowing the data subject granular choices regarding the neural data that will be processed, which could involve allowing the data subject to customize which personal information contained in their EEG data they want to share or limiting data acquisition to a specific time of day, location or context. If a BCI application is a neurogame whose business model is monetizing the collected neural data of the players instead of charging a fee to participate, the application could offer a paid alternative for any privacy-conscious users who want to opt out of the monetization of their brainwaves. A novel Inform control could involve layering the privacy notice based on the type of personal data under discussion, with a dedicated section or page for neural data processing details. Ordering the notice in this manner would allow concerned data subjects to quickly learn about the neural data processing practices of the BCI provider without having to consume the whole privacy notice. Other potential Enforce controls include pursuing strong neuroprivacy as a strategic goal for the organization, treating neural data as a business asset and regularly reviewing the privacy policy for alignment with the organization's neuroprivacy strategy.

5. CONCLUDING REMARKS

The future is full of research opportunities for protecting neuroprivacy, with the majority of controls proposed in the literature without publicly known implementations. Several future work possibilities were identified for implementing and improving upon cybersecurity and privacy solutions for protecting the brain. Those in charge of neuroprivacy risk management could benefit from extensive control catalogs, including a vibrant ecosystem of competing providers and open source solutions when deciding how to allocate resources for privacy controls. Brain-computer interface users who want to mitigate their own neuroprivacy risk could also benefit from a plethora of open source solutions or paid options. Future work could adapt insights from various areas of technical privacy – including smartphone privacy controls, medical device privacy controls and the broader privacy-enhancing technologies literature – to brainstorm novel neuroprivacy controls and eventually establish a comprehensive library of neuroprivacy controls.

More research could be done to evaluate how current BCI providers handle neuroprivacy risks. For instance, in-depth studies could be performed to determine whether the reading level of current BCI privacy notices is accessible to the majority of BCI users and whether those notices accurately reflect the risks of neural data processing. Researchers could also investigate the current usability of neuroprivacy mechanisms such as opt-outs

within BCI apps to better understand if maintaining neuroprivacy is challenging from the data subject's perspective. Additionally, more clarification is needed regarding what traditional data protection rights, such as those involving data portability, access and deletion, consist of in the context of neuroprivacy. The right to have inaccuracies in one's personal data corrected is particularly challenging to conceptualize in the context of neural data [20]. Dialogues such as these could help spur innovations in neuroprivacy whilst illuminating the most pressing neuroprivacy challenges of today.

Preemptively addressing neuroprivacy risks through the development and deployment of BCI privacy controls is key to protecting the rights and freedoms of the data subjects. Perhaps most importantly, it will help avert brain hacking versions of traumatic privacy violations, such as the Snowden revelations or Facebook's Cambridge Analytica scandal, whatever a "neuro" version of such events would look like. Brain-computer interface providers, as part of their privacy risk management efforts, will need neuroprivacy control taxonomies and catalogs to operationalize privacy protections for their users. Regulators will need assurance that data protection regimes are being adhered to with regards to neural data processing, and data subjects will want assurance that their minds and brains are safe from privacy violations. This paper inventoried the existing neuroprivacy controls, serving as a starting point for researchers, BCI companies, data subjects and regulators to understand what mechanisms currently exist for mitigating brain privacy risks and indicated gaps where further innovation and research is necessary.

