Brain-Hack: Remotely Injecting False Brain-Waves with RF to Take Control of a Brain-Computer Interface

Brain-Hack: Remotely Injecting False Brain-Waves with RF to Take Control of a Brain-Computer Interface

Alexandre Armengol-Urpi
armengol@mit.edu
Massachusetts Institute of Technology
Cambridge, MA, USA

Reid Kovacs
rkovacs@mit.edu
Massachusetts Institute of Technology
Cambridge, MA, USA

Sanjay E. Sarma
sesarma@mit.edu
Massachusetts Institute of Technology
Cambridge, MA, USA

ABSTRACT

The promise of Brain-Computer Interfaces (BCIs) is counterbalanced by concerns about vulnerabilities. Recent studies have revealed that EEG-based BCIs are susceptible to security breaches. However, current attack approaches are challenging to execute in real-world settings because they need access to, at a minimum, the EEG data stream. In this work, we introduce an unexplored vulnerability of current EEG-based BCIs that consists of remotely injecting false brain-waves into the recording device. We do this by transmitting amplitude-modulated radio-frequency (RF) signals that are received by the physical structure of the EEG equipment. We demonstrate the versatility of our system by successfully attacking three different categories of EEG devices: research-grade (Neuroelectrics), open-source (OpenBCI), and consumer-grade (Muse). We test our attack system by taking control of three different BCIs: a virtual keyboard speller, a drone-control interface, and a neurofeedback meditation interface. Our system was successful in each case, forcing the input of any desired character with the virtual keyboard, crashing the drone, and reporting false meditative states, respectively. To the best of our knowledge, this is the first time that an EEG device is remotely hacked at the physical layer. This work shows the risks that can arise from this type of attacks, which can not only be dangerous by seizing control of a BCI, but could also lead to severe misdiagnoses in clinical EEG tests.

CCS CONCEPTS

• **Security and privacy → Hardware attacks and countermeasures**; • **Human-centered computing → Human computer interaction (HCI)**; • **Hardware**;

KEYWORDS

Brain-Computer Interface; BCI; EEG; Radio-Frequency; SDR; SSVEP; Speller; Drone; Meditation; Vulnerabilities; Attack; IEMI

ACM Reference Format:
Alexandre Armengol-Urpi, Reid Kovacs, and Sanjay E. Sarma. 2023. *Brain-Hack*: Remotely Injecting False Brain-Waves with RF to Take Control of a Brain-Computer Interface. In *Proceedings of the 4th Workshop on CPS&IoT Security and Privacy (CPSIoTSec '23), November 26, 2023, Copenhagen, Denmark.* ACM, New York, NY, USA, 14 pages. https://doi.org/10.1145/3605758.3623497

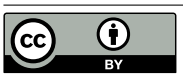

CPSIoTSec '23, November 26, 2023, Copenhagen, Denmark

ACM ISBN 979-8-4007-0254-9/23/11.
https://doi.org/10.1145/3605758.3623497

1 INTRODUCTION

A Brain-Computer Interface (BCI) provides users with a direct connection between their brain and an external device, such as a computer, wheelchair, or robot. The electroencephalogram (EEG) has become the most common input signal for BCIs, because of its low cost and ease of use. BCIs initially emerged as tools to help disabled people restore their damaged hearing, sight, movement, or communication capabilities. Current development of BCIs is focused on expanding its use cases to enable new and innovative forms of communication between brains and devices. For example, some BCIs allow users to correct robot mistakes [1], play video games [2], improve commercial eye trackers [3], or create hands-free interfaces for virtual [4] or augmented reality [5]. Recent advances in hardware design have led to the appearance of consumer-grade BCI devices, tailored to assist users with meditation [6], focus [7, 8], and IoT interaction [9].

This trend toward the development of broader BCI applications seems to indicate that BCIs will be increasingly present in our lives. This projection, however, is in contrast with the surprisingly little amount of research addressing the security of these systems. Security attacks on BCIs can cause harmful consequences, especially if the BCI is devoted, for instance, to vehicle safety [10], wheelchair control [11], drone control [12], prosthetics [13], or patient communication [14, 15].

Figure 1 illustrates the typical building blocks of a BCI. First, the neural activity is captured by the EEG recording device. Then, the signals are transmitted to a device that processes them by extracting useful features and classifying them into a specific command. Finally, the command is sent to an external device or application and the action is manifested. Usually, the loop is closed when visual feedback is given to the user. As shown in Figure 1 and as explained later in Section 2, all the literature related to BCI security vulnerabilities focuses on the latter blocks of the BCI systems, and none of them addresses potential attacks to the physical layer, such as the hardware dedicated to signal acquisition. Moreover, existing research assumes the attacker has access to either the raw EEG data, the classification model, or the external device that provides stimuli or feedback to the user. Our work bypasses all these assumptions by remotely compromising the physical layer of the BCI.

In this work, we show an unexplored vulnerability of current BCIs by remotely injecting fake brain-waves into the EEG recording device. We do this by sending Radio-Frequency (RF) signals that are received by the EEG equipment and inserted into the BCI system. To the best of our knowledge, this is the first time that a BCI is remotely compromised at the physical layer. Although we frame this article in the context of BCIs, the same vulnerability exists in

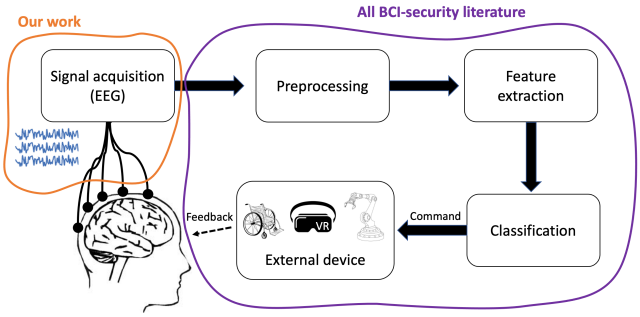

Figure 1: Workflow of a common Brain-Computer Interface. In purple we enclose the parts of the BCI whose security vulnerabilities have been studied by related work. Our work, in contrast, focuses on interfering in the physical layer of the system, that is the signal acquisition block (encircled in orange).

other EEG-based setups, such as clinical EEG recordings, where an attack using our system could result in severe misdiagnosis.

In particular, we successfully attacked three different EEG systems: a research-grade device (Neuroelectrics Enobio [16]), an open-source device (OpenBCI Ganglion [17]), and a consumer-grade device (Muse 2 [18]). Additionally, we tested our hacking technique on three different BCI applications: a virtual keyboard speller, a drone-control interface, and a neuro-feedback meditation interface. With these examples, we show that our system can successfully take control of BCIs by typing any phrase in the speller, crashing the drone, and inducing a signal indicating a false meditative state, respectively.

To achieve this injection, amplitude-modulated (AM) signals are wirelessly transmitted to the EEG headset. The wires of the EEG headset act as unassuming antennas that receive the signal. Due to non-linearities in the amplifier response, the modulating frequency will be captured by the recording device and the BCI will read the injection as a neurological signal. If the transmitter is of sufficient power, the injected signal will exceed the power of the real neural activity of the user, obfuscating the true signal.

2 RELATED WORK

2.1 Privacy and Security of BCIs

In contrast with the growing popularity of BCIs in medical and non-medical areas, few studies have been conducted on the privacy and security of these systems. In terms of privacy, Martinovic et. al. [19] showed for the first time that private information – such as PIN numbers or area of living – can be retrieved from BCI users by analyzing their brain-wave responses to tailored visual stimuli. Frank et. al. [20] proposed a subliminal attack that can deduce confidential information by probing the victim below their cognitive perception and analyzing their brain activity. Earlier works have used EEG signals to study sexual orientation, religious beliefs and deviant sexual interests of the user [21–23].

In terms of BCI security, there are several studies that show how to modify the raw EEG data with small perturbations to take control of the BCI. Zhang et. al. [24] showed that P300 and Steady-State Visually Evoked Potentials (SSVEP) spellers are very vulnerable to adversarial perturbations. They show how an attacker could spell anything they want by adding custom noise to the EEG data. Similarly, Bian et. al. [25] propose to use square wave signals as adversarial perturbations to attack and control SSVEP-based BCIs. Meng et. al. [26] introduced the idea of using backdoor attacks to poison the training data of EEG-based BCIs. Test samples with the backdoor key are then classified into the target class determined by the attacker. Finally, studies from Zhang et. al. [27] and more recently from Liu et. al. [28] have shown the vulnerability of CNN classifiers in EEG-based BCIs. They demonstrated that these deep learning models can easily be fooled with EEG data contaminated with small deliberate perturbations.

There is another group of related work that identifies potential privacy and security flaws of BCIs but without actually implementing the attacks. They are mostly in the form of reviews or white papers and seek to raise awareness of the emerging risks of malicious brain-hacking. They analyze the different building blocks of BCIs and study how each of them could be attacked. Some of this related research also proposes countermeasures to prevent such attacks or strategies to mitigate the risks involved [29–36].

As explained above, all BCI attacks proposed in the related literature assume the attacker has access to either the EEG data in training, the EEG data in testing, the EEG feature classification model or the external device that provides feedback or visual stimuli to the victim. Our work, in contrast, does not require any of these assumptions, since the attack is performed remotely to the physical layer of the system.

2.2 IEMI Attacks on Analog Sensors

There are a variety of Electromagnetic Interference (EMI) noise sources that may interact with electronic devices through the induction of voltages on conductors. Systems designed to measure small analog signals are particularly sensitive to such noise sources, as the noise may obscure the desired measurement. In the case of EEG devices, EMI noise can impact the quality of the measurement [37]. This type of noise is often produced by power electronics, which are commonly subject to extensive EMI filtering. It is also common to include an input filter to remove unwanted frequencies from the sensor, but it is difficult to build a filter to remove all out-of-band signal [38].

