

Subliminal Probing for Private Information via EEG-Based BCI Devices

Mario Frank[#], Tiffany Hwu[#], Sakshi Jain[#], Robert Knight[#], Ivan Martinovic^b,
Prateek Mittal[†], Daniele Perito[#], Dawn Song[#]

[#]UC Berkeley, ^bUniversity of Oxford, [†]Princeton University

December 23, 2013

Abstract

Brain-Computer-Interfaces (BCIs) are becoming popular low-cost consumer devices for use in networked applications such as gaming or in learning programs with neuro-feedback loops. Martinovic et al. [26] proposed a BCI-based attack in which an adversary is able to infer private information about a user, such as their bank or area-of-living, by analyzing the user's brain activities. However, a key limitation of the above attack is that it is intrusive, requiring user cooperation, and is thus easily detectable and can be reported to other users.

In this paper, we identify and analyze a more serious threat for users of BCI devices. We propose a *subliminal attack* in which the victim is attacked at the levels below his cognitive perception. Our attack involves exposing the victim to visual stimuli for a duration not exceeding 13.3 milliseconds – a duration usually not sufficient for conscious perception. The attacker analyzes subliminal brain activity in response to these short visual stimuli to infer private information about the user. If carried out carefully, for example by hiding the visual stimuli within screen content that the user expects to see, the attack may remain undetected. As a consequence, the attacker can scale it to many victims and expose them to the attack for a long time.

We experimentally demonstrate the feasibility of our subliminal attack via a proof-of-concept study carried out with 27 subjects. We conducted experiments on users wearing Electroencephalography-based BCI devices, and used portrait pictures of people as visual stimuli which were embedded within the background of an innocuous video for a time duration not exceeding 13.3 milliseconds. Our experimental results show that it is feasible for an attacker to learn relevant private information about the user, such as whether the user knows the identity of the person for which the attacker is probing.

1 Introduction

Brain-Computer Interfaces (BCIs) are becoming increasingly popular devices for use in networked applications such as entertainment and gaming, or in learning and cognitive enhancement, such as attention and relaxation training applications, as well as a medium for hands-free writing [10, 28, 33]. The Emotiv device [13] and the Neurosky device [29] are examples of low-cost commodity BCIs, and are intended for home usage with applications written by third-party developers and available for download from application markets (see, e.g., [14, 30]). A popular technology used in BCI for recording brain activities is Electroencephalography (EEG), which uses external scalp electrodes to capture fluctuations of the electrical potentials in the brain. The EEG signal is then processed by the application which extracts salient brainwave features and translates them into certain computer instructions or generates a feedback.

Martinovic et al. [26] recently noted that BCI devices may offer the raw EEG signal to the said potentially untrusted third-party applications. If such an application is malicious, it could abuse the BCI device to infer private information about a user, such as her/his preferred bank or area-of-living. The general idea of this attack is similar to a polygraph where physiological reactions of an interrogated person are used to reason about his/her knowledge. However, a fundamental limitation of the attacks proposed by Martinovic et al. is that they are very intrusive, requiring a cooperative user, and are thus easily detected. Even more problematic for a wide deployment of this attack, after a

few users realize such an abnormal behavior of their new downloaded application, they would report it and flag it as not functional or even malicious. It is a fair assumption that all users of the application are well connected via the application market and the market enables sharing such warnings or flags. This would prevent the attacker from carrying out a large scale attack.

Based on these observations, we raise the question: is it possible to infer private information about users wearing BCI devices in a fully concealed way? In this paper, we propose a *subliminal* attack that infers private information about a victim by attacking the victim at a level below his/her cognitive perception. Similar to *subliminal advertising* (see, e.g., [22]), our key idea is to hide the visual stimuli within the screen content that the user expects to see, for a duration of time that does not exceed 13.3 milliseconds. As an example, for a video-based application, we propose to implement small snippets of visual stimuli within a few frames of the video. The attacker succeeds if it can infer private information about the user without arousing the users' suspicion. This is a challenging task. If the stimuli are shown too prominently, then this increases the chance of the attack being detected. If, in contrast, the attacker does too good a job of hiding the stimuli, then the user's subliminal detection may not be strongly affected, reducing the probability of inferring relevant information about the user. Thus, the attacker must operate within this narrow regime of the user's input channel.

We experimentally demonstrate the feasibility of our subliminal attack via a proof-of-concept study with 27 subjects. We conducted experiments on users wearing EEG-based BCI devices. Our study implements an attack scenario where a user watches a video and the attacker tries to infer whether the user knows a particular person by hiding pictures of this person in the video for a time duration shorter than 13.3 milliseconds. To analyze users' subliminal reactions to the embedded visual stimuli, we use machine learning techniques on the recorded EEG signal. For 18 out of the 27 subjects the attacker was able to guess the secret correctly (8 out of 9 in a variant of the attack). Thereby, the success chance did not vary significantly between subjects who were able to detect the person hidden in the video and subjects who did not notice anything abnormal. These experimental results show that the subliminal attack is feasible – attackers can make probabilistic inferences on the users' knowledge of the person depicted in the visual stimuli in a manner that is concealed from the user.

Our main **contributions**:

- We propose a new attack against users wearing BCI devices. Our attack is subliminal, and infers users' private information by exploiting brain activity in response to visual stimuli that are not cognitively perceived by users. If carried out carefully, for example by embedding the visual stimuli for a very short duration within screen content that the user expects to see, the attack can remain undetected.
- We experimentally demonstrate the feasibility of our subliminal attack via a user study of 27 subjects. Our experimental results show the feasibility of subliminally learning private information about a user, such as whether the user knows the identity of the person depicted in the visual stimuli.

The remainder of the paper is organized as follows: we provide background information about BCI devices and neuro-physiological terms in Section 2. Next, we present our threat model in Section 3. We discuss our experimental attack setup and our analysis methodology in Section 4 and Section 5, respectively. We present our experimental results in Section 6, and discuss broader implications in Section 7. Finally, we discuss related work in Section 8 and conclude in Section 9.

2 Neuroscientific Background on EEG-based BCI

Electroencephalography (EEG) monitors electrical activity at the scalp that corresponds to changes in ion concentrations of neurons in a functioning brain. A typical use of EEG involves the attachment of one or several electrodes to coded locations of the scalp and monitoring changes in potentials. The signal of each pair of electrodes is amplified through a differential amplifier, filtered, recorded at a high sample rate (typically in the range of 128Hz-16kHz), and saved for later analysis.

EEG is widely used in a medical setting to monitor neurological diseases. For instance, patients with epilepsy often undergo EEG to observe and categorize seizures, which aids in the appropriate choice of treatment. Yet another medical use is to diagnose the condition or possible brain death of comatose patients.



Figure 1: The communication channels between the computer, (parts of) the user, and the BCI headset. The investigated side-channel attack is based on malware that the user has downloaded and installed. It records data from the channel between the BCI headset and the computer (dashed) and analyses it in order to guess secrets of the user. This attack differs from pure side-channel attacks in that it also contributes a signal to one of the channels (dotted). However, it is at the heart of the attack to modify this channel in a way that prevents the user to consciously notice the modifications.

In neuroscience research, EEG serves as a non-invasive, cost-effective method of measuring brain activity. Among other methods such as functional MRI, EEG is low in spatial resolution, such that it is not used to locate specific areas of brain function, but high in temporal resolution ideal for capturing minute temporal changes in the brain on the scale of milliseconds.

