Privacy in the Age of Neurotechnology: Investigating Public Attitudes towards Brain Data Collection and Use

Emiram Kablo
ekablo@mail.upb.de
Paderborn University
Paderborn, Germany

Patricia Arias-Cabarcos
pac@mail.upb.de
Paderborn University
Paderborn, Germany

ABSTRACT

Brain Computer Interfaces (BCIs) are expanding beyond the medical realm into entertainment, wellness, and marketing. However, as consumer neurotechnology becomes more popular, privacy concerns arise due to the sensitive nature of brainwave data and its potential commodification. Attacks on privacy have been demonstrated and AI advancements in brain-to-speech and brain-to-image decoding pose a new unique set of risks. In this space, we contribute with the first user study (n=287) to understand people's neuroprivacy expectations and awareness of neurotechnology implications. Our analysis shows that, while users are interested in the technology, privacy is a critical issue for acceptability. The results underscore the importance of consent and the need for implementing effective transparency about neurodata sharing. Our insights provide a ground to analyse the gap in current privacy protection mechanisms, adding to the debate on how to design privacy-respecting neurotechnology.

CCS CONCEPTS

• **Security and privacy** → *Usability in security and privacy*; *Social aspects of security and privacy*; **Human and societal aspects of security and privacy**.

KEYWORDS

neurotechnology, neuroprivacy, brain data, contextual integrity, user study

ACM Reference Format:
Emiram Kablo and Patricia Arias-Cabarcos. 2023. Privacy in the Age of Neurotechnology: Investigating Public Attitudes towards Brain Data Collection and Use. In *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security (CCS '23), November 26–30, 2023, Copenhagen, Denmark*. ACM, New York, NY, USA, 14 pages. https://doi.org/10.1145/3576915.3623164

1 INTRODUCTION

The usage of Brain Computer Interfaces (BCIs) has evolved far beyond the medical field to enter domains like entertainment, wellness, and marketing. To give some examples, we can nowadays find brain-controlled games or meditation applications to be enjoyed with consumer-grade electroencephalogram (EEG) readers. VR vendors are introducing brainwave sensors in their headsets to improve user experience [1], and integrations in everyday wearables, such as common earbuds are expected. In fact, the BCI market is growing at a fast pace with projections predicting an expected value of $3.93 billions in 2027 [17].

However, while neurotechnology gadgets open new horizons for providing rich user services to enhance the lifestyle of the general public, sensing the brains of consumers opens the door to unprecedented privacy abuses [10]. Brainwaves correlate with our mental states, cognitive abilities, and medical conditions, which can be a basis for inferring emotions, prejudices, interests, health disorders, personality traits, or other private data that can be used perniciously [22]. Indeed, the feasibility of some of these attacks has been already demonstrated [30**?**], and new attack possibilities appear with the swift advancement of AI and its application to brain data processing. Just recently, scientists from the University of Texas successfully applied Large Language Models to decode human thoughts with 82% accuracy using non-invasive EEG technology [46]. Similar improvements have been published regarding the extraction of imagined images [45].

It becomes apparent that looking into mind-related data poses risks to individual freedom of thought, and therefore to society, which calls for new legal frameworks to regulate "Neuroprivacy" [23]. In this space, the ethical and legal debates have been intense, leading to Chile's pioneering modification of its constitution to introduce new neurorights [21]. While research so far has concentrated on addressing ethical, legal, and BCI cybersecurity issues, such as vulnerabilities and attacks, there is little knowledge about users' privacy needs. In the face of the unique risks and possibilities brought about by consumer neurotechnology, it is crucial to understand users' privacy expectations regarding BCI devices to build ground knowledge for designing privacy-respecting applications. To address this current knowledge gap, we design and conduct a user study (n=287) driven by the following research questions:

- **RQ1 [Neuroprivacy Expectations] Under which conditions do people consider sharing neurodata acceptable?** We seek to understand with whom, for which services, and to which extent users are willing to share their brain data.
- **RQ2 [Neuroprivacy and Neurotechnology Awareness] How aware are people of neurotechnology privacy implications? How would they use this technology?** We seek to explore the level of comprehension, specially regarding privacy, that users have about neurotechnology, identifying barriers and enablers for adopting BCIs.

CCS '23, November 26–30, 2023, Copenhagen, Denmark

ACM ISBN 979-8-4007-0050-7/23/11…$15.00
https://doi.org/10.1145/3576915.3623164

We contribute with the first study on user-centered neuroprivacy, using Contextual Integrity (CI) theory [37] as the ground framework to collect privacy norms around neurodata. Our analysis shows that people have strong concerns against the use of brain data for advertising, while they are interested in health and research use cases, as well as using neurotechnology for self-knowledge and life improvement. Generally, while our participants are positive about the technology, privacy is seen as a critical issue and informed consent is placed as a core determinant for acceptability. These results underscore the need to design better transparency, neuroprivacy protections, and enforcement mechanisms before the commodification of neurodata sharing normalizes brain surveillance.

Methodologically, by analyzing how users rate the acceptability of neurodata-sharing scenarios, we uncover limitations in the current application of the CI survey methodology, which we unpack in our discussion to inform the design of future studies.

2 BACKGROUND

This section sets the terminology and reviews related work, highlighting how we expand on it and contribute to this research.

2.1 Neuroprivacy and Contextual Integrity

Neuroprivacy or "brain privacy" refers to the rights people have regarding the extraction and analysis of neural data (also known as neurodata) from their brains [7]. It encompasses the privacy issues raised by the use of neurotechnology and its relevance increases with the advancement of consumer BCIs and applications outside the more controlled medical scenarios. Current BCI technologies can be classified as invasive and non-invasive methods. The first type, invasive methods, record signals within the cortex by directly implanting electrodes near the surface of the brain, such as NeuraLink's implantable neural threads[1]. In turn, non-invasive neurotechnology operates by measuring the brain from the outside, which is less risky, and therefore more convenient for non-critical applications. The most portable and commonly used of these non-invasive techniques is electroencephalography (EEG), which records electrical activity through sensors placed on the scalp surface. We focus our privacy study on this type of consumer neurotechnology, for which multiple wearables (e.g., Muse[2], Emotiv[3]) and services are currently available in the market, signaling the potential for more immediate increased adoption.

When it comes to studying privacy, **Contextual Integrity (CI)** is a relevant instrument developed by Helen Nissenbaum [37]. This theory provides a framework for understanding privacy expectations and norms, on the argument that they are shaped by the specific context in which personal information is shared, rather than being governed by universal principles or laws. CI specifies 5 parameters that define an information flow: (1) the information *sender*, (2) the information *recipient*, (3) the *attribute* (i.e., type of information) being sent, (4) the *subject* of the information, and the (5) *transmission principle* or condition that governs the information flow. Any change of parameters can lead to a privacy norm violation, and result in feelings of discomfort or mistrust. For example, a person (sender/subject) may be comfortable sharing personal brain-related (attribute) information with their doctor (recipient) in the context of a medical exam with guaranteed confidentiality (transmission principle). However, they would feel violated if that same information were shared with a third party without their consent (change of subject/transmission principle).

CI has been widely used in different research communities as a framework to define and reason about privacy, as a tool to evaluate conformance between expectations and privacy regulations, or to find gaps in state-of-the-art technical protections to match user needs [3, 4, 19, 38, 41].

The study of neuroprivacy expectations, as one of the focuses of the present paper, is an application of CI to understand user norms in emerging scenarios involving brain data sharing.

2.2 Related Work

2.2.1 On BCI Privacy and Security. Brain activity is rich in information. Indeed, given the unique features of these signals, they can be used to identify individuals, as a type of biometric [20]. But beyond the identification capabilities underlying neurodata, brainwaves correlate, among others, with our mental states, cognitive abilities, and medical conditions [44]. Examples of concrete inferences from brain data are alcoholism detection [25] or emotion tracking [27]. The type of privacy leaks that can come from neurodata have been recently systematized [29], highlighting its potential consequences, especially regarding the risk of user tracking for targeted advertisement.

Several works cover attack vectors. Martinovic et al. [30] and Frank et al. [18] demonstrated how by manipulating visual stimuli (even subliminal) presented to BCI users, their EEG signals could reveal private information such as PINs or the area where they live. Tarkhani et al. [47] analyzed wearable BCIs to detect vulnerabilities from an operating system and adversarial machine learning perspective, proposing mitigation solutions. A survey by Hanisch et al. [22] on behavioral data privacy highlights that brain data protection research is significantly immature with regard to that on other types of behavioral information. Another comprehensive review, by Bernal et al. [9] categorizes the state-of-the-art on attacks to BCIs, calling for the need to raise user awareness. Our work contributes to this area, by expanding knowledge of users' privacy needs as a core to better understand attack vectors and design protection mechanisms.