However, this phenomenon can be leveraged to manipulate sensors with intentional electromagnetic interference (IEMI) [39]. Due to hardware required for analog sensing, this approach can provide an unimpeded path to the manipulation of an otherwise protected system. Many systems make critical choices based on such sensor readings, often using sensors as feedback for some form of actuation. Such sensors include devices to measure ambient conditions (temperature, humidity, pressure), for localization (altitude, GPS position), for human interaction (microphones, touch screens), and medical diagnostics (electrophysiology). IEMI attacks have been applied broadly to these types of systems. For example, they have been used to fake physical interactions with smartphone touchscreens [40], also to inject silent voice commands in smartphones

[41], to manipulate the image information captured by cameras [42, 43], or to induce spurious serial communications [44].

Some of these attacks deserve particular attention because they impact systems that are crucial for safety. For example, IEMI attacks have been used to spoof Anti-Lock Breaking Systems (ABS) in cars [45], to paralyze drones [46], to manipulate the temperature sensor measurements of infant incubators [47], to cause pacemakers to stop pacing [48], or to take control of PWM-controlled actuators and change the flight trajectory of unmanned aerial vehicles[49], for example.

Once the signal is induced into the system via the sensor, or the physical layer, the signal is typically amplified, conditioned, quantized, and recorded in the analog and digital layers[38]. If the frequency of the noise source is known, filters can be employed to remove the interference. It is common for sensing systems to employ filters to mitigate expected noise sources. However, if the IMEI is within the frequency range of the desired signal, it is impossible to remove through filtering, as a filter would also remove the desired signal. If the interference is at a significantly higher frequency than the measured signal, it is possible that the interference could bypass an input filter through high-frequency coupling. Finally, the corrupted signal will reach the amplifier. The high-frequency injection will experience non-linear effects within the amplifier [38]. In the case of an amplitude modulated signal, the non-linear effects of the amplifier leads to a form of demodulation, ultimately adding the modulating frequency to the baseband signal [48].

3 WORKING PRINCIPLE

In this section, we describe the different elements that take part in our attack approach and explain the effects that take place for it to succeed in the task of injecting fake brain-waves remotely.

3.1 Functioning of an EEG device

An electroencephalogram (EEG) is an electro-physiological measurement technique used to observe brain activity. Information is propagated and processed in the brain through electro-chemical processes that yield externally observable electric fields. These fields can be measured as the voltage induced in small metal electrodes positioned around the head.

The fields generated by the communication of these cells are quite small, on the order of tens of micro-volts when measured outside the head. Due to amplitude of these signals, amplifiers with significant gain must be employed. These systems also typically have many electrodes, each with their own wires which are highly susceptible to noise. This property is exactly what we take advantage of for this injection technique.

Typical EEG systems employ instrumentation amplifiers for each channel to overcome noise challenges and improve signal quality. This amplifier configuration has a differential input, meaning that the difference between the two input signals is amplified and the common-mode is rejected. In an EEG, the two inputs would be a recording electrode and a reference placed on bone. Typically, these amplifiers exhibit some non-linearities that we will exploit to inject our false brain-waves, as explained in Sections 3.2 and 3.5. All recording electrodes are typically amplified with respect to the same reference electrode.

In order to disrupt an EEG-based BCI system, the measured brain signals must be overshadowed by a larger injection. Due to the small magnitude of typical EEG signals, typical systems can be fooled with a relatively low-power injection that is harmless to a BCI user. It is demonstrated that this can be achieved wirelessly with radio-frequency transmission.

EEG electrodes are most commonly positioned on the head using a cap or helmet, but intracranial sensors have also been used since they offer improved signal quality [50–52]. All EEG devices used in this experiment are non-invasive, scalp electrodes. Although, it may be possible to achieve the same effect in intracranial electrode arrangements or fully implanted systems.

3.2 Non-Linear Amplifier Response

Instrumentation amplifiers such as those utilized in EEG devices exhibit non-linearities for frequencies outside their operating range. This is the characteristic that we will exploit in order to inject the false brain-wave signals into the EEG device. If the input signal is x, then the output of the amplifier can be expressed as follows:

$$x_{out} = \sum_{n=1}^{\infty} A_n x^n = A_1 x + A_2 x^2 + A_3 x^3 + ... \quad (1)$$

In theory, the non-linear output can be represented as an infinite power series. However, the third and higher order terms are negligible and can be disregarded. We find opportunities to exploit the system via the second order term, as seen in Section 3.5.

3.3 Reception of Radio-Frequency Signals

The main principal behind our hacking approach is to utilize the physical structure of the EEG device as a receiving antenna that captures remotely transmitted radio-frequency signals. In particular, the primary elements that are best-suited to act as receiving antennas are the exposed cables that link the EEG electrodes to the recording device. According to basic antenna theory [53], the length of an emitting or receiving antenna (cable length in our case) should be in the same order of magnitude as the wavelength of the desired signal to send or receive. Since brain signal frequencies are on the order of tens of Hertz [54], we can compute the required cable length by the electromagnetic radiation relation $\lambda = c/f$, where λ is the radiation wavelength, f is its frequency, and $c = 3 \times 10^8 m/s$ is the speed of light. If f is on the order of tens of Hertz, the resulting wavelength and, as a consequence, the required cable length to capture such signals is on the order of thousands of kilometers. Hence, sending RF energy at the frequency of brain activity is unfeasible because the length of the EEG cables is fixed and on the order of centimeters. Therefore, we need to resort to a different technique which is very common in the radio-communications field: amplitude modulation.

3.4 Amplitude Modulation

It is clear that in order to capture electromagnetic radiation at the EEG cables, the wavelength of such signals should be in the same order of magnitude as the length of the cables, as explained above. If these are, for example, 1 meter long, this translates to sending electromagnetic waves on the order of 300 MHz ($f = c/\lambda$). Since the goal is to inject signals into the device at about 1 to 60 Hz, we

implemented the solution of encoding the desired low-frequency signal by multiplying it with a fast oscillating wave (or carrier wave). In other words, the carrier wave is amplitude-modulated at the frequency of the desired signal injection. Mathematically, the radio signal sent can be modeled as:

$$x(t) = (1 + M(t))cos(2\pi f_c t + \phi) \tag{2}$$

where $M(t)$ represents the signal that modulates the carrier and that we want to inject into the EEG recording, and $cos(2\pi f_c t + \phi)$ is the carrier signal with frequency f_c. Since we want to inject periodic signals that are interpreted as oscillating brainwaves, $M(t)$ can be written as $M(t) = cos(2\pi f_i t + \phi)$. Therefore the signal that will be transmitted takes the form:

$$x(t) = cos(2\pi f_c t) + \cos(2\pi f_c t + \phi)\cos(2\pi f_i t) \tag{3}$$

and can be expanded to:

$$x(t) = cos(2\pi f_c t) + \frac{1}{2}\cos\big(2\pi(f_c - f_i)t + \phi\big) + \frac{1}{2}\cos\big(2\pi(f_c + f_i)t + \phi\big) \tag{4}$$

where f_i is modulating frequency and the frequency of the signal that is injected into the EEG device. In the frequency domain, $X(f)$ would be represented as seen in Figure 2.

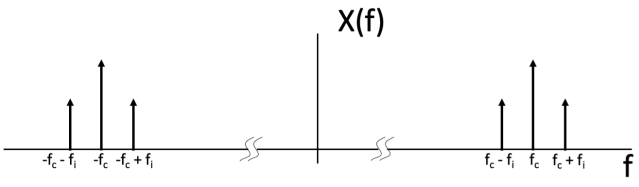

Figure 2: Representation of the transmitted amplitude modulated radio signal in the frequency domain.

3.5 AM signals in Non-Linear Region of Amplifiers

As described in section 3.2, the amplifiers within an EEG device exhibit a non-linear response outside of their intended operating range. After passing through the amplifier stage of the EEG device, the measured signal can be modeled as follows:

$$x_{out} = A_1\left[\big(M(t) + 1\big)C(t)\right] + A_2\left[(M(t) + 1)C(t)\right]^2 \tag{5}$$

where $M(t)$ refers to the modulating signal and $C(t)$ refers to the carrier signal. The first order term yields the expression in equation 3. The frequencies from this term are well beyond the operating range of a typical amplifier for this application. The second order term contains a multiplication of the included signals, resulting in frequency components at $2f_c$, $2f_i$, $2f_c \pm 2f_i$, $2f_c \pm f_i$, and, most importantly, $\pm f_i$. Demonstrated mathematically,

$$\begin{aligned} x_{out,2} &= A_2\left[(M(t) + 1)C(t)\right]^2 \\ &= A_2\Big[\cos(2\pi f_c t) + \frac{1}{2}\cos\big(2\pi(f_c - f_i)t + \phi\big) \\ &\qquad + \frac{1}{2}\cos\big(2\pi(f_c + f_i)t\big)\Big]^2 \\ &= \frac{A_2}{2}\cos\big(2\pi(f_c - (f_c - f_i))t\big) + \text{(higher frequency terms)} \\ &= \frac{A_2}{2}\cos\big(2\pi f_i t\big) + \text{(higher frequency terms)} \end{aligned}$$

All of the frequencies in the received signal are filtered out or aliased due to the relatively slow sampling of the EEG device, except for the signal at f_i. Thus, the modulating frequency is resolved by the EEG device and can be easily sampled, as these devices have a sampling rate on the order of 500Hz.

This resulting injection may still be present despite the implementation of an input EMI filter. It is difficult to design an EMI filter that effectively removes all frequencies above a desired cutoff. Due to high-frequency coupling and parasitic elements, the filter may still pass signals that are orders of magnitude above the desired cutoff [55], as in this attack.