EEG recording devices vary in sampling frequency, number and location of electrodes, and general quality, depending on the needs of the user. Most recently, affordable, portable EEG devices have appeared on the market, availing themselves in the form of lightweight devices for gaming and personal EEG monitoring. Cognitive states picked up by such devices can be utilized by games to allow consumers to control on-screen avatars, monitor and train their own mental states, and improve gaming experience by collecting information on reactions and emotions.

Stimuli and event-related potentials (ERP) In many applications EEG signal is analyzed in conjunction with the presentation timing of visual or auditory stimuli. The consequent visible waveforms caused by the presentation of stimuli can be categorized into ERPs, which are combinations of negative and positive spikes occurring at different times after the stimuli, on the scale of milliseconds. These time intervals after a stimulus are often called **epochs**. One example of an ERP is the p300, a positive amplitude response that occurs around 300 milliseconds after certain stimuli, including images or sounds considered novel or threatening [15, 24]. The p300 has been successfully used in EEG devices to enable users to spell letters [19]. Another ERP of relevance is the N200, a negative amplitude response that often correlates with face stimuli and emotion [3].

3 Threat models and attacks

In this section we will describe the attack scenario and introduce the attacker capabilities. There are three agents in this game: the vendor, the user, and the attacker. The vendor of the BCI device runs a platform where developers can deploy their applications to be downloaded from users. The vendor provides the API of the device and distributes the devices (e.g., the API for designing third-party BCI games is already available for NeuroSky and Emotiv devices, see [14, 30]). The user has bought a device and uses it at home with various applications downloaded from the third-party developer platform. The user trusts her computer and the BCI device. The user actively supports setting up the device and also calibrates it, if necessary. We also assume that the user has an incentive to use the applications that she has downloaded.

The attacker is an application developer. Through the device’s API, the application can access the raw EEG signal recorded by the device. When the user executes the application, it can modify screen

content and audio and read input from all interfaces of the computer. Let us assume, for instance, that the attacker has modified a benign game or video viewer by inserting his malicious code and has posted the application on the platform with a slightly modified application name. When the user plays the game or uses the viewer to watch a video, the application can modify the game or video content or simply show additional content on top of it. The goal of the attacker is to obtain private user information. This could be any of the scenarios proposed in [26] such as guessing the banking provider, PINs, or month of birth. Generally, the target can be any memory of the user that could be useful for an attacker. For instance, the attacker could blackmail users who seem to be familiar with the logos of particular porn sites. Or a repressive political regime could try to identify whether users are familiar with some of the key persons of the underground opposition.



Figure 2: Example for a stimulus hidden in the video. The example shows two subsequent frames of the same video. At the second frame, the attacker inserts a face of a person that is suspected to be known by the user.

Attack Preparation: The attacker uses event-related potentials (ERPs) to run the following strategy. The attacker designs a number of visual stimuli that correspond to alternative answers of the questions to which he wants to find an answer. For instance, the attacker wants to know if the user knows a particular person. He collects an image of the suspected person. Then the attacker places this image and other visual stimuli, such as images of other people, at random times and random positions of the application and offers the application in the online application store. In our setting, the application is a video viewer, so the attacker includes images of people at random frames and at random screen locations of the video. See, for instance, the two video frames depicted in Figure 2. In this example, the attacker wants to find out if the suspect knows Barack Obama. Note that we have chosen this attack in our experiment because we already know the answer and can thus validate our method. In a real-world scenario it would be uninteresting to test for Obama except for internally validating a particular configuration of the attack.

Attack Execution: After the user has downloaded and installed the application and starts using it, the attacker collects the EEG signal recorded while the user is exposed to the different images displayed on the screen. Hoping that the person known to the user triggers the strongest event-related potentials, the attacker analyzes the EEG signal in a comparative way. This analysis works best if the recorded data can be contrasted with prior recordings of the victim, where it is known what the most relevant stimulus was. The ideal prior observations for this purpose is EEG data that has been generated by calibrating the BCI device. The usual calibration step consists of a sequence of numbers that are being flashed randomly. The user must count the occurrence of a particular number. Waiting for this number to appear is a very good way to provoke p300 artifacts. This calibration step of the device is often used by benign applications that rely on ERPs, too. Thus, ideally the attacker implants his attack in an application that requires this calibration step anyway.

4 Experiments

In this section we describe the setup of our proof-of-concept experiment. Investigating whether the proposed subconscious side-channel attack is feasible is a very challenging task due to many factors that can affect the results. A negative outcome could have many reasons. For instance, the equipment used could be suboptimal, the video used to hide the attack could have many still images where it is

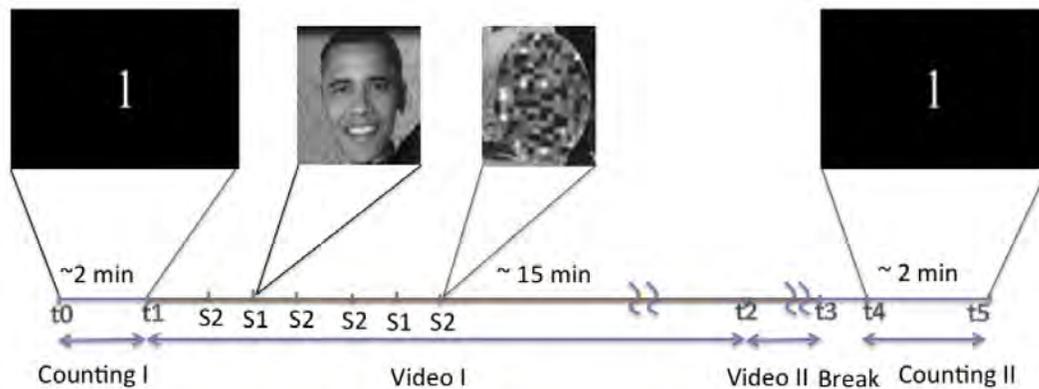


Figure 3: The experimental protocol is divided into 4 sub-experiments marked by t_i timestamps, namely Counting I, Video I, Video II and Counting II. The embedded stimuli S_j are depicted above the timeline.

hard to hide a stimulus, the secret that is being attacked could be too complex, and so on. For instance, let us suppose that we try to probe for the 16-digit credit card number of the user using a cheap BCI device and a video attack based on a still videolectures video. In case the attack does not work, it is hard to perform a root cause analysis of why such an attack might be unsuccessful. Therefore, our experiment is designed to investigate whether the attack is possible at all instead of starting off with sophisticated variants. In particular, we make design decisions that minimize the chance that the attack fails due to factors that we can control. For instance, we used a good EEG device and a video with many flickering artifacts to hide the attack. The following sections detail other design decisions.

4.1 Test Population and Setup

After obtaining approval from the Institutional Review Board, 29 undergraduate and graduate students (21 males and 8 females) in the Computer Science department were recruited to participate in our experiment, 2 of which had unusable data due to recording problems. All subjects were self-screened for neurological disorders and metal implants which could potentially interfere with recording. Prior to the experiment, subjects were informed of the basic EEG procedures, but not yet informed of the subliminal nature of the stimuli. The participants signed informed consent and received compensation in the form of a \$40 gift card. The experiment took 90 minutes total for each user, including setup time. This was the main limiting factor for population size. ActiveTwo BioSemi equipment [8] was used for the collection of EEG data. Participants were measured and fitted with a tight cap, and 64 Ag/AgCl electrodes were attached to the cap with conducting gel. All electrodes were then attached to a low-noise DC coupled post-amplifier, with a sampling rate of 1024 Hz. All stimuli were presented in a dim room on a CRT monitor (75Hz refresh rate) using presentation software [27].