2.2.2 On Users' (Neuro)-Privacy Expectations. A rich stream of research explored user privacy attitudes, perceptions, and reactions with regard to new forms of data collection and emerging technologies, such as social networks [41], Virtual Reality [2], IoT or smart homes [3, 4]. In the field of neuroprivacy and neurotechnologies, however, there is little work on understanding user concerns. Chuang et al. [31], as a side study to gather the acceptability of brain biometrics, conducted a survey with 200 participants asking about privacy perceptions on different biosignals. They concluded that brainwaves are seen among the most revealing biosignals in their ability to reveal the inner workings of a person's mind. However, they did not capture the contextual factors shaping users'

[1] https://neuralink.com/approach/
[2] https://choosemuse.com/
[3] https://www.emotiv.com/

neuroprivacy norms. We comprehensively extend this knowledge, by investigating under which conditions users are willing to share neurodata.

The most relevant work in the application of CI to capture privacy expectations is the research by Apthorpe et. al [3, 4] on IoT privacy norms. In their first study, they introduce a survey method to apply CI for quick and efficient discovery of privacy norms at scale and apply it (n=1,731) in the smart home context. In their second study, they apply this same survey method to analyze how the COPPA privacy regulations meet the privacy norms of parents whose children use smart IoT toys, proving alignment. We replicate their survey methodology, adapting it to the neuroprivacy context, to collect user expectations regarding 116 possible neurodata sharing scenarios. During the study design, we discovered and solved methodological limitations that became lessons learned to refine the application of this instrument in the future.

3 METHODOLOGY

Here we describe our survey-based methodology, based on Apthorpe et al.'s instrumentalization of CI to scale the collection of privacy norms [3], and adapting it to the neuroprivacy case.

3.1 Defining Neurodata Information Flows

In this section, we describe how we selected the CI parameters to define neurodata information flows, summarized in Table 1. We started by fixing the *sender* to be a BCI device. BCI devices are the most common type of gadget available in the market, while integrations of brain data acquisition sensors in other objects (e.g., VR headsets) are still under development. We also established the *attribute* parameter as "brain signals" because this is the basic raw type of information sent by BCI devices. There are other possible data types that a BCI can collect, e.g., devices usually include other sensors, such as gyroscopes to detect head movement or Photoplethysmography (PPG) sensors to record heart rate. However, given that brain data privacy is a complex enough type of data and is currently understudied, we decided to focus on neurodata privacy alone, which can set a baseline to later analyze the effect of incrementally adding other attributes. Furthermore, as BCI devices are wearable and always measure their user, the *subject* parameter is also fixed. The remaining two CI parameters were selected as described below.

Recipients. We crafted our list of neurodata recipients based on practical BCI application scenarios that have been already deployed, consulting academic literature and recent books that comprehensively review the state of neurotechnology [15, 36]. We also reviewed the most popular consumer devices available in the market [13, 24, 34, 35] and the services/apps offered on top of them as described in their websites and associated marketplaces.

Based on our analysis, the most frequent usage of neurotechnology so far comes from the medical realm, where it has been crucial for disease diagnosis [42]. Additionally, research with BCIs has been continuous and varied, from the development of health-related applications to enable disabled people to communicate using brain-to-speech technology [40], to the study of cognitive processes, such as learning, in psychology research, or the exploration of brainwaves as a new form of biometric authentication in the field of computer security [5, 12, 20]. On these grounds, we added doctors and researchers as options in the CI recipient parameter.

The review of state-of-the-art reports on neurotechnology yielded additional recipients. Employers are already using (or envisioning) neurodevices, for example, to detect drowsiness in truck drivers or to track focus. BCIs have been used in schools to allow teachers to understand if students are focused. Early adopters post their EEG data screenshots to social media accounts to exchange experiences and get comments from other members, so we also deemed it plausible that data could be shared with family members. Finally, there is a growing number of government-led research initiatives, which makes government agencies another relevant recipient. An example of this latter scenario is the increased use of neuroscientific evidence in criminal proceedings.

When reviewing the current neurotechnology market, the manufacturer of the BCI devices came up as an obvious recipient of brain data. Other types of recipients are the online services that provide applications based on brain data and that are not necessarily linked to the manufacturer. In the app marketplaces for different BCI devices, gaming was the bigger category, so we additionally included entertainment companies as specific recipients.

Transmission Principles. We defined the transmission principles based on common conditions under which general data information flows occur in similar scenarios. More specifically we adapt the list of transmission principles in the IoT CI-based privacy study by Apthorpe et al. [3], considering the neurodata use cases described earlier. Some of these principles are specifically mentioned in many device privacy policies, such as "data is kept confidential and secure". Others involve common practices in data collection and storage, for example "if the user has given verifiable and revocable consent". We also added compliance to the GDPR, to gauge the importance of this regulation for European users. Furthermore, as seen in [3], we added a *null* transmission principle to create unconditional information flows as well.

Finally, when combining transmission principles with recipients, we discarded certain information flows that were not applicable at the authors' discretion, based on the use cases investigated in the related work.

This process resulted in a total of 116 neurodata information flow descriptions to be rated regarding acceptability, including the *null* principle. This is the core component of our survey questionnaire described in the next section.

3.2 Survey Design & Implementation

Structure. The survey procedure starts by asking users to provide informed consent. Then, we give a brief definition of what Brain Computer Interfaces are, explaining that EEG wearables are already available in the market and showing examples. After this introduction, participants fill in the questionnaire part. All materials about the survey instrument can be found in Appendix A.

Our survey questions are organized in the following categories:

- **Neuroprivacy Expectations**. We start by asking about the acceptability of sending neurodata to the 10 different receivers specified in Table 1. These answers establish the acceptability baseline, i.e., the judgment of users regarding data sharing without considering a specific transmission principle (*null* principle).

Table 1: Contextual Integrity Parameters to define Neurodata Flows

Sender	Recipients	Subjects & Attributes	Transmission Principles
A BCI device	its manufacturer	the user's brain signals	if the user has given verifiable and revocable consent
	online service providers		if the user is directly notified before data collection
	academic researchers		if data is kept confidential and secure
	entertainment companies		if data is stored online for a limited period
	the user's medical doctor		if data is only stored until required
	the user's social media accounts		if data is used to improve device performance
	government agencies		if data is used to minimize safety concerns
	immediate family members		if data is used to enhance cognitive abilities
	the user's employer		if data is collected for medical assessment and monitoring
	the user's teacher/professor		if data is used for marketing
			if data is used for academic research
			if data is used to improve entertainment experience
			if data is subjected to privacy-preserving techniques
			if it complies with the EU Data Protection Regulation
			if the information is used to provide a personalized service
			null (no principle)

Subsequently, participants rate the acceptability of sending neurodata to each receiver under the specific applicable principles defined in Table 1. The order of the information flows in each of these blocks was randomized per respondent. In total, our respondents rated a total of 106 conditional flows and 10 unconditional flows. Rates are given per flow on a 5-point Likert scale with choices: Completely Acceptable (2), Somewhat Acceptable (1), Neutral (0), Somewhat Unacceptable (-1), Completely Unacceptable (-2). We have added the matrix-like visual presentation to rate neurodata flows in the questionnaire to the Appendix. To further explore how users rate acceptability, we included two additional questions. First, we asked if the transmission principles were evaluated independently or if consent was always implicitly assumed. Then, we added an open question for participants to share any important information on their decision-making process. This part of the survey addresses research question RQ1.

- **Neuroprivacy & Neurotechnology Awareness.** Regarding neuroprivacy awareness, we want to capture if people have any knowledge of what types of sensitive data can be revealed from brainwaves. For this purpose, we asked participants what information they think can be inferred from brain data captured with a commercial BCI. Regarding neurotechnology awareness, in case they are BCI users, we ask about their usage scenarios and frequency, as well as if they would use this technology in different situations, providing a rationale. For non-users, we ask whether they would use a BCI headset if they had the opportunity, inquiring also about their intended usage scenario(s) and reasons. This set of qualitative questions was designed to help in understanding the *why* behind acceptability scores, contributing to answering research question RQ2.
- **Demographics & Background**. The questionnaire ends with questions to collect demographic data (as listed in the Appendix: gender, age, education, country) and information about participant's backgrounds that we deem relevant to analyze their neuroprivacy responses in the main part of the questionnaire (IT knowledge, BCI knowledge, technology usage, privacy attitudes). We use the standard Internet Users' Information Privacy Concerns (IUIPC) scale by Malhotra et al. [28] to measure privacy attitudes. The information collected in this part of the survey allows us to understand the representativeness of our sample and to account for demographic variables in our analysis.

Design Considerations. We followed best practices in survey design to guarantee the quality of the instrument. To avoid priming our participants, the study title and introduction did not mention or even allude to privacy. Instead, respondents participated in a "*Brain Computer Interface Technology Survey*". To account for potential dishonest or careless answers, we included attention-check questions.

Deployment and Ethical Assessment. The study was approved by our university's Institutional Review Board. We recruited participants using Prolific [39], an online platform that explicitly caters to researchers. We applied Prolific screening settings to select participants living in European countries. We decided to focus on people from Europe to get the perspective of technology users who are subject to the EU General Data Protection Regulation (GDPR) [14]. Participants were informed that the survey was anonymous, voluntary, and that all collected data would be processed according to the GDPR, before asking for consent and confirmation of being over 18 years old. The questionnaire was administered via the LimeSurvey web-based survey tool[4], whose servers are located in Germany and comply with the GDPR.