3.6 Geometric Dependence of the Optimal Carrier Frequency

When EEG headsets are worn, the electrode cables are not always perfectly straight, in fact, they usually form small loops or get tangled. This changes the geometry of the "receiving antenna" of the device, and, as a consequence, the optimal carrier frequency also changes. Therefore, the configuration of the cables dictates the required carrier frequency for a successful injection. It follows that each time the device is used, the optimal carrier frequency will slightly vary. In general, the functioning carrier frequency is not unique and there is a range of usable carriers. This calibration is not needed for devices that have a fixed structure without cables, such as the Muse headband. In Figure 3, we show an example of how shifting the carrier frequency changes the amplitude of the injected waves for a particular cable arrangement of the Neuroelectrics Enobio headset. We can see there is a wide band where the transmission is successfully injected into the device. Specifically, the system has a greater response (greater power of injected signal) at around 500 - 550 MHz. We determined that this optimal range did not change greatly between experiments, but some adjustment was necessary if the maximum power transmission is desired.

4 SYSTEM DESIGN

Here we present the design and usage of the *Brain-Hack* system. We describe the components of the system and the EEG devices that are compromised. Additionally, we provide an overview of the system and the technique used for identifying optimal settings.

4.1 Materials

4.1.1 EEG Devices. This section describes the EEG systems used to test and validate the injection technique described in this work.

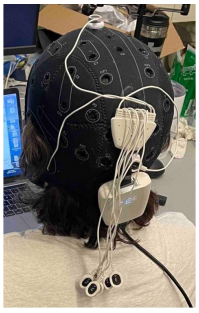

(a) Neuroelectrics EEG device with a particular arrangement of its electrode cables.

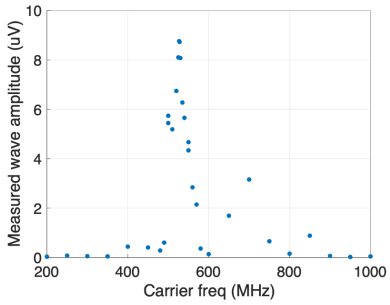

(b) Injected wave amplitude as a function of the carrier frequency. We can see there is a large band where the sent signal is successfully injected.

Figure 3: Geometric dependence of a usable carrier frequency. On the EEG cap, only channel 1 is connected to position Oz. Reference and ground are placed in FCz and left mastoid respectively.

Neuroelectrics Enobio. The Neuroelectrics Enobio EEG headset is a device designed as a research tool [16]. It may be found as a part of a psychology or neuroscience study. Compromising this device would enable an attacker to influence the results of a study or inject concerning signals into a clinical EEG analysis. In particular, we used Enobio 8, which has eight signal channels, and its electrode cables measure 32 cm.

OpenBCI Ganglion. OpenBCI designs and sells low-cost and open-source EEG equipment [17]. These headsets are typically used in simpler research applications or employed as a BCI. The Ganglion board has four signal channels, and one-meter electrode cables are used in these experiments. The OpenBCI platform may be used in low-cost BCI applications due to its price. For example, an attacker could assume control of household items [56] or even a wheelchair [11].

InteraXon Muse 2. The Muse is a device designed to provide consumers with neuro-feedback during meditation sessions [18]. A typical user would meditate with the device and app, allowing for the system to play sounds that indicate the level of mental activity. By compromising this system, an attacker would be able to influence the result of the meditation. For example, an attack could cause the app to suggest to the user that they never fell into a meditative state. The headband has 4 signal channels, with two sensors located in the frontal region (AF7 and AF8) and the other two in the temporal regions (TP9 and TP10). It has 3 reference sensors around location Fpz.

4.1.2 Software-Defined Radio. A Software-Defined Radio (SDR) is a radio-frequency (RF) communication system in which the components have been implemented in software using digital signal processing. Without an SDR, such RF transmissions would be more difficult to produce and vary as they would need to be implemented in analog circuitry. Typically, the software is deployed on a computer or embedded system. The SDR employed here is a Nuand

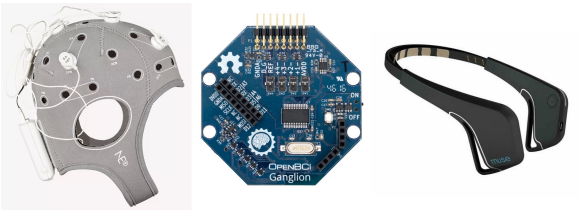

Figure 4: The 3 EEG recording devices used (and hacked) in this work. Neuroelectrics Enobio (left), OpenBCI Ganglion (middle), and InteraXon Muse 2 (right).

BladeRF micro xA9 [57]. The BladeRF is a portable, relatively low-cost, and consumer-available SDR. The BladeRF is well-supported and would be straightforward for an attacker to obtain and implement. In this work, all experiments were carried out with a transmitting power of 12.41 dBm (measured at 480 MHz). The BladeRF can be easily programmed using GNU Radio[58], an open-source radio toolkit. In the case of this experiment, GNU Radio is used to produce a simple user interface for varying the carrier frequency and modulating signal. This allowed for quick testing and exploration of possible carriers.

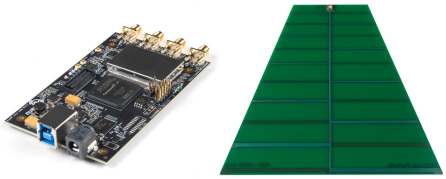

Figure 5: Nuand BladeRF SDR and a log-periodic antenna.

4.1.3 Log-Periodic Antennas. To transmit the RF signals, log-periodic antennas were used. The antennas were implemented in FR4 (Flame-Retardant Glass-Reinforced Epoxy Resin), as a printed circuit board (PCB) with a single SMA connector for signal input. The PCB antennas are designed to work in the frequency range of 400 MHz to 1000 MHz. Log-periodic antennas were ideal for this application due to their directionality. Also, due to the PCB-based implementation, this solution is extremely flexible. The antenna could be redesigned to meet the required frequency range of a given EEG system, if necessary.

4.1.4 Power Amplifier. In order to be able to receive the radio-frequency signals from longer distances, we used a power amplifier placed between the SDR and the antenna. In particular, we used the BT-100 Bias-tee power amplifier from Nuand [59]. This amplifier provided a gain of 8.37 dB (measured at 480 MHZ).

4.2 System Overview

Brain-Hack system consists of a computer running GNU Radio to choose the desired carrier and modulating frequencies. The computer is connected to the BladeRF SDR, followed by the bias tee power amplifier, which is plugged into the log-periodic antenna.

The antenna transmits the radio-frequency signals defined in software by GNU Radio, which are received by the physical structure of the EEG recording device. As explained in Section 3.5, the low-frequency modulating signals are inserted into the recording device, becoming part of the recorded signals. Hence, the received signals contain simultaneously real EEG waves as well as fake injected signals. These are processed by a second computer linked to the EEG device, and the BCI algorithm outputs a specific command based on the captured signals. If *Brain-Hack* is successful, the output command will be correlated to the information contained by the fake brain-waves injected.

4.3 System Calibration

It is clear that the carrier frequency determines the optimal transmission. In section 3.6, the geometric dependence of the carrier is explored. Due to this variation, a working carrier frequency must be chosen empirically or from prior knowledge of the system. In a lab setting, such as if the attacker were to procure the target system ahead of an attack, a working carrier frequency can be determined by sweeping through a reasonable range while monitoring the output of the target system. When the output of the target system shows a local maximum of injection power, a working carrier has been found. To determine the carrier in real-time, the attacker can transmit a jamming frequency and sweep the carrier over a reasonable range while observing the action of the target system. An expected range of the carrier f_c can be determined mathematically and is related to the length L of the electrode wires used in the system: $f_c = c/L$, where c is the speed of light. This will give an approximate number of a working carrier frequency.

5 SSVEP-BASED BCIs

There is a wide range of possible attacks that can be performed by our *Brain-Hack* system, which range from injecting false brain-waves during clinical trials that can lead to severe misdiagnoses to jamming a brain-controlled wheelchair to paralyze and shut down the system. However, for simplicity, we will focus most part of this work on attacking brain-computer interfaces based on detecting SSVEPs to execute commands. SSVEPs stands for Steady-State Visually Evoked Potentials and they are sinusoidal brain-waves generated in the visual cortex when a person attends at a flickering visual stimuli [60–62]. The key aspect of these type of brain signals is that the frequency of the SSVEP is the same as the frequency of the flickering stimuli. Hence, by flashing each stimuli at different frequencies, one can create a BCI that detects what stimuli is receiving the attention of the user. These stimuli can be rendered in a screen, allowing users to select the flickering elements in the screen just using their visual attention. If each visual stimulus in the screen is related to a different command, the user can select the desired command to execute by gazing at the corresponding stimulus.

5.1 Taking Control of a SSVEP-based BCI

An SSVEP-based BCI detects what command the user is willing to execute by processing the EEG signal recorded and identifying the frequency with highest amplitude of those available for the user. Therefore, in order for an attacker to take control of such type of BCIs, they just need to inject a false SSVEP with an amplitude larger than the real brain-wave. This way, the BCI will execute the command desired by the attacker instead of the one desired by the user. SSVEP amplitudes may range from $0.5\mu V$ to around $10\mu V$ [60].

6 ATTACK MODEL

Figure 6 shows an overview of our attack model. We discuss the components of the attack model in more detail below.

Intent of attacker: The primary intention of the attacker is to remotely assume control of a brain-computer interface and cause harm by forcing the system to behave at the attacker's will. This is accomplished by injecting signals into an EEG measurement system via amplitude-modulated radio-frequency signals transmitted from a remote antenna. As a result, the injected signals are erroneously interpreted as brain waves.

Attack Target: The attacker can target any EEG system. In the case of this work, the attacker can target an SSVEP-based BCI system. As described above, an SSVEP-based BCI uses frequency peaks from measured EEG signals as a control input. The attacker can transmit a modulated pure tone at a control frequency to inject a false frequency peak. This false peak is interpreted as a valid SSVEP response, causing the system to react to the attacker rather than the user of the BCI. Further, the attacker can transmit a series of frequencies or a fixed frequency that is not a possible control input to simply block the user from interacting with the BCI.