4.2 Experimental Protocol

After the setup described above, the participants were asked to try to remain relaxed for the entire duration of the experiments. The interaction with the participants was kept as short and concise as possible. There were four parts to the complete experiment, two parts pertained to a counting task discussed below, and the other two tasks involved observing videos. Figure 3 semantically shows the four parts on a timeline for better understanding. The experiment lasted approximately one hour including setup. The following sections give details of each part of the experiment.

Counting I and II: In these parts of the experiment, the participant was presented with a randomly permuted sequence of numbers from 0 to 10. Each number except 1 appeared exactly 16 times. The digit 1 could appear anywhere between 14-18 times, chosen uniformly at random. The participant

was asked to count the number of occurrences of the number 1. Each stimulus lasted for 250 ms, and pauses between stimuli were randomly chosen to be between 250 ms and 375 ms long. At the end of this step of the experiment, the participants were asked for their count to check for correctness. This part of the experiment lasted for about 2 minutes. It was carried out in the beginning of the experiment (counting I) and at the end (counting II). These counting tasks are standard tasks used to calibrate BCI devices to users to ensure correct functionality of BCI applications.

Video I: In this phase, the participant was instructed to watch a 15 minute long black and white video extracted from Charlie Chaplin’s ”The Gold Rush” (1925). They were asked to pay attention to the plot of the video to make sure they concentrated on watching the video through its entire duration. Two kinds of stimuli (S1 and S2) were used, one with a black and white portrait of Barack Obama (Figure 3) (S1) and the other being a completely scrambled and blurred form, (S2). We choose these stimuli because we wanted to make sure that every subject was familiar with S1 and would not recognize S2. A stimulus was shown every 5 seconds, making a total of 180 stimuli over 15 minutes. Every 4th stimulus was S1 and was displayed at the top right corner of the image frame. The position of S2 rotated along the remaining three corners. A stimulus lasted for about 13.3 ms. The limiting factor of this time was the screen refresh rate. Once the video ended, the participant moved on to the next part of the experiment.

Video II: This phase is similar to the previous one for most of its structure. The participant was asked to watch a 1-minute continuation of the previous video. This video was embedded with a single stimulus, a black and white portrait of Michael Jackson (Figure 3). The stimulus appeared every 4 seconds and rotated along the four corners of the image frame. A total of 15 presentations of the stimulus were shown. The data from this portion was used for another study.

Recognition survey: While dismantling the electrodes, the users were asked if they noticed anything odd in the video. No further questions were asked if they negated. However, in the case that they did mention something, they were asked for details of what they saw. We categorize their answers as follows: participant recognized nothing, participant saw something, participant saw a face, participant saw ”Barack Obama”.

5 Data analysis

The methods described in this section serve to investigate whether subliminal side-channel attacks with BCIs are feasible if the attacker can control screen content. We are carrying out a proof-of-concept by running an attack under controlled conditions. As described above, the user has calibrated the BCI by actively participating in a counting experiment. Then, at a later point in time, the user watches a video while still wearing the device. The attacker can manipulate the video and insert small oval images of two kinds at random screen locations and at random times: a portrait photo of Barack Obama (see Figure 2 as an example) and a blurred image of a random person. With these two images we created a situation where most likely the subject knows the person in the first image (Obama) and does not know the person in the second image (Blur).

The attacker does not know which person is most relevant to the subject. It is the goal of the attacker to identify which person is known to the user by analyzing the EEG data that is collected while the user watches the video. From a technical perspective, this means that the classifier must analyze the different EEG sequences (the epochs) that have been recorded during the video and, based on this analysis, must decide for one out of three hypotheses: i) Obama, ii) the scrambled face, iii) the plain video sequences without any picture being shown.

5.1 Data Acquisition and Preprocessing

In this section, we describe how the raw data from EEG looks and what preprocessing steps need to be carried out before training the classifier. The raw data consists of wave signals from a number of different electrodes, called channels. Each channel is sampled at 1024Hz. We mark the EEG signals with exact positions of stimuli using the timestamps and indicators obtained from the software. This helps us to correlate the EEG signal with stimuli. For pre-processing, we first extract the signal 200 ms before and 1000 ms after every stimulus into epochs. Each such epoch is associated with the respective stimulus that triggers it. For each epoch, we then calculate the mean of the first 200 ms to get a baseline and subtract this baseline from the entire epoch. We reduce the high frequency noise

by passing the signals through a low pass filter with a pass band of [0.35, 0.4] in normalized frequency units. Finally, we apply a median filter that extracts the median from each four consecutive measurements. We use the data thus obtained for classifier training and stimulus prediction as described in the next section.

5.2 Classification

The attacker tries to identify if the stimulus is relevant for the user. This is a binary classification problem. Classification is a supervised learning technique and requires a training set of data containing observations whose class membership is known already. Each observations class is analyzed as a function of features describing the observation. It is this function which is later evaluated on a new observation to estimate if it belongs to a particular category. We now describe exactly how we adopt this general classifier framework to predict stimulus relevance from EEG signals.

In our setting an observation is an epoch. Each epoch corresponds to a single stimulus and contains the signals from all the EEG channels for a time period of [signal - 200ms, signal + 1000 ms]. If C denotes the number of channels being used and f denotes the sampling frequency, we have $(1000 + 200)f = S$ measurements per channel per epoch. We concatenate the signals from all the channels for each epoch into a feature vector of dimensionality $K = C \times S$. The classification algorithm consists of two phases, training and testing. In the training phase, the input samples provided to the classifier are of the form: $\{x_i \in X, y_i\}$ where $X = \{x_i \in \mathbb{R}^K, i = 1, \dots, N\}$, N denotes the total number of input epochs. The set $Y = \{y_i \in \{0, 1\}, i = 1, \dots, N\}$ denotes the class labels, i.e. whether epoch x_i corresponds to the target stimulus relevant to the subject ($y_i = 1$) or not ($y_i = 0$). Since the system used to display the stimuli captured the indicator of each stimuli and the corresponding timestamp, for each epoch, the value of y_i could be obtained. Given this set of inputs for training, the classifier learns the function that maps the feature vector (epochs) to the stimulus indicator: $f(x_i) : \mathbf{x}_i \in \mathbb{R}^K \rightarrow \mathbf{y} \in \{0, 1\}$.

In the testing phase, the classifier is provided with a set of fresh observations \mathbf{x}_i for which it must output label predictions \mathbf{y}_i . In other words, the classifier must predict for each epoch, if the corresponding stimulus shown to the participant is relevant or not.

We use a logistic regression method to make predictions on the stimulus relevance given the EEG signal. We train this classifier by minimizing the negative Bernoulli log-likelihood of the corresponding model in an iterative fashion as proposed in [17, 18]. A variant of this technique was used for p300 spelling in [19] with MATLAB code being available online. Also, Martinovic et al. used it for their attack [26] and it showed good performance for guessing user secrets from event-related potentials. For these reasons, we also chose this classifier for our experiment. As follows, we briefly describe how the classifier works. We will use the same notation as in the original paper [19] to simplify following up on details.