At the end of the survey, participants were debriefed on the actual purpose of exploring privacy concerns around BCIs and provided the option to contact us in case they had questions.

3.3 Data Analysis

Quantitative analysis methods. Closed-ended responses on acceptability were mapped from the Likert scale to quantitative values, averaged, and treated as continuous data. We analyzed for significant differences in acceptability scores contingent upon a) specific

[4] https://www.limesurvey.org/

transmission principles, and b) impact on sharing data with different recipients. In all cases, we performed non-parametric Wilcoxon signed-rank tests [49] for repeated measurements, accounting for the Bonferroni multiple-testing correction [6]. Our objective for both analyses was to assess the effect of recipients and transmission principles. Here, we followed the methodology of Apthorpe et al. [3]. We compared the acceptability score of each recipient against the *null* principle, which represents unconditional data flows, and *null* transmission principles that portray conditional information flows. The *null* principle served as a baseline for assessing the influence of the recipients and transmission principles on the acceptability score across all participants. We performed multiple tests for each of the 10 recipients and the 15 transmission principles. Due to the inapplicability of some principles to certain recipients, a total of 106 tests were conducted. To account for multiple comparisons, we adjusted the standard threshold of 0.05 to a new threshold of 0.00047 (0.005 / 106) based on the Bonferroni method.

Furthermore, we report the results of six Wilcoxon tests to compare different demographic groups: age, education, IT background, wearable device usage, BCI knowledge, and consent interpretation. In this case, we set the significance threshold to 0.0083 (0.05 / 6) to account for the Bonferroni correction method since we used six demographic groupings.

Qualitative analysis methods. Open-ended responses were analyzed following an iterative, inductive coding approach [32]. One member of the research team read responses and created the codebook with thematic codes, and a second researcher independently coded the full set of data. We measured the inter-coder reliability using Cohen's Kappa [11]. The high scores obtained on this metric for all questions (κ > 0.75) suggest an excellent level of agreement [16]. The cases where the coders differed in their final codes were discussed and reconciled. The final codebooks with category descriptions and code distribution are detailed in Appendix A. Additionally, we tested for significant differences in the open-ended answers between participants giving high (>=0) and low average acceptability scores (<0), using a Chi-square test (χ^2). This allows us to look at the interplay between the quantitative and qualitative data.

Pilot test. Before publishing the main survey, we conducted a pilot test with 15 participants, requesting their feedback and explicitly querying for perceived ambiguities. The most important observation during piloting was that several participants mentioned they did not know how to rate the acceptability of the principles. More specifically, some people assumed that for transmission principles to be "valid" they should all contain consent, e.g., "*if data is kept confidential and secure*" should be *"if the user has given consent* ***and*** *data is kept confidential and secure"*. This discussion led us to conclude that participants might rate transmission principles in two different ways: either independently, assuming that consent is not given; or assuming that consent is implicit. This different understanding can potentially distort the acceptability scores and the correct interpretation of the results. On these grounds, we decided to include new questions to capture insights into the individuals' rating procedure.

Additionally, the pilot was useful to sharpen the wording and arrangement of questions through minor modifications. It confirmed that the data gathered was precise and consistent with the anticipated structure and it helped establish an estimated completion time (roughly 20 minutes).

Piloting the study also helped us to estimate if the number of questions in the survey was too tiring. The design decision of asking users to provide ratings for 116 flows (including the *null* principle) was based on previous work. Apthorpe et al. successfully carried out their CI study with participants rating 82 flows, therefore we followed the same matrix-like question design that reduces cognitive fatigue (see Fig. 1). Since we added two more matrices, we explicitly asked pilot participants to give us feedback, and their answers signaled that it was adequate.

4 RESULTS

The study was conducted in March 2023. We received 435 responses and the final sample is n=347. We had to filter duplicates, answers that failed the attention checks, and participants that did not conform to the screening filter on the EU location. In the first baseline acceptability question, we added the response option, "Doesn't Make Sense" and 60 participants responded that it does not make sense to share brain signals with one or more recipients, as seen in Figure 2. Despite their participation in the study, these individuals were not included in the subsequent quantitative acceptability analysis as it was apparent that they did not fully comprehend some information flows. This resulted in a total number of 287 participants. Participants took 20 minutes on average to complete the survey, and we paid them 3£.[5].

Participants Background. Our study participants who rated all flows (n=287) are 62% men, 36% women, and the rest either chose not to disclose their gender or identify as non-binary. Regarding age and education, participants are mostly young adults (80% <35 years old), and the majority has completed a bachelor's (35%) or master's degree (24%). It is also relevant to note that 42.5% of the respondents have a background in IT. When it comes to knowledge about BCIs, while the majority (55%) has heard about this technology, only 7 participants (2%) own or have used a neurodevice. Appendix A provides detailed demographic information.

4.1 RQ1 - Neuroprivacy Expectations

Our investigation involves examining the acceptability of sharing brain signals with different recipients, as well as analyzing how this practice relates to various factors like the participants' familiarity with BCI technology, age group, and level of education.

Figure 4 shows the overall acceptability score of participants sharing their brain signals with different recipients. The color scale ranges from red, indicating -2 (Completely Unacceptable), to green, representing 2 (Completely Acceptable). Our findings suggest that participants are more amenable to sharing brain signals with medical doctors and academic researchers compared to sharing their data with government agencies, social media platforms, or online service providers. Participants appear to be ambivalent about whether to disclose brain data to their family members, professors/teachers, or manufacturers of BCI technologies. In the following subsections, we will present the most meaningful insights from our results. The

[5]This compensation was settled based on the ethical reward scheme recommended by Prolific, i.e., at least 9£/hour

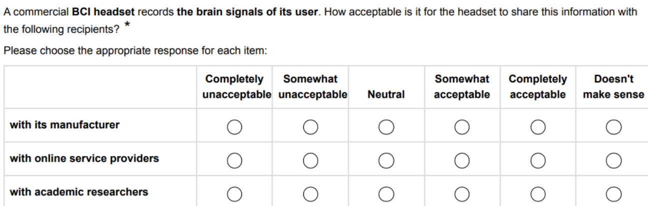

A commercial **BCI headset** records **the brain signals of its user**. How acceptable is it for the headset to share this information with the following recipients? *

Please choose the appropriate response for each item:

	Completely unacceptable	Somewhat unacceptable	Neutral	Somewhat acceptable	Completely acceptable	Doesn't make sense
with its manufacturer	○	○	○	○	○	○
with online service providers	○	○	○	○	○	○
with academic researchers	○	○	○	○	○	○

(a) Unconditional Brain Data Flows to Recipients

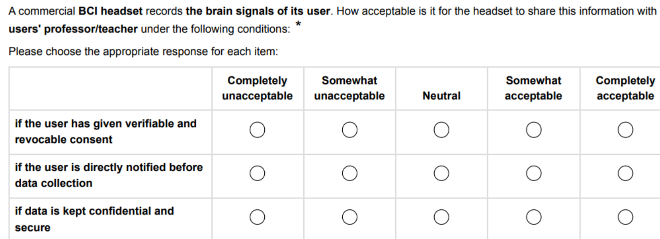

A commercial **BCI headset** records **the brain signals of its user**. How acceptable is it for the headset to share this information with **users' professor/teacher** under the following conditions: *

Please choose the appropriate response for each item:

	Completely unacceptable	Somewhat unacceptable	Neutral	Somewhat acceptable	Completely acceptable
if the user has given verifiable and revocable consent	○	○	○	○	○
if the user is directly notified before data collection	○	○	○	○	○
if data is kept confidential and secure	○	○	○	○	○

(b) Brain Data Flows for Specific Recipient under Transmission Conditions

Figure 1: Matrix-like blocks of questions used for rating acceptability of brain data flows. (a) captures acceptability of recipients under no specified transmission condition (*null* principle). (b) captures acceptability ratings for a specific recipient (teacher/professor) under different conditions. Rows are partially shown to illustrate the question visualization concept; the full set of principles and recipients is described in Table 1.

Figure 2: Number of "Doesn't Make Sense"-responses for different recipients

appendix provides further relevant graphs and a detailed overview of the average acceptability scores over all CI answers for different demographic groups and the resulting p-values to indicate statistically significant differences.

4.1.1 Principles Show an Impact on Participants' Perception of Acceptability. As explained in Section 3.3 we performed significance tests to compare the effect of conditional and unconditional information flows and the effect of having different recipients.

The graph in Figure 5 shows the percentage of instances for which including a particular transmission principle (vs. the *null* transmission principle) leads to a statistically significant difference in average acceptability score. We can observe that the principles "if data is used for marketing" and "if data is used to improve device performance"[6] result in significantly different scores with regard to unconditional flows in 100% of the instances. The principle "if the user has given verifiable and revocable consent" resulted in significantly different scores in 90% of the cases. Participants showed a greater aversion towards sharing their brain data when the data collection was intended for marketing purposes while giving consent resulted in a more positive attitude about sharing brain signals. Apthorpe et al.'s study showed similar results in regards to the effect of transmission principles and recipient on information flow acceptability of sharing data from different IoT devices [3]. From this, we can infer that principles have a major impact on how participants perceive the acceptability of sharing brain signals. On the other hand, we noted that the principle "if data is stored online for a limited period" results in significantly different scores in 50% of the instances.