Capabilities of attacker: The attacker must have a software defined radio, a computer to control the SDR, and an antenna for transmission. To minimize cost, open-source SDR platforms can be used with a single-board computer, such as a Raspberry Pi, and a low-cost antenna to bring the price to roughly $400. To successfully execute the attack, a working carrier frequency must be known. An approach to determining these values is described in section 4.3. To force undesired outputs from a BCI, valid control signals must also be known. These can be determined from prior knowledge, such as experimenting with target hardware, or from real-time sweeping during an attack. A control frequency sweep is feasible because of the limited frequency range of SSVEP.

Access level of attacker: The attacker may gain access to the target in one of two ways. **First**, the attacker can deploy the system in a fixed location, likely indoors. As shown in this work, successful transmission through walls is possible, and the system could be stored in a nearby room or closet if the EEG or BCI use is indoor. If control is desired, the system must have an internet connection for remote access. **Second**, the attacker can carry the system with them for mobile use, likely outdoors. The software defined radio can be powered by a laptop or battery allowing the entire system to fairly discrete. In this scenario, the attacker must be in proximity to the target, but can remain unassuming in public spaces where BCIs may be used for smartphone or wheelchair control. Additionally, having a line-of-sight to the target allows for easier real-time characterization of the system.

Attack application: The attacker will turn on the SDR and enable AM transmission. In the case of a pre-installed, fixed-position deployment of the system, the system will disrupt any EEG reading within range of the device. In the case of a mobile deployment, the

attacker will discretely position themselves such that they are close enough to the target to disrupt their activity.

Attack outcomes: The attacker changes the apparent brain state of the target, causing a physician or a BCI system to view and interpret erroneous EEG patterns. In the case of this work, the attacker forces a BCI-controlled drone crash, a BCI-based speller type undesired words and a neuro-feedback meditation tool output false meditative states. Additionally, this technique could lead to clinical misdiagnoses, malicious control of BCI-based assistive technologies, and the revealing of personal information via BCI-based smartphone interactions. As BCI technology becomes increasingly ubiquitous, the possible applications of this technique increase in number.

7 SYSTEM CHARACTERIZATION

Once having proven that our approach could successfully inject false brain-waves into an EEG monitoring device, we wanted to further characterize the system. Therefore, we conducted several experiments to understand how the system behaves under different conditions. All experiments were carried out with a transmitting power of 12.41 dBm (measured at 480MHz).

7.1 Range

In this experiment, we wanted to measure how the amplitude of the injected waves decays with distance to the transmitting antenna. For this reason, we connected the electrodes of ground, reference and channel to a board that simulates the impedance of a real head at each electrode (about 4 $k\Omega$) [63]. The cables and device were placed in no particular orientation and left on the table (see Figure 7a). The transmitting antenna was oriented with its plane (and polarization) parallel to the ground. For this experiment we used the Neuroelectrics Enobio device.

In Figure 7 we can see the results. Notice that the plotted results follow the characteristic trend of electric field decay with distance at rate $1/d$. We can see that with the specified transmitting power, the maximum range is up to 3 meters approximately.

7.2 Angle of Incidence

This experiment consisted of measuring the amplitude of the injected signals as a function of angle of incidence. We define angle of incidence as the angle formed by the direction of maximum antenna radiation and the direction faced by the subject wearing the EEG headset. This way we could retrieve the impact of our *Brain-Hack* system depending on the direction of incidence. See Figure 8a to visualize a schematic of the experiment setup. The antenna was placed 1 meter away from the head of the user, and the user's chair rotated from 0° to 360°. We took measurements both with the antenna positioned horizontally and with the antenna positioned vertically, see Figures 8b and 8c with the respective results. The channel measured with the Neuroelectrics and OpenBCI devices was the occipital Oz, and the left temporal channel TP9 with the Muse headband. It appears that positioning the antenna horizontally leads to stronger injected waves, and a larger coverage angle. Note the bi-modality of the Neuroelectrics device, which, with the horizontal antenna, can receive the signals both facing (0°) and facing away (210°-240°) the antenna. The shift of the resulting coverage lobes towards the first quarter angles may be due to the fact that the reference electrode of the Neuroelectrics device and of the OpenBCI device, as well as the measuring channel of the Muse headband were placed by the left ear. Since the antenna was facing the left region of the user's head while $\alpha \in (0°, 90°]$ (user rotated clockwise, see Figure 8a), this resulted in stronger signals for those angles, hence the slight shift of the lobes towards the first quarter of the polar plot.

7.3 Probabilities of Attack Success Versus Distance

The purpose of this experiment was to conduct an statistical analysis of the effectiveness of our attack as a function of distance between the antenna and the BCI user. This would allow us to compute the probabilities of a successful attack for each of the different distances. A successful attack is defined as recording an RF-injected wave with higher amplitude than the amplitude of the real SSVEP brain-waves, as explained in Section 5.1. This way, the BCI system would execute the command associated to the frequency of the wave injected by the attacker, instead of the command desired by the user. For this experiment, the user sat in front of a computer screen that rendered a visual stimulus consisting of a flickering white square on a black background. This is one of the most common schemes used for SSVEP-based brain-computer interfaces [64], and it was similar to what we would later use for a real attack demonstration (see Section 8). The stimulating square comprised a visual angle of 5° and flickered at 20 Hz. In order to obtain a distribution of the amplitudes of real SSVEP waves, we recorded a window of 5 seconds while the user attended at the flickering square, and repeated the measurement 10 times. This way, we obtained an average amplitude of $3.8\mu V$ with a standard deviation of $1.07\mu V$. Similarly, we set up our *Brain-Hack* system to remotely inject signals of 20 Hz and recorded 10 5-second windows with the EEG device for different distances. This allowed us to obtain an average and standard deviation of the amplitudes of the injected signals for each distance. We repeated the measurements placing the antenna in the room next door to explore the performance of the attack if an attacker was hiding in another room. You can see the results in Figure 9 (Top). The solid lines represent the average and the colored shadows the standard deviations. We can see that the standard deviation is considerably small for the injected signals, which means that the attack results are very repeatable. Once these statistics have been obtained, we can compute the probability of attack success for each distance. If we assume that the amplitudes of the signals are normally distributed we can easily determine that the probability of a successful attack is given by $Pr(A_i > A_r)$, where A_i and A_r are random variables normally distributed that represent the amplitudes of the injected and real brain-waves. If $A_i \sim \mathcal{N}(\mu_i, \sigma_i^2)$ and $A_r \sim \mathcal{N}(\mu_r, \sigma_r^2)$, we can define $Z = A_i - A_r$ and then, $Pr(A_i > A_r) = Pr(A_i - A_r > 0) = Pr(Z > 0)$, where $Z \sim \mathcal{N}(\mu_i - \mu_r, \sigma_i^2 + \sigma_r^2)$. This allows us to obtain the plot in Figure 9 (Bottom), where we show the probabilities of attack success for each distance. The fact that the standard deviation of the signals is very narrow, results in considerable jumps of probability from 0 to 1 very quickly. Moreover, we can see that at a distance of about 2 meters, the probabilities of attack success drop to zero. We want to

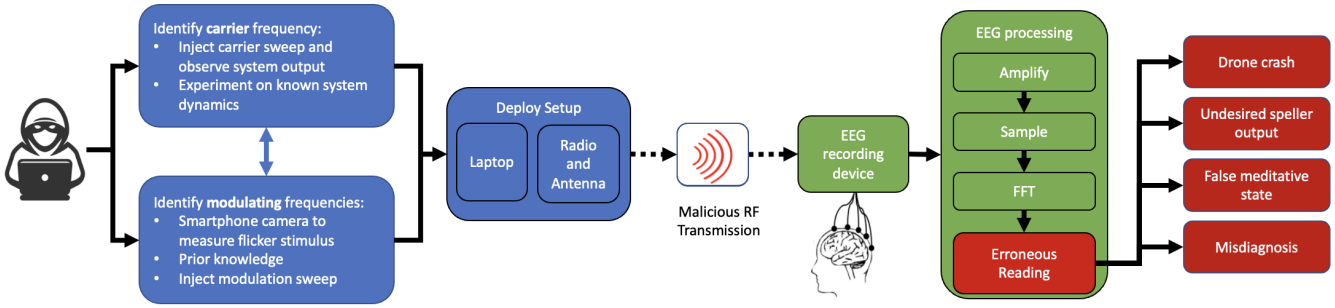

Figure 6: A brief overview of *Brain-Hack*'s attack model. Blue blocks represent elements or actions performed by the attacker, green blocks represent elements or processes performed by the attacked system, and red blocks indicate negative outcomes of the attack.

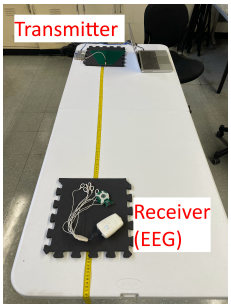

(a) Setup to obtain the measurements voltage vs. distance. EEG connected to "fake head" board.

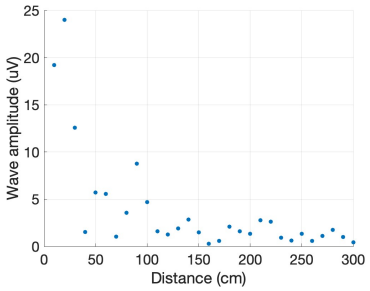

(b) Wave amplitude versus distance plot. Notice the decay rate $1/d$.

Figure 7: Amplitude of wave received versus distance to transmitting antenna. In this case, the EEG electrodes were connected to a "fake head" board that simulates the electrode impedance of a real head.

point out that these results represent the probabilities of success under the conditions that these particular experiments were carried out. The effective range of the attack could be easily increased with more transmitting power (current transmitting power is just of 17.4 mW).