The boosted logistic regression classifier (BLR) variant that we employ is an ensemble method. This means it consists out of a set of $M \in \mathbb{N}$ individual classifiers f_m with $m \in \{1, \dots, M\}$ that all output individual classifier scores. Each classifier has a linear form $f_m(\mathbf{x}_i; \mathbf{w}_m) = \mathbf{w}_m^T \mathbf{x}_i$ with coefficients \mathbf{w} that differently weight the recorded channels at different points of time. These individual classifiers are incrementally blended into a single classifier F_m . At step m , the probabilistic model underlying classifier F_m is

$$p_m(y_i = 1 | \mathbf{x}_i) = \frac{\exp(F_m(\mathbf{x}_i))}{\exp(F_m(\mathbf{x}_i)) + \exp(-F_m(\mathbf{x}_i))} \quad (1)$$

and computes the probability that the stimulus of the current epoch i is a target stimulus given the concatenated EEG signal \mathbf{x}_i . The likelihood of this model for all epochs of the training data (i.e. the data where labels Y_i are given) is

$$L(F_m; \mathbf{X}, \mathbf{Y}) = \prod_{i=1}^N p_m(y_i = 1 | \mathbf{x}_i)^{y_i} (1 - p_m(y_i = 1 | \mathbf{x}_i))^{1 - y_i} \quad (2)$$

In the training phase, we iteratively optimize this function. At every optimization step m , the classifier of the last step is updated by adding a new weak classifier: $F_m = F_{m-1} + f_m$. The weights of this

additional classifier are computed by minimizing the least-squares distance of the gradient of the log-likelihood:

$$f_m = \underset{f}{\operatorname{argmin}} \sum_{i=1}^N \left(\left[\frac{\partial L(F(\mathbf{x}_i))}{\partial F(\mathbf{x}_i)} \right]_{F=F_{m-1}} - f_m(\mathbf{x}_i; \mathbf{w}_m) \right)^2 \quad (3)$$

At each update step, the new weak classifier f_m is added with a weight γ_m such that, finally, the ensemble classifier is $F_M = \sum_{m=1}^M \gamma_m f_m$. These weights are computed after the optimization step of the respective f_m with Eq. 3. It is selected such that Eq. (2) becomes maximal. Please see the full details of this algorithm and an experimental evaluation in [19].

In the described training phase, we provide all EEG data of the counting experiment together with the class labels to Eq. (2). Here epochs triggered by number 1 are accounted to class $y_i = 1$ and epochs with other numbers get the class label $y_i = 0$. For each weak classifier $C \cdot S$ coefficients $\mathbf{w} \in \mathbb{R}^K$ must be learned. If we record many channels at a high frame rate, then the optimization in Eq. (3) can become under-determined if too few epochs are available for training. We approached this problem of a low observation-to-dimension ratio by taking only those channels into account that are located along the z-axis, parietal, and occipital areas of the scalp where p300 ERPs are usually the strongest. In particular we took channels ‘Fz’, ‘Cz’, ‘Pz’, ‘P3’, ‘P4’, ‘PO7’, ‘PO8’, ‘Oz’. Also, for each channel we took the median of every 4 consecutive measurements to reduce the dimensionality by a factor of 4. This adds the side effect of denoising the signal by a median filter.

6 Results

In this section, we present and discuss the results of the conducted attacks. The attacks were explored using a series of different experiments, and we analyze them separately.

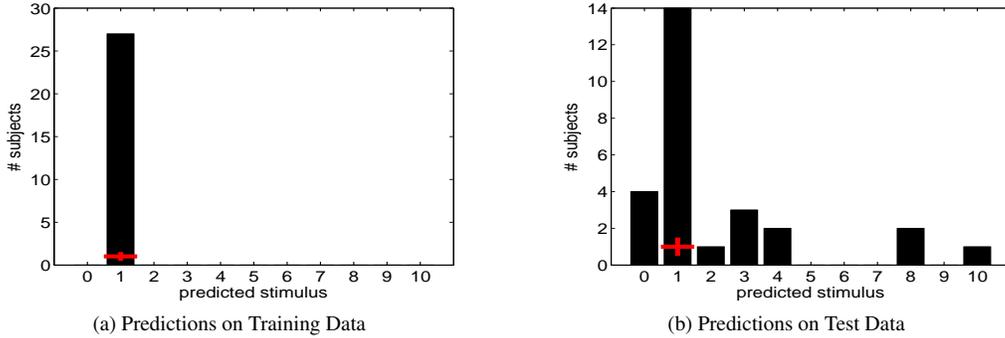


Figure 4: Predictions of a classifier trained on a counting experiment. Users counted the occurrence of number 1. The classifier must estimate which number between 0 and 10 was counted. On the train data, the predictions are perfect. On a second counting experiment, conducted 15 minutes later, some mis-classifications are apparent.

6.1 Classifier Training and Validation

Since the success of our attack depends fundamentally on the strength of the BLR classifier described in Section 5.2, we begin by affirming its reliability using a sanity check before challenging it with the more difficult subliminal attack problem.

In this experiment, we first train BLR with data collected from the first counting phase of the experiment. This is the phase where the user calibrates the device by counting the occurrence of the number 1 while a sequence of random numbers between 0 and 10 is repeatedly flashed on the screen. We extract the corresponding epochs from the EEG recordings and label all epochs corresponding to Stimulus 1 as the positive class and all other epochs as the negative class. This dataset serves as a training set to train BLR.

Next, we test BLR on the epochs of the second counting sequence, in which the user had to repeat

the counting task again. The only difference here may have been that the user was slightly less concentrated after having conducted the calibration phase and after having seen the videos. This is supported by the absolute counting errors of the users which were on average larger in the second counting phase (0.72) than in the first phase (0.38).

We test the classifier on both the training data from the first counting phase and the test data from the second counting phase. In the testing step, the classifier outputs a score for each stimulus. This score quantifies the classifier’s belief that the corresponding stimulus is the one that the user counted. There are 11 candidates, all numbers from 0 to 10. For each candidate number we average the score of all its corresponding stimuli. The classifier chooses the candidate number with the highest average score as its final prediction.

The predictions for both phases are depicted in Figure 4. As expected, the predictions on the training data are accurate. For all subjects the correct number (1) was guessed by BLR. On the test data, the predictions were correct for most users, though for some predictions were wrong. We conclude that the classifier works, but already observe that predicting the relevant stimulus for a user in a counting experiment is not straightforward. Even though counting the occurrence of a defined object evokes strong ERPs, the classifier does not show 100% accuracy on test data. This suggests that, on a technical level, predicting relevant stimuli is a difficult learning task, such that any success we have in our subliminal proof-of concept attack is significant. It also shows that there is room for improving our proposed attack from the machine learning perspective. In that respect, it is exciting to see that the machine learning community is advancing this field. For instance, recent work shows that the accuracy in p300 spelling is improved by learning across subjects by blending their classifiers via transfer learning [23].

In this validation analysis, we have separated training data and test data temporally because we were interested to see how persistent the classifier is over time. The longer the attack is, the more important this factor gets. When the attack is deployed on a user’s computer, a second counting phase might not be available. In this case, the attacker could just randomly sample training and test data from the calibration sequence that the user carries out prior to using the device. In this case, the attacker must rely on earlier results about temporal persistence or wait for a later calibration step to come. In Section 6.4, we will analyze alternative methods for training and validation that are applicable even if calibration data is unavailable.