4.1.2 Recipients Ranking Changed in Unconditional vs. Conditional Information Flows. Comparing the ranking of recipients for unconditional (i.e., *null* principle) and the average acceptability scores for conditional information flows, sorted by their means, we notice a change in the least acceptable recipients. Before asking the participants about the acceptability of data flows for different recipients, we observed that government agencies were the least acceptable recipients, followed by social media accounts and online service providers. After naming transmission principles explicitly, participants find it least acceptable to share brain signals with their social media accounts, employer, and entertainment company, respectively, only then followed by governments. After presenting conditional flows determined by transmission principles, participants gained a more realistic comprehension of purposes and cases for brain data collection and processing.

4.1.3 Doctors and Researchers are Acceptable Brain Data Recipients. Figure 3 displays a heat map showcasing the acceptability of different transmission principles for various recipients of brain signals from BCIs. The y-axis shows the transmission principles, while the x-axis displays the recipients. The last row consists of the values for the *null*-principle which indicate the overall acceptability score as seen in Fig. 4 for the unconditional information flow.

When it comes to analyzing the willingness to share BCI data across all participants, considering transmission principles, the results demonstrate that medical professionals receive the highest positive value which indicates a strong trust towards sharing data for medical purposes as well as with academic researchers who are also favored among all participants. Sharing data with immediate family members and professors or teachers is considered acceptable but rated lower. It is striking that users do not find it acceptable to

[6] Note that this principle is only applicable to the recipient "device manufacturer"

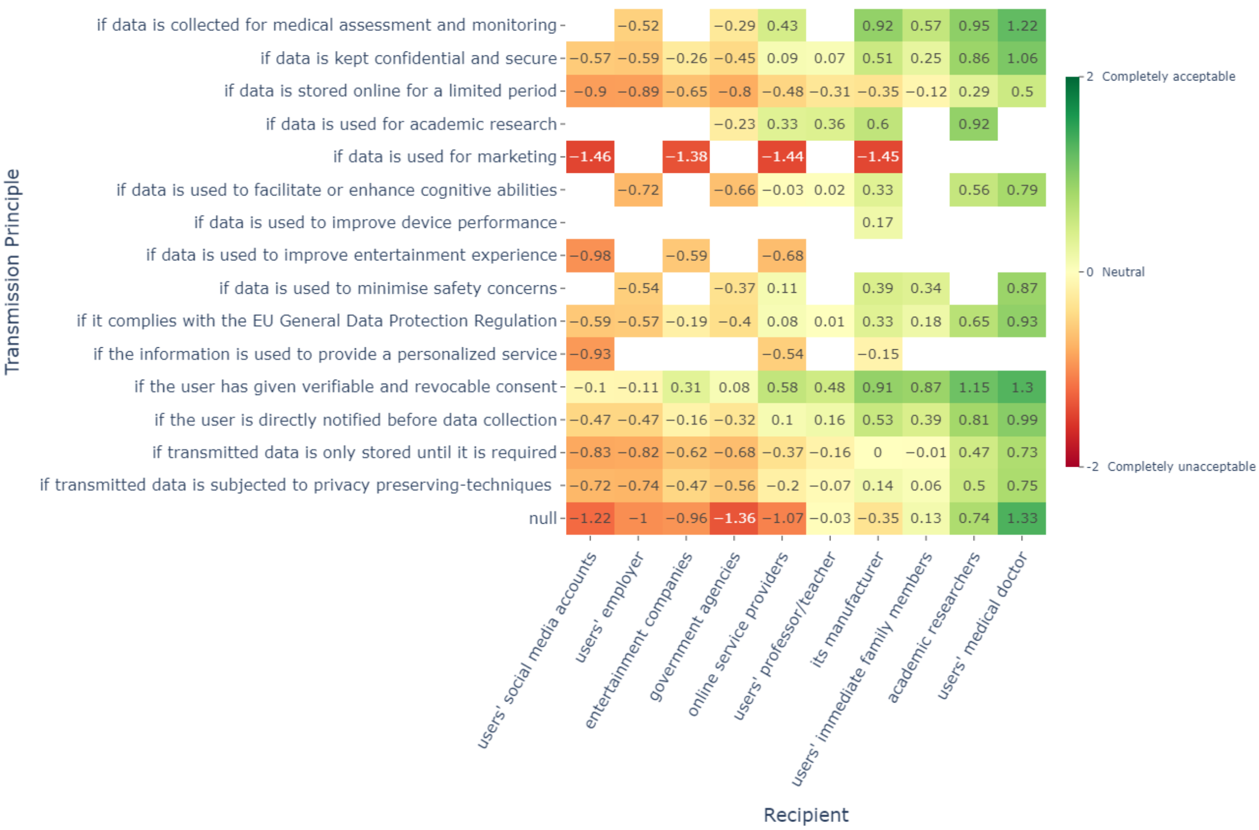

Figure 3: Heat map showing average acceptability scores of information flows by transmission principles and recipients

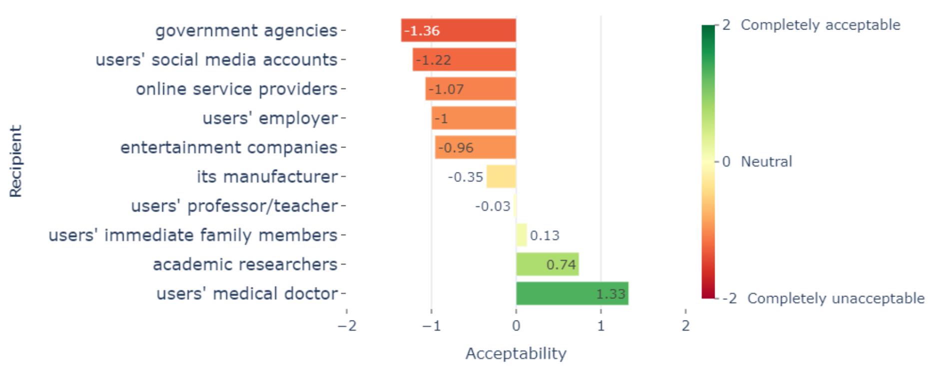

Figure 4: Overall acceptability score of participants sharing brain signals with different recipients without any condition

share their data with their employers under any condition, including safety.

4.1.4 Using Brain Data for Marketing is Unsuitable. Users are more inclined to share their data if it is not used for marketing purposes, regardless of the recipient, such as their social media accounts, online service providers, entertainment companies, or device manufacturers. Among all age groups, particularly those aged 25-54 years old, displayed a strong aversion to sharing brain signals if the data is used for marketing purposes. This negative attitude may stem from

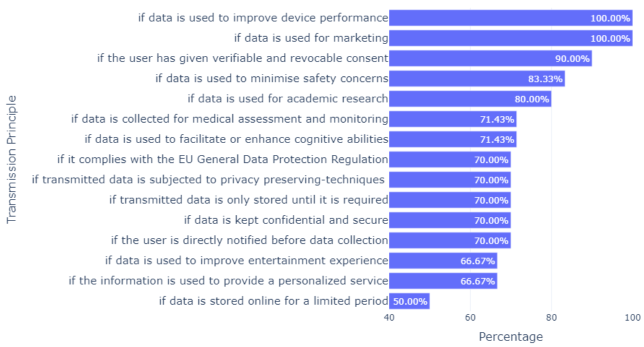

Figure 5: Percentage of instances where the inclusion of the specified transmission principle resulted in a statistically significant difference in acceptability scores with regard to the baseline *null* principle

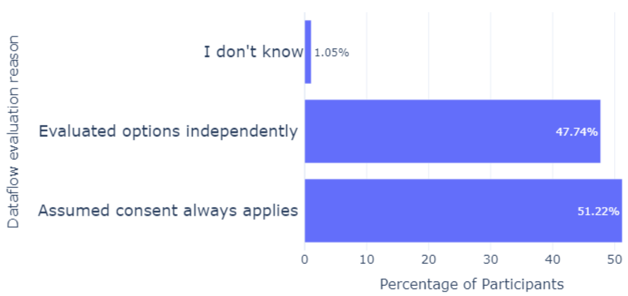

Figure 6: Distribution of participants' reasons for evaluating transmission principles as a percentage

concerns about privacy, data misuse, and the commercialization of sensitive information.

Regarding data collection for medical assessment and monitoring, the respondents find data monitoring for health reasons acceptable, but they are not comfortable with any recipient carrying out the monitoring. They believe that it is unsuitable for governments and employers to monitor health data.

4.1.5 Consent Makes a Difference. When conducting the survey, we inquired about the interpretation of the transmission principle, i.e., if participants assumed that verifiable and revocable consent always applies under conditional data flows. As an example, participants would infer that "the user is notified" means "the user has given consent and is notified". An alternative interpretation of the data flows is that the participants evaluated the principles individually, i.e., "the user is notified" is interpreted as stated.