7.4 Walking Axially

In this experiment, we wanted to demonstrate the robustness of our attack to the body and head movement of the EEG device user. For this reason, we set an experiment where the user wearing the EEG device walked away axially from the transmitting antenna. The path followed by the user was a straight line aligned with the direction of maximum power transmission of the antenna. The antenna height was set so it matched the height of the user's head. The user walked up to a distance of 2.5 meters from the antenna. The injected frequency was 20 Hz. The result can be seen in Figure 10, where we show an spectrogram of the recorded signal with the EEG device. We can clearly see that the 20 Hz wave was captured at the EEG, decreasing its power with time and distance. This result shows that even with the head and body movement of the user, our system can still successfully inject the desired signal into the device.

8 ATTACK MODEL DEMONSTRATION

In this section we show how our attack approach can be used to take control of three different brain-computer interfaces: a speller, a drone-control and a neuro-feedback tool for meditation. The speller and drone-control applications functioned using SSVEP, as explained in Section 5.

8.1 Virtual Keyboard Speller

As said above, we designed a speller application that functioned using SSVEPs. The application flashed each letter key at a different frequency, and the user could type just attending the desired key. SSVEP-based spellers are well-known in the BCI community since they allow patients with reduced mobility to communicate [65–68].

The demo consisted of two parts. First, the user typed the words 'I LOVE MIT' without any remote attack. We can see in Figure 11a that the task is completed successfully. In the second part of the demonstration, the user wished to type the same words but, in this case, the attack system was activated. While the user was attending the corresponding letters to type 'LOVE', the Software-Defined Radio was transmitting signals with encoded frequencies to type 'HATE' instead. The radio was programmed so that the power of the injected fake brain-waves was higher than that of the real evoked SSVEPs, as seen in Figure 11b. This way, the classifier algorithm –a simple FFT that chooses the frequency with highest amplitude– would wrongly choose the key desired by the attacker as the attended key. The result can be seen in Figure 11c. A link to a video of the full demonstration can be found in the Appendices. This example shows that our attack approach can make a user type undesired sentences remotely, even hiding the system in a room next-doors as seen in Section 7.3.

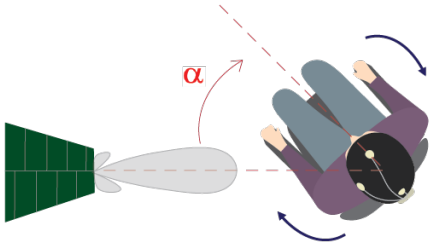

(a) Schematic representation of the experiment, transverse view. When $\alpha = 0^{\circ}$, the subject was facing the direction of maximum radiation of the antenna. The subject rotated clockwise.

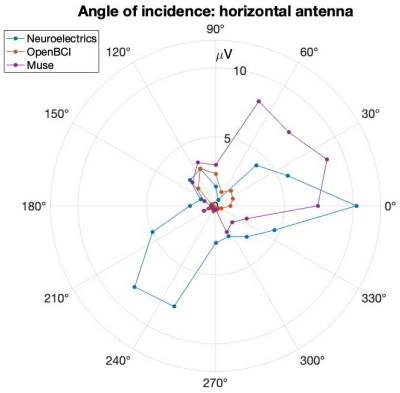

(b) Polar plot of the injected waves amplitude in μV as a function of the angle of incidence α. In this case, the antenna was placed horizontally (parallel to the ground).

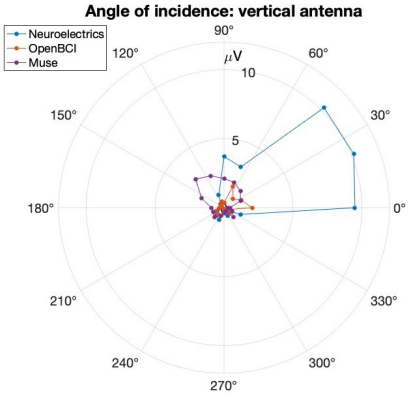

(c) Polar plot of the injected waves amplitude in μV as a function of the angle of incidence α. In this case, the antenna was placed vertically (perpendicular to the ground).

Figure 8: Angle of incidence experiment.

8.2 Drone-Control

This demo consisted of a brain-computer interface to control a drone (or small helicopter) in real time. The interface functioned also with SSVEP, as seen in the speller demo, Section 8.1. The user was seated in front of a computer screen which showed different flashing

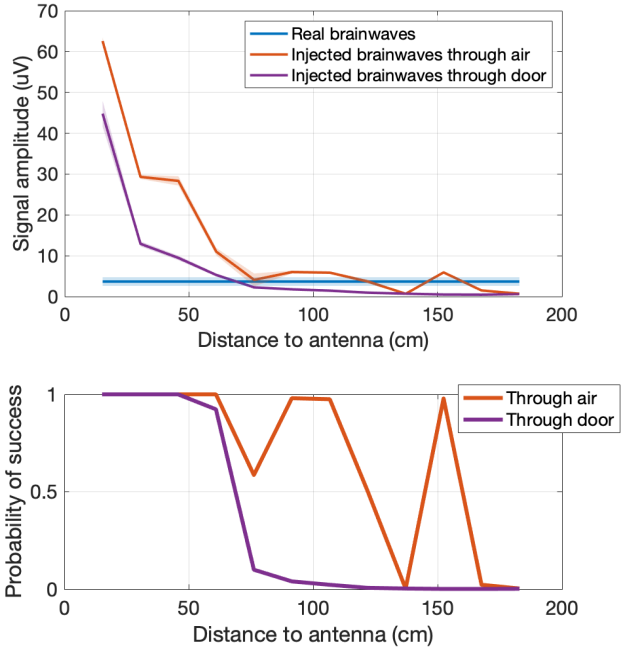

Figure 9: (Top) Signal amplitudes of real (blue) and injected (orange and purple) brain-waves as a function of distance between the antenna and the user. Solid lines represent average values, shadows represent standard deviation. (Bottom) Probabilities of attack success as a function of distance between the antenna and the user. We assume that the amplitudes of the signals are normally distributed.

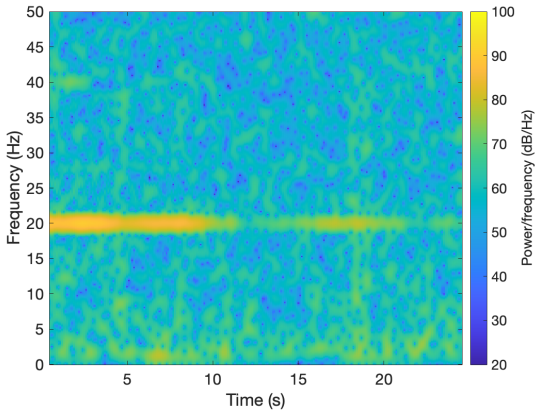

Figure 10: Spectrogram of the EEG signal recorded when the user walked away axially from the transmitting antenna. We can see that the desired 20 Hz wave is successfully injected into the EEG device. The power of the injected signal decreased as the user walked away until reaching 2.5 meters from the antenna.

(a) In the first part of the demonstration, the user typed the desired sentence (I LOVE MIT) without any remote attack. The user was wearing the Neuroelectrics device.

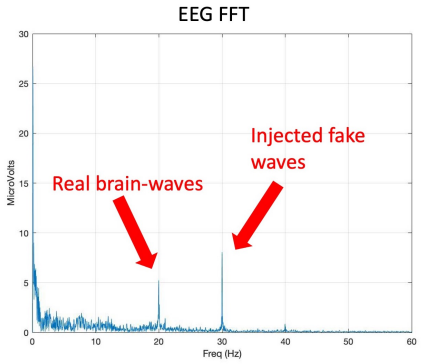

(b) Fast Fourier Transform (FFT) of the EEG recorded during the typing of the letter 'L'. Attending the 'L' key generated a 20 Hz SSVEP. Simultaneously, our SDR was transmitting radio-waves with the frequency of 30 Hz encoded, which corresponded to letter 'H'. As seen, the power of the injected fake wave is larger than that of the real SSVEP. Therefore, the speller classifier wrongly chose the letter 'H' as the attended key.

(c) In the second part of the demo, the remote attack caused the speller to type the word 'HATE' instead of 'LOVE'. See the SDR and the log-periodic antenna at the bottom right corner of the picture.

Figure 11: Demonstration of our attack to an SSVEP-based speller.

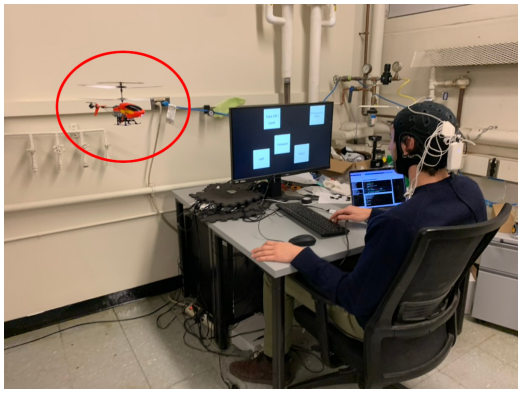

(a) In the first part of the experiment, the user could freely control the drone gazing at the visual stimulus associated to the desired command. The user could complete the task since no hacking took place. The user was wearing the Neuroelectrics device.

(b) In the second part of the experiment, the SDR hacking was turned on (top right), which activated the Emergency stop command. Consequently, the drone crashed to the ground (bottom left).

Figure 12: Demonstration of our attack to an SSVEP-based drone-control interface.

visual stimuli. Each stimuli was linked to a particular command: Take off/Land, Go forward, Turn left, Turn right and Emergency stop, as seen in Figure 12a.