6.2 Probing the Victim for Relevant Stimuli

After we have confirmed that the BLR classifier works, we run the actual attack on the data that is acquired from the Video Sequence 1. To carry out the attack, we use the BLR classifier that has been trained on the first counting sequence. We take the EEG data recorded while the user watches the video. From this data, we extract all epochs triggered by hidden images of Barack Obama, all epochs triggered by the blurry image and equally many epochs taken from random frames where the video was not manipulated. This defines three classes which we call ‘Obama’, ‘Blur’, and ‘Blank’. Again, we let the classifier output a score for each epoch of this dataset. Recall that, based on the training data used, this score outputs the classifier’s belief that the user has ‘counted’ the respective stimulus. Even though the user did not actively count the target stimulus (she should not even realize that it is on the screen), the classifier is searching for the same artifacts in the EEG signal.

As in the counting experiment, the prediction of the classifier is the candidate stimulus that gets the highest average classifier score. This time, there are three possible outcomes. Since we assume that all our users know Barack Obama, we can compare the classifier output against this ground truth. We depict the prediction statistics in Figure 5. For 18 users the classifier predicted the correct answer. For 5 users, BLR predicted ‘Blur’ and for 4 users BLR predicted ‘Blank’. From a machine learning perspective, it appears that the attack works, as the classifier is able to distinguish a relevant stimulus from irrelevant stimuli. From an attacker’s perspective, we now know that the user very likely recognizes Obama.

As we are ultimately interested in carrying out the attack subconsciously, we asked the users, after they saw the video, if they recognized anything strange in the video. If they answered affirmatively, we further asked what they saw. Some users could not further specify what they had seen. They had just realized that there was some flickering in the video. Some users described that they could see faces popping up here and there. Finally, some users were able to tell that they saw images of Barack

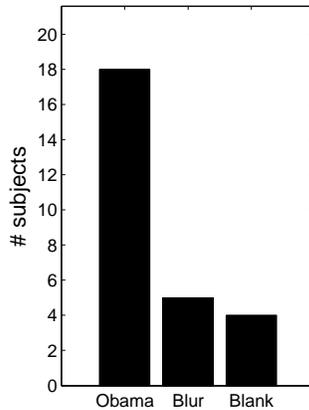


Figure 5: Detecting a hidden face in the video. In some of the frames small snippets of Barack Obamas face are hidden as well as a face that has been blurred. The classifier must predict if the target stimulus (Obama) was shown or not. The three possible candidates for the target stimulus are: the true answer (Obama), the blurred face (Blur), and phases where no images are shown (Blank).

Obama. For 4 users we do not know the answer. We resolve the classifier output by these different levels of user awareness and depict the results in a matrix in Figure 6. The subjects are grouped by five different columns, stating the four different recognition levels of the users plus the group of users without an answer to the recognition question. The rows represent the three different classifier outputs. For example, if a user saw ‘a face’ and the classifier predicted ‘blur’ as the target stimulus then the user populates the matrix entry (‘blur’, ‘I saw a face’).

prediction	Obama	3	3	2	6	4
	Blur	0	1	0	2	2
	Blank	1	1	1	0	1
		N/A	..nothing	..sth.	..a face	..Obama
				I saw...		

Figure 6: Classifier predictions vs. recognition level of the user. We distinguish between users of different awareness levels. Those that i) have not noticed anything in the video, ii) have noticed something iii) noticed that there is a face iv) noticed that Obama was shown.

Given that the stimuli were screened for 13.3 ms, it is interesting to see that for many users the stimuli were not subliminal. Neuroscience literature speaks of 10 ms to 55 ms as a suggested presentation time range for a stimulus to be subliminal (a good overview of designing experiments with subliminal stimuli can be found in [31]). This means that, in order to scale the attack, more sophisticated ways to hide stimuli are needed. In addition to using screens that allow for a higher frame rate (which is hard to control for the attacker) we elaborate other possible ways to achieve this in Section 7.

However, as can be seen from the matrix, the attack works almost independently of the extent to which the victims realize that the video has been manipulated. In each recognition group, the classifier found the correct answer for most users. For 5 participants the effect was completely subliminal, which means that subliminal attacks are a definitely a possible attack scenario. 3 users noticed ‘something’ but were not able to tell what it was. Clearly, this positive result can still be improved along several directions. As we have seen in the validation experiment, the classifier already provides some wrong answers for the rather clean counting data. This limitation will improve with the next generation of BCIs or as soon as machine learning methodology for this kind of data becomes more tailored for these kinds of tasks. In the next section we present an analysis method to obtain more reliable attacks already today.

6.3 Measures of Confidence

In this section we propose a method to make the prediction of the attack more reliable. We begin by distinguishing two attack scenarios and then introduce a way to compute confidence scores.

Targeted probing versus agnostic probing The experiments with the modified video are targeted towards a specific person. This means, the attacker suspects that the user might know a particular person (here Obama) and wants to probe for this hypothesis. He contrasts EEG data showing images of this person against images that are most likely irrelevant for the user. If the classifier ranks the epochs highest that correspond to the person of interest (here Obama), then the attacker concludes that the user knows this person. In an alternative attack scenario, the attacker does not know for which alternative to test for. In this agnostic attack, the attacker confronts the user with a number k of stimuli that are equally likely relevant for the user. Then the attacker must make a one-out-of- k decision based on the BCI traffic recorded. This task is more difficult than the targeted attack. In the targeted attack the attacker can simply discard the hypothesis if the target stimulus did not achieve the highest score. In agnostic probing, there will always be a stimulus that achieves the highest score. The attacker must decide if he trusts this outcome. In the next paragraph we propose a method to decide with confidence, and report on the results of an attack where we pretend that the attacker does not know in advance that he is probing for Obama.

Confidence scores for improved attacks In both the targeted attack and the agnostic attack, the attacker must decide whether to trust the classifier outcome and accept the hypothesis or to reject it. A technical way to base this decision is to compute a measure of confidence for a given classifier outcome and request a minimal confidence score for accepting the hypothesis.

For targeted attacks, we propose to compute the difference between the average classifier scores of the epochs that show the target class and the average score of the best non-target class. For instance, if we probe for Obama’s image and this class achieves the highest score and the blurry face achieves the second best score, then the measure of confidence is the difference between these two scores. Only if this difference exceeds a predefined threshold does the attacker accept it. In all results that we have reported so far we used a threshold of 0, i.e. the Obama hypothesis is accepted if only the corresponding epochs score the highest. In a scenario, where the attacker wants to avoid false positives, he can implement a higher threshold to get more prudent estimates. Interestingly, if the attacker is less risk-averse, he could even apply a negative confidence threshold to accept the hypothesis even if the target class is not ranked first by the classifier.

For agnostic attacks, the confidence score can be applied similarly. If the class that ranks highest has a classifier score that exceeds the second best class by a predefined confidence score, then the attacker accepts the hypothesis. If not, the outcome means that none of the presented stimuli is relevant for the victim.

In Figure 7, we report the results of these modified attacks. The x-axis depicts the applied threshold of the confidence score. The y-axis depicts the number of users for which the attack hypothesis was accepted by the attacker, given the current threshold. With agnostic probing, wrong hypotheses can also be accepted (we aggregate all wrong outcomes, ‘blur’ and ‘blank’, under a single error rate). With targeted probing, it can happen, that the hypothesis is rejected even though the victim knows the target (Obama). For a negative confidence threshold the attacker even accepts the Obama hypothesis if its epochs do not achieve the highest average classifier score.