Figure 6 is a bar chart illustrating participants' approaches to evaluating transmission principles in BCIs. The findings indicate that approximately 51% of participants assumed that consent was given when answering questions related to the flow of brain data, indicating a potential lack of clarity around data privacy and consent. Around 48% of the participants evaluated options independently and did not assume that consent applies to every principle. Looking at the heat maps in Figure 7, the respondents who assumed that consent always applied are more willing to share their data. Their average acceptability score over all CI flows lays at 0.13, while the score for the group of participants evaluating principles independently is -0.25. This difference is significant ($p < 0.0083$). Hence, conditional information flows are more acceptable for individuals who consider that consent always applies.

4.1.6 Young Adults are Less Concerned. In our study, we analyzed the average acceptability scores of information flows for different recipients in relation to the age groups of our participants. Our findings indicate that there are notable differences in attitudes toward sharing brain signals among different age groups. Specifically, we observed that participants younger than 45 are slightly more open to sharing their brain signals in comparison to participants between the ages of 45-64. The average acceptability scores for all rated flows of the first group is -0.04 whereas the second group gave a score of -0.19. The difference in the acceptability scores is significant with $p < 0.0083$.

To elaborate further, individuals aged 18-44 years exhibited a relatively favorable perspective on sharing brain signals if the information is obtained for medical assessment and treatment, kept confidential and secure, or used for academic research. This trend suggests that younger and middle-aged adults recognize the potential benefits of BCI technology in healthcare and research. Additionally, they seem to be more favorable towards sharing brain signals if the user has given verifiable and revocable consent, highlighting the importance of user control and consent in the acceptability of BCI technology.

Interestingly, people with a higher age tend to express more apprehension about sharing brain data overall. These findings underline the importance of considering age as a crucial factor influencing trust, privacy concerns, and sharing preferences.

4.1.7 Academic Degrees show More Openness to Researchers as Recipients. The analysis provides interesting insights into how individuals from different educational backgrounds perceive sharing of brain signal data. Respondents who held advanced academic degrees, such as a Doctorate or Master's degree, demonstrated a more favorable attitude towards sharing their brain signals with academic researchers than participants with a lower level. The participants' negative attitude towards sharing their brain signals with entertainment companies may stem from their perception that such entities are less concerned with ensuring personal well-being and safety, as compared to medical or academic settings.

Comparing the average score over all neurodata flows of participants with a higher academic degree (Bachelor's, Master's, or Doctorate) and the average score of those with other educational backgrounds, we observe that the first group is more concerned (-0.12 versus 0.07). The difference is statistically significant ($p < 0.0083$).

4.1.8 Concerns Among Non-IT Individuals. We further analyzed the average acceptability score of transmission principles with different

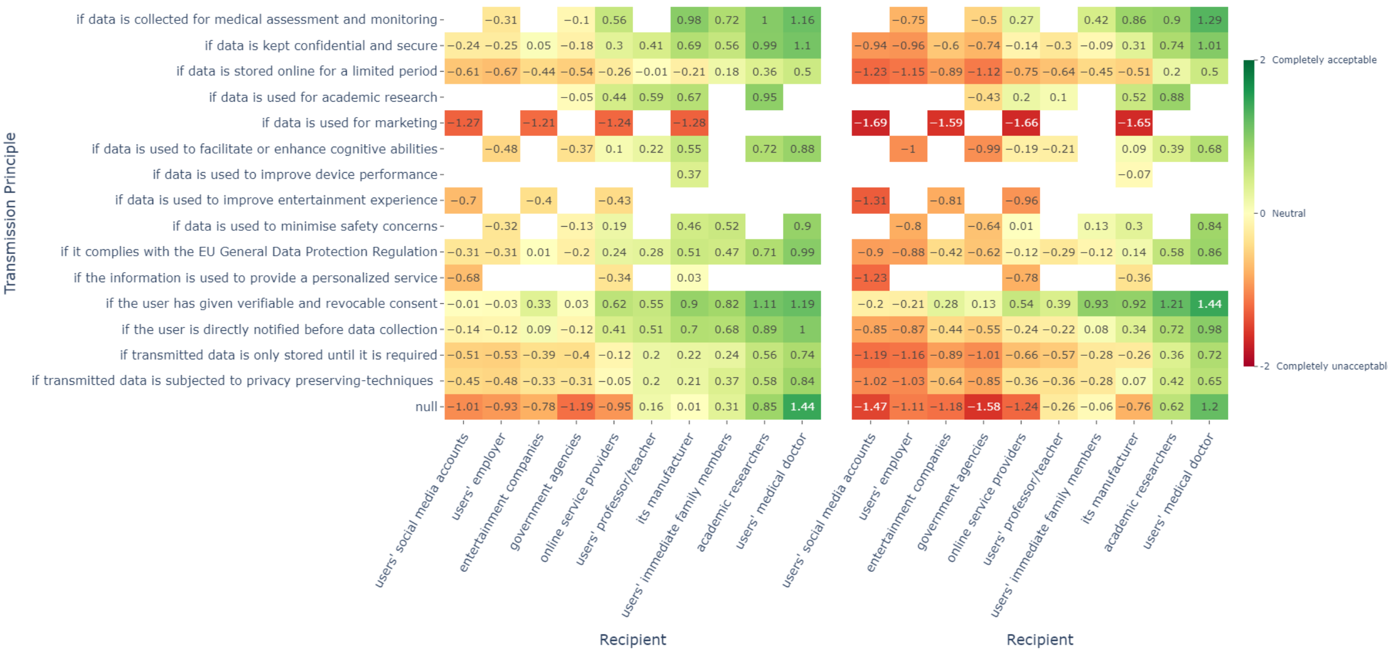

Figure 7: Heat maps that show average acceptability scores of data flows by transmission principles values and recipients for the different responses in regards to the information flows evaluation. The *left* heat map shows average scores if data flows from participants who assumed that consent was always given. The *right* map displays the values corresponding to participants who indicated that they evaluated the choices independently and did not assume that consent was given for every data flow

recipients in relation to the participants' educational or professional fields. The participants were grouped into two categories based on their educational or professional background. As mentioned above, more than 40% of the participants reported having an education or working in the field of IT. More than 50% of people without an IT background appear to be more concerned about sharing their brain data in total since we can observe more negative values. The average acceptability score over all neurodata flows for participants who have an IT background is 0.001, while the score of the participants with an IT background is -0.08. This difference is statistically significant ($p < 0.0083$). In summary, participants who have a strong affinity for technology are more concerned about sharing brain data with IT companies.

4.1.9 Wearable Device Users show a Greater Level of Openness. As indicated in Appendix A, 43% the of participants reported using wearable devices. When comparing this group to participants who do not use wearable devices, the scores reveal that wearable device users generally exhibit a greater willingness to share their brain signals. Their average acceptability score for all data flows is 0.03, whereas the score for non-wearable device users is -0.11. The difference is statistically significant ($p < 0.0083$). Hence, wearable device users show a greater level of openness towards neurodata sharing.

Looking at BCI knowledge, 55% of the participants reported that they had not heard about BCIs before. However, comparing this group and the participants who own, use, or had heard about BCIs before, the acceptability scores of -0.04 and -0.07, respectively, are very similar, and not statistically significant.

4.1.10 Regional Factors Influence Perceptions. In this section we examine the acceptability of sharing brain signals among participants from diverse countries, considering various receivers. We indeed observed variations between countries, comparing acceptability scores for conditional and unconditional data flows. However, we did not conduct statistical testing given that the number of participants for some countries was too low.

Differences can be observed in the scores of data flows between Greek residents and individuals from Poland. The heat maps show that respondents from Greece have a highly negative attitude towards sharing brain data, especially if data is stored online for a limited period or collected by government agencies while the unconditional acceptability score for medical doctors is significantly high, as well as when consent is given. In contrast to that, participants from Poland are more inclined to share their brain signals in general and if it is not used for marketing. Residents of Italy show a greater aversion to the flow when data is used for marketing purposes.

Participants from Spain showed a higher willingness to share brain signals with their immediate family members than participants from other countries. Individuals from Germany exhibit a low level of willingness to share their brain signals with social media platforms, entertainment companies, and government agencies.

These observations signal that it is essential to consider country-specific backgrounds to account for differences in acceptability when designing BCIs and privacy policies, as cultural and regional factors may influence users' perceptions of privacy and trust.

4.2 RQ2 - Neuroprivacy and Neurotechnology Awareness

Neuroprivacy Awareness. All the n=347 participants answered the question *"What do you think can be derived about a person from their brain signals collected by a commercial BCI?"*. We discarded 13 responses that did not provide an actual answer and then coded the rest of the dataset, obtaining a Cohen's κ = 0.85. Our final taxonomy of brain data inferences as perceived by participants comprises nine categories. The most salient inference types per category are detailed in the following. We give the frequency of the category and code appearance between parentheses.

The most frequent type of data participants think can be derived from brainwaves are **Mental States (26%)**. They refer to mental/cognitive states, moods, feelings, and emotions in general (13.8%), and mention 19 concrete feelings. The most common of these specific inferences are likes/dislikes and interests of the BCI user. The next top category is **Mental Processes (18.8%)**. Participants believe that brain reactions to stimuli can be tracked (6.8%) and they specified examples of other mental processes that they consider deducible from brainwaves, such as the processes of thinking and decision-making, desires, attention, and the possibility to determine if a person is lying. Remarkably, that several participants speak about inferring *true* desires or *real* opinions (as opposed to what a person shows to the outside) and detecting lies, all of which pose an important threat against mental freedom [15].