The task given to the user consisted of maneuvering the drone from the ground and landing it on a red surface at approximate 1.5 meters high. In the first part of the experiment, the user completed the task successfully since there was no hacking attack, as seen in Figure 12a. Conversely, in the second part of the experiment the SDR was programmed to send radio-waves modulated at the frequency linked to the Emergency stop command. Hence, during the normal flight of the drone, the command was activated and the drone crashed to the floor, as seen in Figure 12b. A video of the complete demo can be seen following the link available in the Appendices. This experiment shows that the use of our system could take complete control of a brain-controlled drone. Instead of crashing the drone, this technique could cause much more harm by directing the drone towards people around it.

8.3 Meditation

This demonstration consisted of hacking the consumer-grade EEG device, Muse. Muse comes with its mobile App that provides feedback during a meditation session in the form of sounds. At the end of each session, a report shows how well the subject performed. A graph shows the state of mind of the user during the session, classifying it in three different mental states: Active, Neutral, and Calm. The goal of this demonstration was to use our system to inject into the device fake brain-waves that are known to appear in deep meditation states. In particular, these are slow oscillatory waves found in the lower band of the brain spectrum: delta (< 4 Hz) and theta (4 - 8 Hz) waves [69, 70]. We discovered that by remotely injecting fake delta waves (3 Hz) into the frontal left electrode (AF7) and simultaneously injecting theta waves (8 Hz) into the left temporal electrode (TP9), the Muse algorithm wrongly classified the state of mind into "Calm", even if the subject remained with the eyes open and no intention to meditate (see Figure 13).

The experiment was conducted as follows. First, the subject (who has no experience with meditation practice whatsoever) completed the calibration requested by the Muse App with the eyes closed and trying to relax, as indicated. When the 5-minute session started, the subject was instructed to remain still and with the eyes open during the whole time. During the first half of the session (first 2.5 minutes), the *Brain-Hack* system was off. We can see in the Muse App report that during the first half of the session the state of mind was classified as "Active" (see Figure 13b). Then, after approximately the first 2.5 minutes, the hacking system was activated. We can see in Figure 13b that the Muse App reported a sudden change in state of mind, jumping from "Active" to "Calm". The App also played many bird sounds (shown as well in the report as blue icons), which appear when the subject is very calm for a long time.

This demonstration showed that *Brain-Hack* can make a Muse device report that a "bad meditator" reaches an outstanding meditation state. This experiment demonstrates the versatility of our system, which can take control of a consumer-grade EEG device without any cables exposed, as opposed to the other two devices. Moreover, this attack successfully proved that two different frequencies can be injected into the device simultaneously, opening the door to generating more complex brain signals, going beyond simple sinusoidal signals.

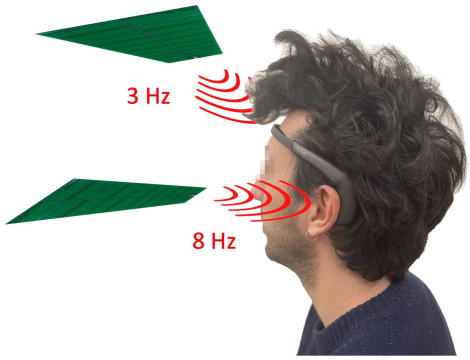

(a) To successfully trick the Muse algorithm into believing that the subject was in a state of deep meditation, we had to simultaneously direct 3 Hz to the frontal electrode of the device (AF7) and 8 Hz to the temporal one (TP9).

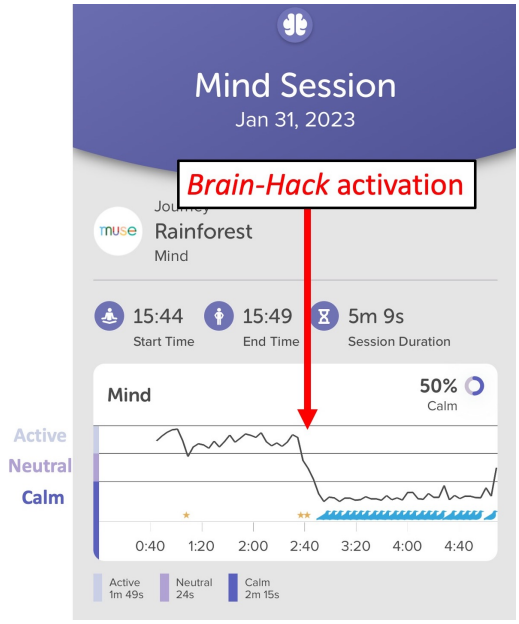

(b) Screenshot of the Muse App report after a 5-minute session. In the first 2.5 minutes, the hacking system was not active, and the Muse algorithm reports an 'Active' mind (the subject kept the eyes open the whole session). In the second half of the session, *Brain-Hack* was activated. As seen, there is a sudden drop of mind state from 'Active' to 'Calm'. The birds depicted appear when the subject is very calm for a long time.

Figure 13: Hacking a Muse EEG headband.

9 ETHICS STATEMENT

This work received IRB approval from our institution (MIT COUHES), and the participant in the experiment gave informed consent. The primary concern for this experiment was the electromagnetic radiation transmitted towards the participant's head. The FCC determines that the only effects that may constitute a human health hazard from exposure to RF energy are those resulting from tissue heating, referred to as "thermal" effects [71]. According to the FCC, the evidence for production of harmful biological impact aside from thermal effects is ambiguous and unproven. Therefore, the FCC specifies that the maximum exposure level allowed to avoid harmful thermal effects is of $580\mu W/cm^2$ [72]. We can compute the power density values to which the participant was exposed with the following formula: $P_D = P_t G_t / 4\pi R^2$, where P_t, G_t and R are the transmitter power, the antenna gain and the distance to the antenna respectively. The measured transmitting power of our radio was 12.41dBm or 17.4mW, and the gain of our antenna was 6dBi. At a distance of 15 cm (the closest we placed the antenna to the participant), the resulting power density is of $36.9\mu W/cm^2$, which is still one order of magnitude lower than the FCC threshold. Therefore, the electromagnetic radiation transmitted by our SDR was always below any harmful level for the participant.

10 DISCUSSION

It is concerning that these BCI devices are easily compromised in this manner. The goal of this paper is to contribute to the conversation in BCI security, and specifically to highlight this security flaw in EEG-based BCI systems. With such systems becoming increasingly common, it will be vital to employ some form of countermeasure for this and other possible attacks.

10.1 Limitations

This work explores the use of amplitude-modulation to inject pure sinusoids, the perfect way to spoof SSVEP-based systems. However, this technique will not allow the attacker to control systems devoted to more complex EEG processes such as motor imagery [73, 74] or Event-Related Potentials (ERP) [75, 76].

Additionally, the determination of the optimal control and carrier frequencies may be challenging or expensive, depending on the scenario. The attacker would be able to determine the control frequencies using a smartphone with high-speed video capabilities, and would need to be close enough to the blinking source to capture the signal. In cases where the frequencies are fixed, the attacker may be able to download relevant applications or purchase relevant devices for analysis prior to the attack. While this prior information aids the process, it is enabled by additional expenses.

Despite these limitations, we consider this technique to be a viable path to disrupt BCI systems and these security holes should be closed for the safety of the BCI users.

10.2 Countermeasures

There are multiple possible countermeasures that can be employed to prevent this effect with minimal additional cost to the manufacturer of EEG systems and BCI devices.

Control Frequency Variation. A simple solution for SSVEP-based applications would be to frequently change the frequencies of the SSVEP control signals. In this case, the attacker would be able to disrupt control, but they would not easily be able to remotely control the system to their whim.

Response Time Personalization. Additionally, an SSVEP-based BCI could take advantage of the individual-specific phase or delay response to blinking lights [77, 78]. With phase recognition, the system could easily identify which signals are from the user and which are falsely injected by the attacker. In other words, the phase of the injected signal would not match that of the BCI user, and the injected signal would be ignored.

Common-Mode Rejection. The antennas used in this work are directional, but not precise enough to target specific regions of the head. As such, all channels see the same injected signal. One possible way to remove the injected signal would be to remove any 'noise' that is common to all channels, rather than only removing noise common to each channel relative to a reference.

Shielding. Further, a simple and effective countermeasure for this attack would be to use shielded cables and to shield the sensitive electronics within the EEG system. This would improve signal resistance to RF interference.

Active Electrodes. Another possible solution is the use of active electrodes [79]. In this configuration, amplifiers are placed proximal to the electrodes, amplifying the signal closer to the source. After amplification, the captured signal will have a significantly higher amplitude and a stronger injection would be required to overpower the true signal.

10.3 Future Work

This work appears to open many possibilities for this setup. For example, the Muse attack demonstrates the use of multiple frequencies in the injection signal. This suggests that the system could successfully deploy more complex, false signals into EEG devices. With the addition of beam-forming antennas, complex injections could be applied at various locations around the head, allowing an attacker to spoof more intricate brain processes. Recordings of known neurological signals would also allow the carrier to be amplitude-modulated with arbitrary waveforms, further disrupting EEG systems.

Additionally, it seems plausible that *Brain-Hack* is capable of disrupting intracranial EEG systems. In the case of an implanted system, the electrode wires will be shorter, the true signal would be stronger, and would be shielded by the skull. The skull can be penetrated by some electromagnetic energy in the microwave range, but this my cause other neurological issues[80]. Compromising implanted devices would be deeply concerning and would call into question the safety of such devices.

Finally, it was observed that the orientation and position of the BCI can cause shifts in optimal carrier frequency. In a future work, we would like to explore solutions to this problem. One possible implementation would be to use a sawtooth carrier signal, creating amplitude-modulated frequency components at integer multiples of the carrier frequency. Additionally, the carrier could be frequency modulated, allowing for a sweep of the AM signal across a desired frequency range. These solutions would lead to having a higher likelihood of hitting a working carrier frequency without the need of calibration.