As can be seen, the number of subjects where the attacker accepts the correct hypothesis declines for all attacks as the threshold of the confidence score increases. Luckily (for the attacker) the empirical chance of accepting a wrong answer declines much quicker than for accepting correct answers such that there is a regime in which users can be attacked with better accuracy. When an attacker deploys the attack, he must reason about risks and costs of false negatives and false positives and select a threshold accordingly. For instance, in Figure 7, we highlight a threshold that would be a prudent choice for an agnostic attack where false positives are costly for the attacker. For negative confidences (full-risk targeted attacks), more users are positively probed for Obama. However, it is not advisable to use such a negative threshold since it would probably also produce many false positives. A negative confidence score does not affect the agnostic attack, as the class that ranks highest is always most likely to be ‘target’.

The detailed outcome of the attack with a prudent confidence requirement is depicted in Figure 8. As can be seen the confidence criterion affects decisions at all levels of user awareness. In all these

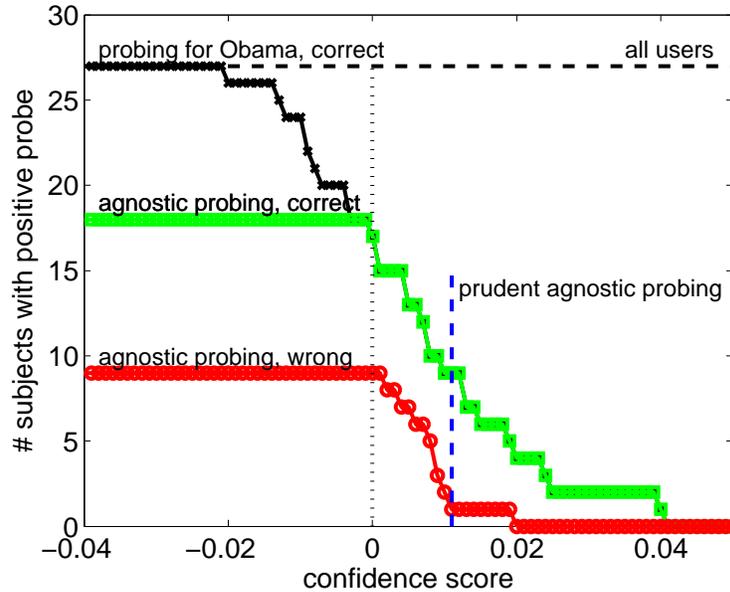


Figure 7: Implementing confidence scores. The 'target-stimulus' label is assigned if a predefined confidence score is met. If this threshold is too large, no victim will be positively tested. If it is too low, all victims get the respective label. It turns out that the wrong classifier decisions are rejected at lower confidence than the correct decisions which means that there is a regime in which confidence scores make the attack more reliable.

categories, the true answer stands out clearly, although overall there is also a loss of correct attacks. If the attacker can afford to loose a few victims and therefore get more reliable results he should apply the confidence criterion. The summary statistics of this result are given in Figure 9.

6.4 Persistence and Subliminal Calibration

In the last section we saw that it is possible to generalize a classifier trained on a counting calibration to a face recognition classifier that works on subliminally collected data. This experiment raises the interesting question as to what extent is this attack limited by the fact that the classifier has been trained on data that was recorded in a different situation than the attack situation. The resulting classifier is trained to detect what stimulus has been *counted* by the user and not directly on what stimulus is *familiar* to the user. Moreover, we have raised the question: how can the attacker validate the persistence of its classifier if no second counting experiment is available after the attack? In this case an attacker could use stimuli that are almost certainly known by the victim to test from time to time if the classifier still works. These stimuli are also useful if no counting phase is available at all, which could be the case if the user uses the BCI device exclusively for applications that do not require ERPs to function. The attacker could resort to training on subliminal stimuli that are known by the user.

prediction	Obama	1	1	2	4	1
	Blur	0	0	0	0	1
	Blank	0	0	0	0	0
		N/A	..nothing	..sth.	..a face	..Obama
				I saw...		

Figure 8: Detailed statistics for an attack using confidence scores.

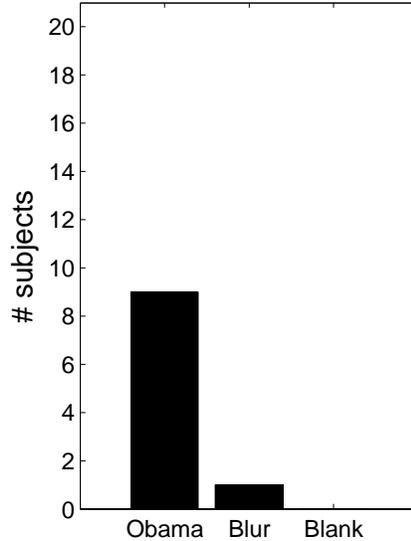


Figure 9: Summary statistics for an attack using confidence scores.

We can address all these questions by carrying out a different experiment where we train the classifier on the first half of epochs that have been recorded while the user watches the modified video. This time, we label all epochs that are triggered by an Obama image as the ‘target’ stimulus and all other epochs, corresponding to both blank video and blurry image, as ‘non-target’. We call this modified classifier BLR_F to account for the fact that it was trained on faces instead of numbers. Hence, the classifier is directly trained in an attack situation. As a result, the attack works significantly better as can be seen in Figure 10. We tested the classifier on the second half of the epochs collected from the modified video. So the test data is similar to the test data of the experiment depicted in Figure 6, except for the fact that the first half of the epochs are missing. We again group the users by the classifier predictions and by their level of recognizing the attack. This time the attack successfully worked for almost all users. It failed only for one user who recognized a face in the video. The statistics of the aggregated users across different levels of consciousness are shown in Figure 11a. This on-the-fly calibration of the attack leads to very accurate results and has the advantage that the user is not required to actively support a calibration phase. This means that the attack can even be deployed in applications where the user would not assume that event-related potentials could be recorded and analyzed.

prediction	Obama	4	5	3	7	7
	Blur	0	0	0	0	0
	Blank	0	0	0	1	0
		N/A	..nothing	..sth.	..a face	..Obama
				I saw...		

Figure 10: The detailed test predictions for a classifier trained on Faces hidden in the video.

7 Discussion

In this section, we critically discuss our attack and possible extensions. As we have opened up new grounds, there are many possible research directions to further analyze variants of the attack. Also, there are many aspects that can be improved beyond a proof-of-principle.

Long-term attacks. Although our experiment has shown that it is feasible within one EEG recording session to acquire information about the subject, the ideal setup in a real-world situation would be to collect EEG data over long periods of time.

The subliminal nature of the stimuli presents a large advantage over the method presented in Martinovic et al. in that the user may never realize anything strange is going on, and proceed to use the application and expose sensitive EEG data for large periods of time [26]. A larger body of data could be more useful in training the classifier and gaining more information about the user by making it possible to present a larger amount of subliminal stimuli. It would therefore be easier to slowly build a profile of various facts about the user which could be used against them.

However, from other EEG security work about attempts to use EEG as a form of authentication, we learn that EEG over multiple sessions may cause a decrease in accuracy, as electrode placement and small aspects of the device setup can change from session to session [4]. To address this problem, our method of subliminal validation with stimuli that the user certainly knows could be used to check if the data situation changes and to re-train the classifier if needed.