> *"I think the real opinion on a subject can be derived from the user, if he is lying or not for example, or his deepest desires." (P488)*

Also notably, several answers mentioned inferences of mental processes/states in relation to products and ads, denoting both awareness and concern about neuromarketing techniques, a trend which is supported by the quantitative acceptability scores in Section 4.1:

> *'I think that someone could get information what someone could feel (like emotion) when seeing or using given product. [...] and that can be harmful, for example for an overweight person who thinks about food most of the time, that person will receive more food ads and that works against that person well-being" (P132)*

The third top category of what participants think can be derived from brainwaves refers to (more complex) **Higher Order Inferences (15.5%)**. This includes health status (6.5%), behavior (1.29%), personality (1.6%), and commands to control a machine (1.1%), as the most commonly mentioned types of data. Additionally, many people responded with **Unspecific Brain Data (11%)**, such as brain signals. Together with the **Physiological Data (9.3%)** answers, we have the 5 categories that cover 80.8% of the codes in the dataset. This latter type of inference encompasses mostly states of stress (3.2%) and fatigue (2.3%), as well as sleep-related information (1.1%).

Few participants consider that it is possible to extract **Mental Constructs (7%)**, such as thoughts, visual representations in our minds, and memories. Interestingly, recent advances in large language models and artificial intelligence are facilitating this type of decoding using non-invasive brain readers [46]. The remaining answers either explicitly mentioned **Don't Know (5.7%)**, referred to the **Amount of Data (5.2%)** – in several degrees from nothing can be inferred to everything – or indicated the potential extraction of **Personal Information (1.1%)**, such as age or gender. Besides the overall feeling of participants that inferences can be used for marketing, their responses show insecurity and fear upon the uncertainty of what can be learned from brainwaves:

> *"That's the point... I don't know very clearly the answer to this question and I am afraid of the output" (P435)*
>
> *"I have no idea but it scares me." (P307)*

Neurotech Awareness. Seven participants reported having used and/or owning a BCI. In the group of non-owners (6), three of them used a BCI device in an academic context (for research, in a lecture), one participant used a BCI-controlled game, another tried the technology briefly during a fair, and the last one reported they used it to record brain activity. Just one person owns and regularly uses a BCI device, in this case for personal mental health purposes. Four respondents mentioned the following specific scenarios where they would use a BCI in the future: health, improve cognitive abilities, research, and if beneficial.

In the case of current non-users, they also reported whether or not, how, and why they would use a BCI. We coded the answers from the 336 participants who responded, obtaining a Cohen's κ = 0.81. The majority of the responses (77.6%) show a positive attitude towards using neurotechnologies, and 33.9% contain a negative opinion [7].

It is to be noted that 5% of the answers are from participants who are unsure if they would use or not the technology because they need more information to make a decision. For positive participants, the most appealing use cases are medical use (for diagnosis, health control) and research, where participants see a clear benefit and trust the recipients. These observations closely match the quantitative acceptability ratings that also favor these data flows. For example, participants would like to contribute to research studies that can help increase knowledge about the brain and benefit others and society.

> *"I would like to use BCI headset. I definetly would use it for medical reasons and science / research. I would agree to use BCI for medical reasons to get better diagnosis and to make doctors understand me better in situations when I don't know how to precisely describe my problem. My reasons for usage of BCI for science / reasearch are simple - i Would love to be part of developing, exploring and making a word a better place." (P148)*

The answers also suggest a strong interest of participants in using BCIs for self-monitoring, to better understand themselves (what stresses them, how they learn) and improve their lives, as well as for entertainment. The main driving reason for participants to use neurotechnology is curiosity.

In summary, use cases for personal and social good are favored. Yet positive participants expressed concern about privacy and the need to safeguard neurodata and make ethical use of the technology. Several respondents pointed out the potential benefits of the technology if they could use it in a private manner:

[7] Note that the percentages do not add up to 100% because some participants voiced both negative and positive attitudes.

> *"I would maybe be interested in using it in certain scenarios, for example at work, to monitor my brain activity (without sharing the data with my employer)." (P262)*

In the group of participants that are negative towards using neurotechnology, the most prominent reason is privacy and that they do not see how BCIs would be useful for them.

> *"Never. I don't want anybody to search my mind and my innermost thoughts and feelings, not even the loved ones." (P421)*

> *"I would not use a BCI headset, because I believe i don't need it in my life. I'm also concered about the safety of my personal data that would get recorded with the headset." (P221)*

Though participants voiced other concerns, such as not being sure about how safe/harmful or mature are current neurotechnologies, these were far less common than the worry about privacy.

To get further insights, we look at the interplay between participants' attitudes about neurotechnology, their perceptions about what can be derived from brain data, and the acceptability scores they gave to neurodata flows. For this purpose, we divide participants into two groups, those who rated all neurodata flows as acceptable (>=0) on average, and those who rated flows as unacceptable (<0). The distribution of perceived inferences on brain data between both groups does not differ significantly: they seem to think similarly, identifying emotions, moods, and feelings as the most common inferable type of data. In turn, when comparing the answers on neurotech awareness, significant differences arise ($\chi^2 = 60.76398, p < 0.05$) between the two groups. Participants who rated neurodata flows on the low side of acceptability are more negative regarding BCI technology, more frequently report not conceiving a use for it, and point at privacy as a reason in higher numbers that the group that rates neurodata flows as acceptable (See Fig. 8). This analysis surfaces the importance of privacy and utility as important factors for acceptance.

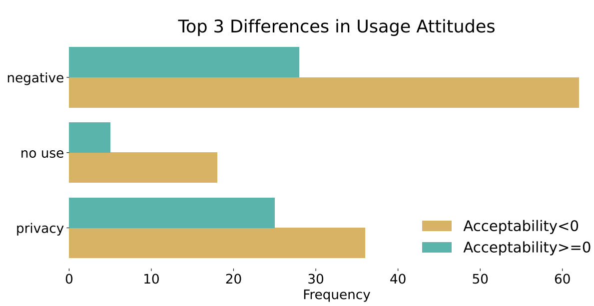

Figure 8: Biggest differences in neurotechnology attitudes for participants who consider neurodata flows as acceptable (on average) versus those who consider them unacceptable

5 DISCUSSION

This section discusses our findings regarding related work, identifies open challenges and recommendations for different stakeholders, and presents the limitations of our study.

5.1 Privacy Attitudes: BCIs vs other Technologies

As mentioned in Section 2.2, several works investigated users' privacy perceptions of collecting and sharing data by different devices. We now reflect on the parallels and differences with our findings.

Apthorpe et al.'s study to discover privacy norms of IoT devices [3] revealed that sharing under the principle of giving consent got the highest average acceptability scores. On the contrary, sharing IoT-collected data about the user for advertising, and sharing with government intelligence agencies or social media accounts had the lowest score. In a second study, Apthorpe et al. [4] analyzed privacy norms of parents whose children use smart IoT toys, following the same approach with similar results.

Comparing the demographics and background of participants, their results showed that younger parents and parents who own smart devices are generally more accepting, but education had minimal influence on the scores of acceptability.

Our findings are aligned with these results: consent is highly relevant to accepting neurodata flows (highest average acceptability score), and "marketing purposes" is the factor that leads to the lowest acceptability. Sharing data with government intelligence agencies and social media also falls on the low end of acceptability. Regarding demographics, similar to [4], young adults are less concerned with sharing their brain signals, and so are wearables device owners. A contrasting finding was that our participants with advanced academic degrees showed greater willingness to share their brain data. The reason could be related to the use cases for which they envisioned BCI usage, especially research, which can be favored by this sector of the population.

In the realm of VR, Adams et al.'s work [2] addressed the security and privacy perceptions of VR users and developers. The majority of participants in their interview-based study had concerns about sharing data from the numerous sensors with the device manufacturer, which is also a common concern for our participants regarding BCI technology.

With a broader scope, Motti & Caine aimed to comprehend the privacy-related worries of wearable devices users [33], for both wrist-mounted devices (WMDs) like smartwatches and head-mounted devices (HMDs), such as smart glasses and AR/VR devices. In both cases, similar to our study and prior work, users were concerned about the lack of control and awareness regarding information access, these concerns being higher for the HMD case. The study did not provide results specifically for EEG devices, so our results complement this line of work, reinforcing the importance of awareness and consent.

Furthermore, the data gathered from sensors of WMDs and HMDs (excluding EEG) differ significantly when compared to the information collected by EEG devices. Our qualitative analysis shows that participants believe that a significant quantity of sensitive information can be derived from brain signals collected by commercial BCIs, especially related to mental states, which makes privacy a strong concern.

On a different type of sensor, Steil et al. [43] conducted a cross-country survey about privacy attitudes towards eye-tracking. Once more, using data for marketing is not liked by users. Interestingly, this study also shows a positive attitude towards health and research use cases.