REFERENCES

- [1] Agarwal, A., Dowsley, R., McKinney, N. D., Wu, D., Lin, C.-T., Cock, M. D., & Nascimento, A. (2018). Privacy-preserving linear regression for brain-computer interface applications. *Proceedings of the IEEE International Conference on Big Data (Big Data)*, Seattle, WA. <https://doi.org/10.1109/BigData.2018.8621861>
- [2] Agarwal, A., Dowsley, R., McKinney, N. D., Wu, D., Lin, C.-T., Cock, M. D., & Nascimento, A. (2019). Protecting privacy of users in brain-computer interface applications. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 27(8), 1546-1555. <https://doi.org/10.1109/TNSRE.2019.2926965>
- [3] Bak, S., Pyo, Y., & Jeong, J. (2019). Protection of EEG data using blockchain platform. *Proceedings of the International Winter Conference on Brain-Computer Interface (BCI)*, Gangwon, South Korea, 1-3. Piscataway, NJ: IEEE. <https://doi.org/10.1109/IWW-BCI.2019.8737260>
- [4] Bernal, S. L., Celdrán, A. H., Pérez, G. M., Barros, M. T., Balasubramaniam, S. (2019). Cybersecurity in brain-computer interfaces: state-of-the-art, opportunities, and future challenges. *ArXiv:1908.03536*.
- [5] Bonaci, T., Herron, J., Matlack, C., & Chizeck, H. J. (2014). Securing the exocortex: A twenty-first century cybernetics challenge. *Proceedings of the 2014 IEEE Conference on Norbert Wiener in the 21st Century (21CW)*, Boston, MA. <https://doi.org/10.1109/NORBERT.2014.6893912>
- [6] Bonaci, T. (2015). *Security and Privacy of Biomedical Cyber-Physical Systems*. University of Washington, *ProQuest Dissertations and Theses*.
- [7] Brigham, K., & Kumar, B. V. K. (2010). Imagined Speech Classification with EEG Signals for Silent Communication: A Preliminary Investigation into Synthetic Telepathy. *Proceedings of the International Conference on Bioinformatics and Biomedical Engineering*, Chengdu, China, 1-4. Piscataway, NJ: IEEE. <https://doi.org/10.1109/ICBBE.2010.5515807>