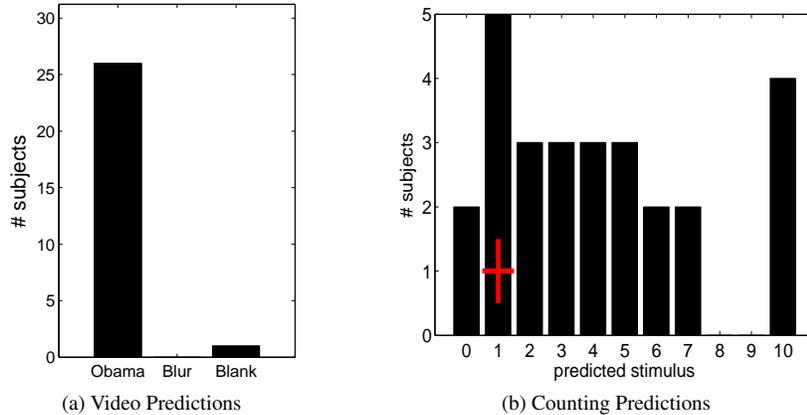


Figure 11: Test predictions for a classifier trained on Faces hidden in the video. We test the predictive power of the classifier on EEG data where users counted the occurrence of number 1 and on EEG data recorded while the user saw another part of the video with faces hidden.

Neuroscientific pitfalls. The improvement of the subliminally trained classifier BLR_F in Section 6.4 over our main result with BLR are striking. However, we also want to analyze an alternative interpretation of this variant of the attack and point to experimental pitfalls when analyzing activity of the human brain. The results in Figure 11 suggest that some of the wrong predictions in the main experiment are caused by the fact that BLR trained in a counting situation generalizes poorly to the attack situation. We investigate this hypothesis by testing the face-trained classifier BLR_F on the epochs of the counting sequence. The results of this validation test are reported in Figure 11b. Even though the BLR_F classifier has the biggest chance to predict the correct answer (number 1), it also predicts other numbers with a similar chance. So BLR_F generalizes worse from subliminal training to testing in a counting experiment than BLR translates from counting to testing on subliminal stimuli.

This means that the face-trained classifier BLR_F might be less useful than our counting based classifier BLR, for the following reason. As detecting and recognizing faces is of utmost importance for human-human interaction, parts of the brain play an important role for these tasks [32]. A face in the visual field triggers many face-detection and face-recognition processes and, as a consequence, lead to strong event-related potentials that might be detected by the BCI. The poor generalization from face-based subliminal training to counting data might stem from the fact that the BLR_F classifier searches for these signals related to processing images of faces. This means the improved result depicted in Figure 11a is in fact the result of a *face detector*, and not of a *relevance-detector*.

In this light, it is interesting to revisit the results of BLR, trained on counting data and tested on subliminal face images. The fact that the predictive power of the classifier translates from the counting scenario to the attack scenario but not vice versa, suggests that both classifier variants detect different neuro-physiological processes. While BLR_F seems to be a face detector in the first place, BLR tests

if the stimulus is relevant for the user. A promising way to further investigate this aspect would be to subliminally train a classifier on stimuli other than faces that are also certainly known to the user such as famous brand symbols or signs. Another way of challenging BLR_F is to train on a set of faces where only one face is known to the user. The results in [26], where a similar method has been probed in a conscious attack suggest that a face-trained classifier can in principle generalize to other kinds of stimuli. Overall, our observations illustrate that it is tricky to analyze brain activity data and that an attacker can always benefit from a neuro-scientific background to prevent misinterpretations.

Dry EEG. Our current experiment uses research-grade EEG recording equipment, which is relatively expensive and has a long setup process involving the injection of gel into the EEG cap. While more testing needs to be done on whether similar effects can be captured with current consumer-grade EEG devices, in the case that only such research-grade equipment works, users would still be protected from such attacks in the present, as it would be implausible to use such devices for daily use in conjunction with application software. However, companies are looking to create EEG devices with dry electrodes that have the same resolution as ones that require gel [5]. Setup time reduces dramatically, with much more capabilities attractive to app developers using this technology. When such devices are made affordable to the public, it becomes more likely that someone deploys subliminal attacks.

Complexity limitations of inferred content. Current methods of attacks through EEG devices are limited to presenting all of the possible values of the variable piece of information we wish to gain from the victim. For instance, while it is feasible to probe for a person being familiar to the user, it is unlikely that our attack can reveal the content of the last conversation of the user and this person. This is due to an explosion of complexity, as the victim must be exposed to all possible hypotheses (i.e. all conversations). Not only does this require more time to carry out the attack; it also requires more work on the part of the attacker to create stimuli for all such possibilities.

Improved strategies to hide stimuli. One direction in which our proof-of-concept could be improved is to better hide the stimuli in the video that is being shown. With respect to presenting the shortest stimulus time possible, we are constrained by the limits of the hardware, with a minimum stimulus time at 13.3ms dictated by the 75Hz screen refresh rate in our experiment.

Still, there are several ways to increase the subliminal effect through non-temporal means. For instance, stimuli which are presented in the foveal, or fixation point of vision require a significantly shorter stimulus presentation in order to maintain subliminality. Parafoveal stimuli, which are approximately one to five degrees from fixation, can maintain subliminal effects with a much longer stimulus time [20, 6]. This could be utilized in our experiment if it is known where the victim’s gaze is focused. This can be achieved provoking the user to look at particular screen locations using additional perceivable stimuli to attract their attention or by gaze tracking devices that might also become more popular in the future.

Another factor affecting the subliminal effect is the continuation of the visual processing of a stimulus, long after the stimulus has disappeared from the user’s sight. After a stimulus is shown on the screen, the image can still remain on the user’s retina for durations up to 30 milliseconds [2]. During this additional time the visual information is still being sent to the brain. This means the retina serves as a buffer that undermines the attackers effort of hiding the stimulus. Particularly, if the face image is shown on a low entropic background that stays calm for an extended period of time, the time that the user can “see” the face exceeds the actual duration at which it was displayed on the screen.

It has been shown that by placing an even stronger stimulus directly after the respective stimulus one can overwrite this visual buffer, a technique named backward masking [16]. Several forms of backward masking have been explored, ranging from bright flashes, patterns, and even noises, which would not block low-level visual processing, but could interrupt higher level vision processing in the brain [20, 16]. In order to not raise the user’s alertness, one cannot simply use another artificial image as a mask, of course. The art of masking a subliminal attack here would rather be to identify frames of the video/game that provide original sudden local changes in contrast, color, or sound. The stimulus could then be deployed directly before such an event occurs.

Yet another way to sidestep the screen refresh rate limit could be to partition the image into multiple parts and show each part separately in a successive frame. It would then depend on the brain to subconsciously reconstruct the image. Further testing would be needed to see whether this technique would work with a high-noise background such as a video.

8 Related Work

In this section, we overview existing literature on neuroscientific aspects of security-relevant applications.