Finally, a unique observation in our survey is that curiosity is a strong driver influencing the will to use BCIs, even when being aware of all potential data contained in EEGs. Therefore, it is important to make people aware of the trade-offs, so they can make fully informed decisions before adopting the technology.

5.2 Lessons learned and Open Challenges

Our qualitative analysis reveals a great interest in neurotechnology but also concern with regard to the purposes for neurodata collection and its usage. Feelings of discomfort, distrust, and lack of control can hinder the adoption of this technology for highly beneficial use cases. Additionally, though our participants demonstrated an intuition on what can be extracted from brainwaves, a share of respondents did not know or reported only vague information. Especially relevant is the low number of people who believe that thoughts, and other specific mental constructs, such as images (e.g., what we are watching while wearing the BCI) can be decoded from brain signals. While commercial technology is not there yet, it does not look far considering recent progress in AI. On these grounds, we posit that developers and researchers should work on the design of appropriate **transparency mechanisms** to communicate to BCI users what the collected data and privacy risks are. For example, extending research on nutrition privacy labels [26] to accommodate the complexity added by neurosensors.

Another future research line to explore is understanding the **alignment of users' privacy expectations with the current industry practices** followed by neurotechnology companies and services. This type of systematic analysis would identify the size of the current gap and the main points where privacy protections are required.

Another challenging research to undertake based on our insights, is the exploration of **novel consent mechanisms** that go beyond unusable privacy policies. Already inefficient in other scenarios, they could become even more complex when including neurodata-related clauses. If no usable approaches emerge, we risk the commodification of neurodata, flowing and being shared with other pieces of information, and threatening our freedom and self-determination to a new extent.

Additionally, one of the main lessons learned from our study came from the practical application of the CI survey methodology. In this process, we observe important differences in how users make judgments when rating information flows. Specifically, we detected that consent is implied and assumed by many raters (around half in our study), affecting the acceptability score (leaning to more positive values). While this might be the result of studying a population (Europeans) that is influenced by being subject to GDPR protection for several years, it is interesting to observe that interpretations of principles can be biased. Based on this, when using this method, **we recommend giving clear instructions to CI-survey respondents on how to rate flows**, or to collect further qualitative information about how they make decisions. This will support an unbiased understanding of privacy acceptability. Another recommendation is to **conduct cross-cultural studies** to understand how cultural and regional differences impact users' neuroprivacy perceptions. Furthermore, the neurodata CI survey can be extended with new principles and recipients to probe for user perceptions regarding novel use cases or to verify if current regulations align or violate users' privacy norms, triggering changes if necessary.

5.3 Factors Influencing Acceptability of BCIs

We have found that sharing brain signals for medical or academic purposes is generally more acceptable to people than sharing them with commercial or government agencies. This is likely due to the trust placed in healthcare professionals and researchers, as well as the potential benefits to personal health and well-being.

The participants show indecisiveness in sharing brain signals with family members, professors/teachers, and BCI manufacturers which may stem from multiple factors. One possible reason could be that participants may not fully comprehend the implications of sharing their data with these parties.

It is also essential to acknowledge that the level of BCI knowledge, age, and education can impact individuals' willingness to share brain signals. For example, those with higher educational degrees tend to be more favorable towards sharing brain signals with academic researchers, while older age groups exhibit stronger negative values for sharing BCI data with various entities, suggesting heightened privacy concerns among older individuals.

Further, consideration must be given to how cultural and regional differences may impact users' perceptions of privacy and trust when sharing brain data.

Understanding the factors influencing the acceptability of sharing brain signals is another crucial aspect for developing future BCI applications and policies that cater to users' needs and concerns. By addressing privacy concerns and building trust in the handling of sensitive BCI data, we can pave the way for more widespread adoption and integration of neurotechnology into various aspects of our lives.

5.4 Exploring Acceptability of Transmission Principles

Participants were found to be more open to sharing BCI data when they had given verifiable and revocable consent. This finding emphasizes the importance of user control and consent in the acceptance of BCI technology. By allowing users to provide informed consent and retain control over their data, trust in BCI technology can be fostered, potentially increasing its adoption and acceptance. Furthermore, participants were found to be more inclined to share data if it was explicitly mentioned that the data would not be used for marketing purposes. This suggests that transparency and specific data usage play a critical role in shaping user attitudes toward sharing brain signals.

In conclusion, understanding these trends in the acceptability of transmission principles can guide the development of user-centered policies and regulations surrounding BCI data sharing. By addressing privacy, consent, and data usage concerns, developers and policymakers can help create an environment that fosters trust in BCI technology and encourages its adoption.

5.5 Limitations

Method. Our study, as with any survey-based research, comes with the limitation that questions are self-reported. We used quality filters (attention checks and time measurements) to remove bad answers.

Sample. We recruited from EU countries to provide a European perspective in addition to the US-centric perspectives mostly dominant in the literature. However, our sample is not representative per country and we can only get general insights from the common point of users that are familiar with the GDPR regulations. This needs to be taken into account, as results may not generalize to other populations. Several studies evidenced differences in privacy sensitivity between citizens of different countries, pointing at Europeans as more privacy concerned [8, 48]. These differences are dynamic, evolve with the social and regulatory environment, and can influence the acceptability scores given to neurodata flows.

Analysis-IUIPC. Although we collected IUIPC to investigate the impact of privacy concerns on acceptability scores, the responses were predominantly provided by users with high levels of privacy concern. The score of 284 out of 287 participants was laying above 3.5 on a 7-point Likert Scale where 1 indicates *strongly disagree* and 7 *strongly agree*, following the analysis methods of Malhotra et al. [28]. As a result, we were unable to examine the potential effect of different privacy attitudes on the acceptability of neurodata flows.

Scope. We capture the privacy norms of people who are mostly non-users (or even never heard) about neurotechnology. For a more comprehensive view, it would be interesting to study early adopters and heavy users of BCIs, analyzing their decision-making process to adopt the technology, usage patterns, and privacy attitudes. Similarly, our study is limited in the scope of the type of brain data being analyzed, we chose brain signals but it would be interesting to conduct further investigations on concerns about specific attributes derived from them, and the combination with other sensor data.

6 CONCLUSION

Emerging neurotechnology-based applications, collect, share, and analyze sensitive brain data, raising privacy concerns. So far, research has focused on ethical aspects, specified the need to establish new privacy rights, and investigated potential attacks on BCIs and technical countermeasures. In this work, we studied the perspective of those who are more directly affected by potential privacy violations in neurotechnology: its users. Our study underscores the importance of purpose, consent, and transparency regarding neuroprivacy. Our collection of privacy norms can serve as a basis to identify gaps in current neuroprivacy practices and support organizations, developers, policymakers, researchers, and privacy advocates in contributing to the privacy-respecting advancement of neurotechnologies.

ACKNOWLEDGMENTS

We thank the anonymous reviewers for their insightful feedback that helped us improve the paper. We also thank the students that supported us in setting up the study: Levi Kletetzka for executing the pilot, and Salman Shabbir, for collecting the final data.