11 CONCLUSIONS

In this work, we exploit the physical structure of EEG devices to inject false brain-waves and impose control signals, rendering the BCI-user under the command of an attacker. This attack was successfully applied to a variety of EEG devices, including a research-grade device, the Neuroelectrics Enobio [16], a consumer-grade device, the InteraXon Muse 2 [18], and an open-source development device, the OpenBCI Ganglion[17].

We find these vulnerabilities concerning, especially considering the growth of EEG-based consumer products and implants. The countermeasures suggested here can be implemented relatively simply, some offering additional benefits to signal quality and noise reduction. We believe that developers of BCI systems have a responsibility to prioritize the safety of the end user. Hopefully, this paper contributes to the conversation of BCI security and improves the robustness of EEG systems.

ACKNOWLEDGMENTS

The authors would like to thank Professor Fadel Adib for his guidance during the initial phases of this work and for providing us with the SDR and antennas needed for this project.

REFERENCES

[1] Andres F Salazar-Gomez, Joseph DelPreto, Stephanie Gil, Frank H Guenther, and Daniela Rus. Correcting robot mistakes in real time using eeg signals. In *2017 IEEE international conference on robotics and automation (ICRA)*, pages 6570–6577. IEEE, 2017.

[2] Bojan Kerous, Filip Skola, and Fotis Liarokapis. Eeg-based bci and video games: a progress report. *Virtual Reality*, 22(2):119–135, 2018.

[3] Alexandre Armengol-Urpi, Andrés F Salazar-Gómez, and Sanjay E Sarma. Brainwave-augmented eye tracker: High-frequency ssveps improves camera-based eye tracking accuracy. In *27th International Conference on Intelligent User Interfaces*, pages 258–276, 2022.

[4] Alexandre Armengol-Urpi and Sanjay E Sarma. Sublime: a hands-free virtual reality menu navigation system using a high-frequency ssvep-based brain-computer interface. In *Proceedings of the 24th ACM Symposium on Virtual Reality Software and Technology*, pages 1–8, 2018.

[5] Hakim Si-Mohammed, Jimmy Petit, Camille Jeunet, Ferran Argelaguet, Fabien Spindler, Andéol Evain, Nicolas Roussel, Géry Casiez, and Anatole Lécuyer. Towards bci-based interfaces for augmented reality: feasibility, design and evaluation. *IEEE transactions on visualization and computer graphics*, 26(3):1608–1621, 2018.

[6] Muse eeg headband. https://choosemuse.com/. Accessed: December 2022.

[7] Neurable. https://neurable.com/headphones. Accessed: December 2022.

[8] Brainco. https://brainco.tech/. Accessed: December 2022.

[9] Next mind. https://www.next-mind.com/. Accessed: September 2022.

[10] Nissan brain-to-vehicle. https://usa.nissannews.com/en-US/releases/nissan-brain-to-vehicle-technology-redefines-future-of-driving. Accessed: December 2022.

[11] Yuanqing Li, Jiahui Pan, Fei Wang, and Zhuliang Yu. A hybrid bci system combining p300 and ssvep and its application to wheelchair control. *IEEE Transactions on Biomedical Engineering*, 60(11):3156–3166, 2013.

[12] Amin Nourmohammadi, Mohammad Jafari, and Thorsten O Zander. A survey on unmanned aerial vehicle remote control using brain–computer interface. *IEEE Transactions on Human-Machine Systems*, 48(4):337–348, 2018.

[13] Dennis J McFarland and Jonathan R Wolpaw. Brain-computer interface operation of robotic and prosthetic devices. *Computer*, 41(10):52–56, 2008.

[14] Xiaogang Chen, Yijun Wang, Masaki Nakanishi, Xiaorong Gao, Tzyy-Ping Jung, and Shangkai Gao. High-speed spelling with a noninvasive brain–computer interface. *Proceedings of the national academy of sciences*, 112(44):E6058–E6067, 2015.

[15] Violaine Guy, Marie-Helene Soriani, Mariane Bruno, Theodore Papadopoulo, Claude Desnuelle, and Maureen Clerc. Brain computer interface with the p300 speller: usability for disabled people with amyotrophic lateral sclerosis. *Annals of physical and rehabilitation medicine*, 61(1):5–11, 2018.

[16] Neuroelectrics' enobio eeg. https://www.neuroelectrics.com/solutions/enobio. Accessed: December 2022.

[17] Openbci. https://openbci.com/. Accessed: December 2022.

[18] Muse 2. https://choosemuse.com/muse-2/. Accessed: December 2022.

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

[20] Mario Frank, Tiffany Hwu, Sakshi Jain, Robert T Knight, Ivan Martinovic, Prateek Mittal, Daniele Perito, Ivo Sluganovic, and Dawn Song. Using eeg-based bci devices to subliminally probe for private information. In *Proceedings of the 2017 on Workshop on Privacy in the Electronic Society*, pages 133–136, 2017.

[21] Michael Inzlicht, Ian McGregor, Jacob B Hirsh, and Kyle Nash. Neural markers of religious conviction. *Psychological science*, 20(3):385–392, 2009.

[22] Rogeria Waismann, Peter BC Fenwick, Glenn D Wilson, Terry D Hewett, and John Lumsden. Eeg responses to visual erotic stimuli in men with normal and paraphilic interests. *Archives of sexual behavior*, 32(2):135–144, 2003.

[23] P Flor-Henry, RA Lang, ZJ Koles, and RR Frenzel. Quantitative eeg studies of pedophilia. *International Journal of Psychophysiology*, 10(3):253–258, 1991.

[24] Xiao Zhang, Dongrui Wu, Lieyun Ding, Hanbin Luo, Chin-Teng Lin, Tzyy-Ping Jung, and Ricardo Chavarriaga. Tiny noise, big mistakes: adversarial perturbations induce errors in brain–computer interface spellers. *National science review*, 8(4):nwaa233, 2021.

[25] Rui Bian, Lubin Meng, and Dongrui Wu. Ssvep-based brain-computer interfaces are vulnerable to square wave attacks. *Science China Information Sciences*, 65(4):1–13, 2022.

[26] Lubin Meng, Jian Huang, Zhigang Zeng, Xue Jiang, Shan Yu, Tzyy-Ping Jung, Chin-Teng Lin, Ricardo Chavarriaga, and Dongrui Wu. Eeg-based brain-computer interfaces are vulnerable to backdoor attacks. *arXiv preprint arXiv:2011.00101*, 2020.

[27] Xiao Zhang and Dongrui Wu. On the vulnerability of cnn classifiers in eeg-based bcis. *IEEE transactions on neural systems and rehabilitation engineering*, 27(5):814–825, 2019.

[28] Zihan Liu, Lubin Meng, Xiao Zhang, Weili Fang, and Dongrui Wu. Universal adversarial perturbations for cnn classifiers in eeg-based bcis. *Journal of Neural Engineering*, 18(4):0460a4, 2021.

[29] QianQian Li, Ding Ding, and Mauro Conti. Brain-computer interface applications: Security and privacy challenges. In *2015 IEEE conference on communications and network security (CNS)*, pages 663–666. IEEE, 2015.

[30] Marcello Ienca and Pim Haselager. Hacking the brain: brain–computer interfacing technology and the ethics of neurosecurity. *Ethics and Information Technology*, 18(2):117–129, 2016.

[31] Ofir Landau, Rami Puzis, and Nir Nissim. Mind your mind: Eeg-based brain-computer interfaces and their security in cyber space. *ACM Computing Surveys (CSUR)*, 53(1):1–38, 2020.

[32] Marcello Ienca. Neuroprivacy, neurosecurity and brain-hacking: Emerging issues in neural engineering. In *Bioethica Forum*, volume 8, pages 51–53. Schwabe, 2015.

[33] Tamara Denning, Yoky Matsuoka, and Tadayoshi Kohno. Neurosecurity: security and privacy for neural devices. *Neurosurgical Focus*, 27(1):E7, 2009.

[34] Sergio López Bernal, Alberto Huertas Celdrán, Gregorio Martínez Pérez, Michael Taynnan Barros, and Sasitharan Balasubramaniam. Security in brain-computer interfaces: state-of-the-art, opportunities, and future challenges. *ACM Computing Surveys (CSUR)*, 54(1):1–35, 2021.

[35] Tamara Bonaci, Ryan Calo, and Howard Jay Chizeck. App stores for the brain: Privacy & security in brain-computer interfaces. In *2014 IEEE International Symposium on Ethics in Science, Technology and Engineering*, pages 1–7. IEEE, 2014.

[36] Hassan Takabi, Anuj Bhalotiya, and Manar Alohaly. Brain computer interface (bci) applications: Privacy threats and countermeasures. In *2016 IEEE 2nd International Conference on Collaboration and Internet Computing (CIC)*, pages 102–111. IEEE, 2016.

[37] Ali Bulent Usakli et al. Improvement of eeg signal acquisition: An electrical aspect for state of the art of front end. *Computational intelligence and neuroscience*, 2010, 2010.

[38] Chen Yan, Hocheol Shin, Connor Bolton, Wenyuan Xu, Yongdae Kim, and Kevin Fu. SoK: A Minimalist Approach to Formalizing Analog Sensor Security. In *2020 IEEE Symposium on Security and Privacy (SP)*, pages 233–248, San Francisco, CA, USA, May 2020. IEEE.

[39] Jayaprakash Selvaraj, Gökçen Yılmaz Dayanıklı, Neelam Prabhu Gaunkar, David Ware, Ryan M Gerdes, and Mani Mina. Electromagnetic induction attacks against embedded systems. In *Proceedings of the 2018 on Asia Conference on Computer and Communications Security*, pages 499–510, 2018.

[40] Kai Wang, Richard Mitev, Chen Yan, Xiaoyu Ji, Ahmad-Reza Sadeghi, and Wenyuan Xu. {GhostTouch}: Targeted attacks on touchscreens without physical touch. In *31st USENIX Security Symposium (USENIX Security 22)*, pages 1543–1559, 2022.