- [8] Cannon, J.C. (2014). *Privacy in Technology: Standards and Practices for Engineers and Security and IT Professionals*. Portsmouth, NH: International Association of Privacy Professionals.
- [9] Chizeck, H. J., & Bonaci, T. (2014). U.S. Patent Application No. 14/174,818.
- [10] Colesky, M., Hoepman, J.-H., & Hillen, C. (2016). A critical analysis of privacy design strategies. *2016 IEEE Security and Privacy Workshops (SPW)*, San Jose, CA, 33-40. Piscataway, NJ: IEEE. <https://doi.org/10.1109/SPW.2016.23>
- [11] Cronk, J. (2018). *Strategic privacy by design*. Portsmouth, NH: International Association of Privacy Professionals.
- [12] Dennedy, M. F., Fox J., Finneran, T. (2014). *The privacy engineer's manifesto: getting from policy to code to QA to value*. New York, NY: Apress
- [13] Denning, T., Matsuoka, Y., & Kohno, T. (2009). Neurosecurity: security and privacy for neural devices. *Neurosurgical Focus*, 27(1), E7. <https://doi.org/10.3171/2009.4.FOCUS0985>
- [14] Dustman, R. E., Shearer, D. E., & Emmerson, R. Y. (1999). Life-span changes in EEG spectral amplitude, amplitude variability and mean frequency. *Clinical neurophysiology*, 110(8), 1399-1409. [https://doi.org/10.1016/s1388-2457\(99\)00102-9](https://doi.org/10.1016/s1388-2457(99)00102-9)
- [15] Emotiv. (2018, May 25). *EMOTIV privacy policy*. EMOTIV. https://id.emotivcloud.com/eoidc/privacy/privacy_policy/
- [16] Emotiv. (2019). *Mobile and Secure EEG Cloud Database*. EMOTIV. <https://www.emotiv.com/emotiv-eeg-cloud/>
- [17] Finn, R. L., Wright, D., & Friedewald, M. (2013). Seven types of privacy. *European data protection: coming of age*, 3-32. Dordrecht: Springer. https://doi.org/10.1007/978-94-007-5170-5_1
- [18] Frank, M. Hwu, T., Jain, S., Knight, R.T., Martinovic, I., Mittal, P., Perito, D., Sluganovic, I., & Song, D. (2017). Using EEG-based BCI devices to subliminally probe for private information. *Proceedings of the 2017 Workshop on Privacy in the Electronic Society (WPES'17)*, Dallas, TX, 133-136. <https://doi.org/10.1145/3139550.3139559>
- [19] Gladden, M. E. (2017). *The Handbook of Information Security for Advanced Neuroprosthetics*, (2nd ed.). Indianapolis, IN: Synthyponic Academic.
- [20] Hallinan, D., Schütz, P., Friedewald, M., & de Hert, P. (2013). Neurodata and neuroprivacy: Data protection outdated? *Surveillance & Society* 12(1), 55-72. <https://doi.org/10.24908/ss.v12i1.4500>
- [21] Hoepman, J.-H. (2014). Privacy design strategies. *Proceedings of the IFIP International Information Security Conference (SEC)*, Marrakech, Morocco. 446-459. https://doi.org/10.1007/978-3-642-55415-5_38
- [22] Hoepman, J.-H. (2019). *Privacy design strategies (the little blue book)*. Groningen: De Privacy Coach. <https://www.cs.ru.nl/J.H.Hoepman/publications/pds-booklet.pdf>
- [23] Ienca, M. (2015). Neuroprivacy, neurosecurity and brain-hacking: Emerging issues in neural engineering. *Bioethica Forum*. 8(2), 51-53. Schwabe.
- [24] Ienca, M., & Haselager, P. (2016). Hacking the brain: brain-computer interfacing technology and the ethics of neurosecurity. *Ethics and Information Technology*, 18(2), 117-129. <https://doi.org/10.1007/s10676-016-9398-9>
- [25] Inzlicht, M., McGregor, I., Hirsh, J. B., & Nash, K. (2009). Neural markers of religious conviction. *Psychological Science*, 20(3), 385-392. <https://doi.org/10.1111/j.1467-9280.2009.02305.x>
- [26] Kokoon. (2019). Privacy policy. *Kokoon*. <https://kokoon.io/policies/privacy-policy>
- [27] Li, Q., Ding, D., & Conti, M. (2015). Brain-computer interface applications: Security and privacy challenges. *Proceedings of the 2015 IEEE Conference on Communications and Network Security (CNS)*, Florence, Italy. <https://doi.org/10.1109/CNS.2015.7346884>
- [28] Martinovic, I., Davies, D., Frank, M., Perito, D., Ros, T., & Song, D. (2012). On the feasibility of side-channel attacks with brain-computer interfaces. *Proceedings of the 21st USENIX Security Symposium*, Bellevue, WA, 143-158. <https://www.usenix.org/conference/usenixsecurity12/technical-sessions/presentation/martinovic>
- [29] Muse. (2019, April 27). Privacy policy. *muse*. <https://choosemuse.com/legal/>
- [30] NeuroSky. (2018, May 25). Effective learner privacy policy. *Effective Learner Cloud*. <https://effectivelearnercloud.com/el/policies/?privacy>
- [31] The Committee on Science and Law. (2005). *Are your thoughts your own? Neuroprivacy and the legal implications of brain imaging*. New York, NY: New York City Bar Association.
- [32] Github. (2018, June 27). Open brain consent. *GitHub Blob*. <https://github.com/con/open-brain-consent/blob/master/docs/source/ultimate.rst>
- [33] P.N. Rao, R. (2013). *Brain-computer interfacing: an introduction*. Cambridge: Cambridge University Press.
- [34] Sempredoni, D., & Viganò, L. (2018). Privacy, security and trust in the internet of neurons. *ArXiv:1807.06077*.
- [35] Solove, D. J. (2006). A taxonomy of privacy. *University of Pennsylvania Law Review*, 154, 477-560.
- [36] Spiekermann, S., & Cranor, L. F. (2009). Engineering privacy. *IEEE Transactions on Software Engineering*, 35(1), 67-82. <https://doi.org/10.1109/TSE.2008.88>
- [37] Stopczynski, A., Greenwood, D., Hansen, L. K., & Pentland, A. (2014, April 21). Privacy for personal neuroinformatics. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssm.2427564>
- [38] Takabi, H. (2016). Firewall for brain: towards a privacy preserving ecosystem for BCI applications. *Proceedings of the 2016 IEEE Conference on Communications and Network Security (CNS)*, Philadelphia, PA. <https://doi.org/10.1109/CNS.2016.7860516>
- [39] Takabi, H., Bhalotiya, A., & Alohaly, M. (2016). Brain computer interface (BCI) applications: privacy threats and countermeasures. *Proceedings of the International Conference on Collaboration and Internet Computing (CIC)*, Pittsburgh, PA, 102-111. Piscataway, NJ: IEEE. <https://doi.org/10.1109/CIC.2016.026>
- [40] Wu, D., Lawhern, V. J., Gordon, S., Lance, B. J., & Lin, C. (2017). Driver drowsiness estimation from EEG signals using Online weighted Adaptation Regularization for Regression (OwARR). *IEEE Transactions on Fuzzy Systems*, 25(6), 1522-1535. <https://doi.org/10.1109/TFUZZ.2016.2633379>
- [41] Zhang, S., Yuan, S., Huang, L., Zheng, X., Wu, Z., Xu, K., & Pan, G. (2019). Human mind control of rat cyborg's continuous locomotion with wireless brain-to-brain interface. *Scientific reports*, 9(1), 1-12. <https://doi.org/10.1038/s41598-018-36885-0>