REFERENCES

[1] [n. d.]. *Galea: Bringing next generation neurotechnology to mixed reality.* https://galea.co
[2] Devon Adams, Alseny Bah, Catherine Barwulor, Nureli Musaby, Kadeem Pitkin, and Elissa M. Redmiles. 2018. Ethics Emerging: the Story of Privacy and Security Perceptions in Virtual Reality. In *Fourteenth Symposium on Usable Privacy and Security (SOUPS 2018)*. USENIX Association, Baltimore, MD, 427–442. https://www.usenix.org/conference/soups2018/presentation/adams
[3] Noah Apthorpe, Yan Shvartzshnaider, Arunesh Mathur, Dillon Reisman, and Nick Feamster. 2018. Discovering smart home internet of things privacy norms using contextual integrity. *Proceedings of the ACM on interactive, mobile, wearable and ubiquitous technologies* 2, 2 (2018), 1–23.
[4] Noah Apthorpe, Sarah Varghese, and Nick Feamster. 2019. Evaluating the Contextual Integrity of Privacy Regulation: Parents' IoT Toy Privacy Norms Versus COPPA. In *28th USENIX security symposium (USENIX security 19)*. 123–140.
[5] Patricia Arias-Cabarcos, Thilo Habrich, Karen Becker, Christian Becker, and Thorsten Strufe. 2021. Inexpensive Brainwave Authentication: New Techniques and Insights on User Acceptance. In *30th USENIX Security Symposium (USENIX Security 21)*. USENIX Association, 55–72. https://www.usenix.org/conference/usenixsecurity21/presentatiogenaron/arias-cabarcos
[6] Richard A Armstrong. 2014. When to use the Bonferroni correction. *Ophthalmic and Physiological Optics* 34, 5 (2014), 502–508.
[7] New York City Bar Association et al. 2005. Are your thoughts your own? Neuroprivacy and the legal implications of brain imaging. *New York: The Committee on Science and Law* (2005).
[8] Steven Bellman, Eric J Johnson, Stephen J Kobrin, and Gerald L Lohse. 2004. International differences in information privacy concerns: A global survey of consumers. *The Information Society* 20, 5 (2004), 313–324.
[9] Sergio López Bernal, Alberto Huertas Celdrán, Gregorio Martínez Pérez, Michael Taynnan Barros, and Sasitharan Balasubramaniam. 2021. Security in brain-computer interfaces: state-of-the-art, opportunities, and future challenges. *ACM Computing Surveys (CSUR)* 54, 1 (2021), 1–35.
[10] Tamara Bonaci, Ryan Calo, and Howard Jay Chizeck. 2014. App stores for the brain: Privacy & security in Brain-Computer Interfaces. In *2014 IEEE International Symposium on Ethics in Science, Technology and Engineering*. IEEE, 1–7.
[11] Jacob Cohen. 1960. A coefficient of agreement for nominal scales. *Educational and psychological measurement* 20, 1 (1960), 37–46.
[12] Rig Das, Emanuele Maiorana, and Patrizio Campisi. 2016. EEG biometrics using visual stimuli: A longitudinal study. *IEEE Signal Processing Letters* 23, 3 (2016), 341–345.
[13] Emotiv Systems. 2023. Emotiv website. https://www.emotiv.com/. Accessed: 03/05/2023.
[14] EU. 2016. General Data Protection Regulation. (2016). Available at: https://gdpr-info.eu/ (accessed: 15.09.2022).
[15] Nita A Farahany. 2023. *The Battle for Your Brain: Defending the Right to Think Freely in the Age of Neurotechnology*. St. Martin's Press.
[16] Joseph L Fleiss, B Levin, and M Cho Paik. 1981. *Statistical methods for rates and proportions*. Vol. 1981218. New York: Wiley.
[17] The Neurorights Foundation. 2023. Market Analysis NEUROTECHNOLOGY. https://neurorightsfoundation.org/. Accessed: 03/05/2023.
[18] Mario Frank, Tiffany Hwu, Sakshi Jain, Robert T Knight, Ivan Martinovic, Prateek Mittal, Daniele Perito, Ivo Sluganovic, and Dawn Song. 2017. Using EEG-based BCI devices to subliminally probe for private information. In *Proceedings of the 2017 on Workshop on Privacy in the Electronic Society*. 133–136.
[19] Frances Grodzinsky and Herman T Tavani. 2010. Applying the "contextual integrity" model of privacy to personal blogs in the blogoshere. (2010).
[20] Qiong Gui, Maria V Ruiz-Blondet, Sarah Laszlo, and Zhanpeng Jin. 2019. A survey on brain biometrics. *Comput. Surveys* 51, 6 (2019), 1–38. https://doi.org/10.1145/3230632
[21] Lorena Guzman H. 2022. UNESCO Website. https://en.unesco.org/courier/2022-1/chile-pioneering-protection-neurorights. Accessed: 03/05/2023.
[22] Simon Hanisch, Patricia Arias-Cabarcos, Javier Parra-Arnau, and Thorsten Strufe. 2021. Privacy-Protecting Techniques for Behavioral Data: A Survey. *arXiv preprint arXiv:2109.04120* (2021).
[23] Marcello Ienca and Roberto Andorno. 2017. Towards new human rights in the age of neuroscience and neurotechnology. *Life sciences, society and policy* 13, 1 (2017), 1–27.
[24] InteraXon. 2022. InteraXon Website. https://choosemuse.com/. Accessed: 03/05/2023.
[25] Nader Karamzadeh, Yasaman Ardeshirpour, Matthew Kellman, Fatima Chowdhry, Afrouz Anderson, David Chorlian, Edward Wegman, and Amir Gandjbakhche. 2015. Relative brain signature: a population-based feature extraction procedure to identify functional biomarkers in the brain of alcoholics. *Brain and behavior* 5, 7 (2015), e00335.
[26] Patrick Gage Kelley, Joanna Bresee, Lorrie Faith Cranor, and Robert W Reeder. 2009. A" nutrition label" for privacy. In *Proceedings of the 5th Symposium on Usable Privacy and Security*. 1–12.

[27] Jingxin Liu, Hongying Meng, Asoke Nandi, and Maozhen Li. 2016. Emotion detection from EEG recordings. In *2016 12th international conference on natural computation, fuzzy systems and knowledge discovery (ICNC-FSKD)*. IEEE, 1722–1727.

[28] Naresh K Malhotra, Sung S Kim, and James Agarwal. 2004. Internet users' information privacy concerns (IUIPC): The construct, the scale, and a causal model. *Information systems research* 15, 4 (2004), 336–355.

[29] Anuradha Mandal and Nitesh Saxena. 2022. SoK: Your mind tells a lot about you: On the privacy leakage via brainwave devices. In *Proceedings of the 15th ACM Conference on Security and Privacy in Wireless and Mobile Networks*. 175–187.

[30] Ivan Martinovic, Doug Davies, Mario Frank, Daniele Perito, Tomas Ros, and Dawn Song. 2012. On the Feasibility of {Side-Channel} Attacks with {Brain-Computer} Interfaces. In *21st USENIX Security Symposium (USENIX Security 12)*. 143–158.

[31] Nick Merrill, Max T Curran, and John Chuang. 2017. Is the Future of Authenticity All In Our Heads? Moving passthoughts from the lab to the world. In *Proceedings of the 2017 New Security Paradigms Workshop*. 70–79.

[32] Matthew B Miles and A Michael Huberman. 1994. *Qualitative data analysis: An expanded sourcebook*. Sage.

[33] Vivian Genaro Motti and Kelly Caine. 2015. Users' privacy concerns about wearables: impact of form factor, sensors and type of data collected. In *Financial Cryptography and Data Security: FC 2015 International Workshops, BITCOIN, WAHC, and Wearable, San Juan, Puerto Rico, January 30, 2015, Revised Selected Papers*. Springer, 231–244.

[34] MyndPlay. 2023. MyndPlay website. https://myndplay.com/. Accessed: 03/05/2023.

[35] Neurosky. 2015. Neurosky Website. https://store.neurosky.com/pages/mindwave. Accessed: 03/05/2023.

[36] NeuroTechX. 2021. *The Neurotech Primer: A Beginner's Guide to Everything Neurotechnology*. Independently published.

[37] Helen Nissenbaum. 2004. Privacy as contextual integrity. *Wash. L. Rev.* 79 (2004), 119.

[38] Elizabeth O'Neill. 2022. Contextual integrity as a general conceptual tool for evaluating technological change. *Philosophy & Technology* 35, 3 (2022), 79.

[39] Stefan Palan and Christian Schitter. 2018. Prolific. ac—A subject pool for online experiments. *Journal of Behavioral and Experimental Finance* 17 (2018), 22–27.

[40] Qinwan Rabbani, Griffin Milsap, and Nathan E Crone. 2019. The potential for a speech brain–computer interface using chronic electrocorticography. *Neurotherapeutics* 16 (2019), 144–165.

[41] Pan Shi, Heng Xu, and Yunan Chen. 2013. Using contextual integrity to examine interpersonal information boundary on social network sites. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 35–38.

[42] Zuoting Song, Tao Fang, Jing Ma, Yuan Zhang, Song Le, Gege Zhan, Xueze Zhang, Shouyan Wang, Hui Li, Yifang Lin, et al. 2021. Evaluation and Diagnosis of Brain Diseases based on Non-invasive BCI. In *2021 9th International Winter Conference on Brain-Computer Interface (BCI)*. IEEE, 1–6.

[43] Julian Steil, Inken Hagestedt, Michael Xuelin Huang, and Andreas Bulling. 2019. Privacy-aware eye tracking using differential privacy. In *Proceedings of the 11th ACM Symposium on Eye Tracking Research & Applications*. 1–9.

[44] Shravani Sur and Vinod Kumar Sinha. 2009. Event-related potential: An overview. *Industrial psychiatry journal* 18, 1 (2009), 70.

[45] Yu Takagi and Shinji Nishimoto. 2022. High-resolution image reconstruction with latent diffusion models from human brain activity. *bioRxiv* (2022), 2022–11.

[46] Jerry Tang, Amanda LeBel, Shailee Jain, and Alexander G Huth. 2022. Semantic reconstruction of continuous language from non-invasive brain recordings. *bioRxiv* (2022), 2022–09.

[47] Zahra Tarkhani, Lorena Qendro, Malachy O'Connor Brown, Oscar Hill, Cecilia Mascolo, and Anil Madhavapeddy. 2022. Enhancing the security & privacy of wearable brain-computer interfaces. *arXiv preprint arXiv:2201.07711* (2022).

[48] Sabine Trepte, Doris Teutsch, Philipp K Masur, Carolin Eicher, Mona Fischer, Alisa Hennhöfer, and Fabienne Lind. 2015. Do people know about privacy and data protection strategies? Towards the "Online Privacy Literacy Scale"(OPLIS). *Reforming European data protection law* (2015), 333–365.

[49] Robert F Woolson. 2007. Wilcoxon signed-rank test. *Wiley encyclopedia of clinical trials* (2007), 1–3.

A APPENDIX: OPEN DATA

The survey instrument, detailed demographics and further analysis material (statistical test results, codebook, heat maps) are available at https://gitlab.com/hitsresearchgroup/neuroprivacy.