[41] Chaouki Kasmi and Jose Lopes Esteves. Iemi threats for information security: Remote command injection on modern smartphones. *IEEE Transactions on Electromagnetic Compatibility*, 57(6):1752–1755, 2015.

[42] Sebastian Köhler, Richard Baker, and Ivan Martinovic. Signal injection attacks against ccd image sensors. In *Proceedings of the 2022 ACM on Asia Conference on Computer and Communications Security*, pages 294–308, 2022.

[43] Qinhong Jiang, Xiaoyu Ji, Chen Yan, Zhixin Xie, Haina Lou, and Wenyuan Xu. Glitchhiker: Uncovering vulnerabilities of image signal transmission with iemi. In *USENIX Security*, volume 23, 2023.

[44] Gökçen Yılmaz Dayanıklı, Abdullah Zubair Mohammed, Ryan Gerdes, and Mani Mina. Wireless manipulation of serial communication. In *Proceedings of the 2022 ACM on Asia Conference on Computer and Communications Security*, pages 222–236, 2022.

[45] Yasser Shoukry, Paul Martin, Paulo Tabuada, and Mani Srivastava. Non-invasive spoofing attacks for anti-lock braking systems. In Guido Bertoni and Jean-Sébastien Coron, editors, *Cryptographic Hardware and Embedded Systems - CHES 2013*, pages 55–72, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg.

[46] Joon-Ha Jang, Mangi Cho, Jaehoon Kim, Dongkwan Kim, and Yongdae Kim. Paralyzing drones via emi signal injection on sensory communication channels. In *NDSS*, 2023.

[47] Yazhou Tu, Sara Rampazzi, Bin Hao, Angel Rodriguez, Kevin Fu, and Xiali Hei. Trick or heat? manipulating critical temperature-based control systems using rectification attacks. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, pages 2301–2315, 2019.

[48] D. F. Kune, J. Backes, S. S. Clark, D. Kramer, M. Reynolds, K. Fu, Yongdae Kim, and Wenyuan Xu. Ghost Talk: Mitigating EMI Signal Injection Attacks against Analog Sensors. In *2013 IEEE Symposium on Security and Privacy*, pages 145–159, Berkeley, CA, May 2013. IEEE.

[49] Gökçen Yılmaz Dayanıklı, Sourav Sinha, Devaprakash Muniraj, Ryan M Gerdes, Mazen Farhood, and Mani Mina. {Physical-Layer} attacks against pulse width {Modulation-Controlled} actuators. In *31st USENIX Security Symposium (USENIX Security 22)*, pages 953–970, 2022.

[50] J Ph Lachaux, D Rudrauf, and P Kahane. Intracranial eeg and human brain mapping. *Journal of Physiology-Paris*, 97(4-6):613–628, 2003.

[51] Karim Jerbi, Samson Freyermuth, Lorella Minotti, Philippe Kahane, Alain Berthoz, and Jean-Philippe Lachaux. Watching brain tv and playing brain ball: Exploring

novel bci strategies using real-time analysis of human intracranial data. *International review of neurobiology*, 86:159–168, 2009.
[52] Dingkun Liu, Xin Xu, Dongyang Li, Jie Li, Xinguang Yu, Zhipei Ling, and Bo Hong. Intracranial brain-computer interface spelling using localized visual motion response. *Neuroimage*, 258:119363, 2022.
[53] Constantine A Balanis. *Antenna theory: analysis and design*. John wiley & sons, 2015.
[54] Saeid Sanei and Jonathon A Chambers. *EEG signal processing*. John Wiley & Sons, 2013.
[55] S. Wang, F.C. Lee, D.Y. Chen, and W.G. Odendaal. Effects of Parasitic Parameters on EMI Filter Performance. *IEEE Transactions on Power Electronics*, 19(3):869–877, May 2004.
[56] V. R. R. Samson, B. Praveen Kitti, S. Pradeep Kumar, D. Suresh Babu, and Ch. Monica. Electroencephalogram-Based OpenBCI Devices for Disabled People. In Suresh Chandra Satapathy, Vikrant Bhateja, P. Satish Rama Chowdary, V.V.S.S. Sameer Chakravarthy, and Jaume Anguera, editors, *Proceedings of 2nd International Conference on Micro-Electronics, Electromagnetics and Telecommunications*, volume 434, pages 229–238. Springer Singapore, Singapore, 2018. Series Title: Lecture Notes in Electrical Engineering.
[57] Nuand bladerf. https://www.nuand.com/bladerf-1/. Accessed: December 2022.
[58] Gnu radio. https://www.gnuradio.org/. Accessed: December 2022.
[59] Bias-tee power amplifier. https://www.nuand.com/product/bt-100/. Accessed: December 2022.
[60] Christoph S Herrmann. Human eeg responses to 1–100 hz flicker: resonance phenomena in visual cortex and their potential correlation to cognitive phenomena. *Experimental brain research*, 137:346–353, 2001.
[61] ST Morgan, JC Hansen, and SA Hillyard. Selective attention to stimulus location modulates the steady-state visual evoked potential. *Proceedings of the National Academy of Sciences*, 93(10):4770–4774, 1996.
[62] Steven A Hillyard, Hermann Hinrichs, Claus Tempelmann, Stephen T Morgan, Jonathan C Hansen, Henning Scheich, and Hans-Jochen Heinze. Combining steady-state visual evoked potentials and f mri to localize brain activity during selective attention. *Human brain mapping*, 5(4):287–292, 1997.
[63] Neuroelectrics' fake head test board. https://www.neuroelectrics.com/solution/spareparts-consumables/testboard. Accessed: December 2022.
[64] Danhua Zhu, Jordi Bieger, Gary Garcia Molina, and Ronald M Aarts. A survey of stimulation methods used in ssvep-based bcis. *Computational intelligence and neuroscience*, 2010:1–12, 2010.
[65] Hubert Cecotti. Spelling with non-invasive brain–computer interfaces–current and future trends. *Journal of Physiology-Paris*, 105(1-3):106–114, 2011.
[66] Aya Rezeika, Mihaly Benda, Piotr Stawicki, Felix Gembler, Abdul Saboor, and Ivan Volosyak. Brain–computer interface spellers: A review. *Brain sciences*, 8(4):57, 2018.
[67] Adrien Combaz, Camille Chatelle, Arne Robben, Gertie Vanhoof, Ann Goeleven, Vincent Thijs, Marc M Van Hulle, and Steven Laureys. A comparison of two spelling brain-computer interfaces based on visual p3 and ssvep in locked-in syndrome. *PloS one*, 8(9):e73691, 2013.
[68] Ivan Volosyak, Hubert Cecotti, Diana Valbuena, and Axel Graser. Evaluation of the bremen ssvep based bci in real world conditions. In *2009 IEEE International Conference on Rehabilitation Robotics*, pages 322–331. IEEE, 2009.
[69] Jean-Paul Banquet. Spectral analysis of the eeg in meditation. *Electroencephalography and clinical neurophysiology*, 35(2):143–151, 1973.
[70] Decho Surangsrirat and Apichart Intarapanich. Analysis of the meditation brainwave from consumer eeg device. In *SoutheastCon 2015*, pages 1–6. IEEE, 2015.
[71] Rf safety by the fcc. https://www.fcc.gov/engineering-technology/electromagnetic-compatibility-division/radio-frequency-safety/faq/rf-safety#Q6. Accessed: September 2023.
[72] Fcc maximum exposure level. https://www.fcc.gov/consumers/guides/human-exposure-radio-frequency-fields-guidelines-cellular-and-pcs-sites#:~:text=In%20the%20case%20of%20cellular%20and%20PCS,levels%20typically%20found%20near%20the%20base%20of. Accessed: September 2023.
[73] Swati Aggarwal and Nupur Chugh. Signal processing techniques for motor imagery brain computer interface: A review. *Array*, 1:100003, 2019.
[74] Cheolsoo Park, David Looney, Naveed ur Rehman, Alireza Ahrabian, and Danilo P Mandic. Classification of motor imagery bci using multivariate empirical mode decomposition. *IEEE Transactions on neural systems and rehabilitation engineering*, 21(1):10–22, 2012.
[75] Reza Fazel-Rezai, Brendan Z Allison, Christoph Guger, Eric W Sellers, Sonja C Kleih, and Andrea Kübler. P300 brain computer interface: current challenges and emerging trends. *Frontiers in neuroengineering*, page 14, 2012.
[76] Jing Jin, Brendan Z Allison, Yu Zhang, Xingyu Wang, and Andrzej Cichocki. An erp-based bci using an oddball paradigm with different faces and reduced errors in critical functions. *International journal of neural systems*, 24(08):1450027, 2014.
[77] Francesco Di Russo and Donatella Spinelli. Electrophysiological evidence for an early attentional mechanism in visual processing in humans. *Vision research*, 39(18):2975–2985, 1999.
[78] Benedetto Falsini and Vittorio Porciatti. The temporal frequency response function of pattern erg and vep: changes in optic neuritis. *Electroencephalography and Clinical Neurophysiology/Evoked Potentials Section*, 100(5):428–435, 1996.
[79] Jiawei Xu, Srinjoy Mitra, Chris Van Hoof, Refet Firat Yazicioglu, and Kofi A. A. Makinwa. Active Electrodes for Wearable EEG Acquisition: Review and Electronics Design Methodology. *IEEE Reviews in Biomedical Engineering*, 10:187–198, 2017.
[80] Yutaka Igarashi, Yoko Matsuda, Akira Fuse, Toshiyuki Ishiwata, Zenya Naito, and Hiroyuki Yokota. Pathophysiology of microwave-induced traumatic brain injury. *Biomedical Reports*, 3(4):468–472, July 2015.

APPENDIX - DEMONSTRATION VIDEOS

A video of the two project demos (speller and drone-control) can be seen here: https://youtu.be/IKgR03NYU9E