Authentication The use of human ability to recognize and remember faces over a long period of time has been used to explore a new method of authentication using the identification of familiar faces within a grid of images [11]. It has been found that this method could relieve users of the need to memorize word-based passwords, and instead rely on the natural ability to remember faces. Even along a timespan of several months, this study shows that users were able to retain memory of their passfaces and correctly authenticate themselves. The use of commodity BCI devices has also incorporated another notion of passfaces [12]. Instead of manually choosing the correct faces out of a grid of images, eyetrackers were used, which consider a 0.5 second fixation on a face as a selection of it. Such uses of facial recognition are used in real-world situations, as in a security verification for Facebook users, which requires people to identify friends in tagged photos [1]. Thus, any information about which faces are familiar to a person can be considered vulnerable information. As our experiment demonstrates that facial recognition can be picked up through subliminal presentation of stimuli, this poses a threat to such security features that involve identifying familiar faces.

Bojinov et al. use *implicit learning* to train coercion-resistant passwords to users [9]. Through repeated execution of specific tasks the user learns behavioral patterns without being aware of what these patterns are. While the user is unable to tell about these patterns a computer system can validate if the user subconsciously 'knows' these patterns.

Subliminal Face Recognition. Given that facial recognition is seriously considered as a method of authentication, it becomes even more important to test the possibilities of extracting information from facial stimuli. There is existing neuroscientific work on the subliminal perception of human faces. ERPs in response to unpleasant expressions on faces have been shown to have a higher positive amplitude than pleasant expressions. This effect shows even through very fast unmasked subliminal presentations of stimuli, at 1ms [7]. This shows that visual information regarding faces can be processed and produce variance in the EEG signal even at a very subliminal level. Although in our experiment several subjects had noticed the stimuli, a presentation time as little as 1ms could still probably reveal enough information in their EEG signal to extract desired information about faces.

An interesting aspect about facial stimuli is that different facial expressions can produce very different ERPs. When expressions of fear are subconsciously exposed to the viewer, there is a large N200 ERP amplitude which seems to be associated with more primal pathways, allowing faster reactions to potentially dangerous stimuli [25]. Conscious exposure to the same fearful stimuli, however, have a stronger p300 amplitude, which is associated with more higher level processing of emotion. This can be important in choosing which stimuli to present to the user, and what information to use for training the classifier. For instance, training on a face with a fearful expression will yield a very different combination of ERP amplitudes, which would lead to inaccuracies if testing on happier expressions.

Although we have achieved some degree of success in our setup concerning face stimuli, it could be worthwhile to explore other types of stimuli, as the neurological response to faces is in many ways unique from other types of visual processing. Subliminal presentation of faces as opposed to words or randomized dots shows a general greater amplitude of ERP, suggesting that faces are processed differently [21].

9 Conclusion

In this work we have examined the question if subliminal attacks to users of EEG-based brain-computer interfaces (BCIs) are feasible. We have designed a proof-of-concept experiment in which the attacker tries to infer if the user knows a particular person or not, without the user noticing that she is being attacked. We hid visual stimuli in form of portrait photos of Barack Obama in a video as well as other visual stimuli that serve as a contrast. In an experiment with 27 subjects we find that our naive attack strategy is able to obtain 66% accuracy in predicting that a subject is familiar with Barack Obama, while an advanced attack strategy that incorporates confidence levels is able to improve the accuracy to 90%. The subjects achieved different levels of recognition in terms of detecting the manipulation of the video. At each recognition level, the attack was successful for most users including the users that did not notice any manipulation.

Our subliminal attacks have been carried out in a controlled setting to demonstrate their feasibility. Future research directions include exploration of different pathways for improving the attack, such as more sophisticated hiding mechanisms and internal subliminal validation techniques. The findings presented in this work suggest that BCI software with the full access to raw EEG data of users constitutes a new attack vector to user privacy and user secrets.

References

- [1] <http://www.facebook.com/help/search/?q=security+verification>.
- [2] Frank Allen. Effect upon the persistence of vision of exposing the eye to light of various wave lengths. *Physical Review (Series I)*, 11(5):257, 1900.
- [3] Truett Allison, Aina Puce, Dennis D. Spencer, and Gregory McCarthy. Electrophysiological studies of human face perception. i: Potentials generated in occipitotemporal cortex by face and non-face stimuli. *Cerebral Cortex*, 9(5):415–430, 1999.
- [4] C. Ashby, A. Bhatia, F. Tenore, and J. Vogelstein. Low-cost electroencephalogram (eeg) based authentication. In *Neural Engineering (NER), 2011 5th International IEEE/EMBS Conference on*, pages 442–445, 2011.
- [5] Ryan Baidya. Mynd by neurofocus. *Effect of Abuse and Neglect on Brain Development during Early Childhood*, page 10.
- [6] John A Bargh, Paula Pietromonaco, et al. Automatic information processing and social perception: The influence of trait information presented outside of conscious awareness on impression formation. *Journal of Personality and Social Psychology*, 43(3):437–449, 1982.
- [7] Edward Bernat, Scott Bunce, and Howard Shevrin. Event-related brain potentials differentiate positive and negative mood adjectives during both supraliminal and subliminal visual processing. *International Journal of Psychophysiology*, 42(1):11 – 34, 2001.
- [8] BioSemi. www.biosemi.com.
- [9] Hristo Bojinov, Daniel Sanchez, Paul Reber, Dan Boneh, and Patrick Lincoln. Neuroscience meets cryptography: Designing crypto primitives secure against rubber hose attacks. In *USENIX Conference on Security Symposium*, pages 33–33, Berkeley, CA, USA, 2012. USENIX Association.
- [10] BrainMaster Technologies. <http://www.brainmaster.com/>.
- [11] Sacha Brostoff and MAngela Sasse. Are passfaces more usable than passwords? a field trial investigation. In Sharon McDonald, Yvonne Waern, and Gilbert Cockton, editors, *People and Computers XIV –?? Usability or Else!*, pages 405–424. Springer London, 2000.
- [12] Paul Dunphy, Andrew Fitch, and Patrick Olivier. Gaze-contingent passwords at the atm. In *4th Conference on Communication by Gaze Interaction (COGAIN)*, 2008.
- [13] Emotiv Systems. www.emotiv.com.
- [14] Emotiv’s App Store. <http://www.emotiv.com/store/app.php>.
- [15] Noa Fogelson, Xue Wang, Jeffrey B. Lewis, Mark M. Kishiyama, Mingzhou Ding, and Robert T. Knight. Multimodal effects of local context on target detection: Evidence from p3b. *J. Cognitive Neuroscience*, 21(9):1680–1692, 2009.
- [16] D Friedman and C Fisher. Further observations on primary modes of perception, the use of a masking technique for subliminal visual stimulation. *The Behavioral and Brain Sciences*, 9(1):1–23, 1986.
- [17] J. Friedman, T. Hastie, and R. Tibshirani. Additive logistic regression: a statistical view of boosting. *Annals of Statistics*, 28:2000, 1998.

- [18] J. H. Friedman. Greedy function approximation: A gradient boosting machine. *Annals of Statistics*, 29:1189–1232, 2000.
- [19] U. Hoffmann, G. Garcia, J.-M. Vesin, K. Diserens, and T. Ebrahimi. A boosting approach to P300 detection with application to brain-computer interfaces. In *2nd International IEEE EMBS Conference on Neural Engineering*, pages 97–100, 2005.
- [20] D. Holender. Semantic activation without conscious identification in dichotic listening, parafoveal vision, and visual masking: A survey and appraisal. *The Behavioral and Brain Sciences*, 9(1):1–23, 1986.
- [21] Minoru Hoshiyama, Ryusuke Kakigi, Shoko Watanabe, Kensaku Miki, and Yasuyuki Takeshima. Brain responses for the subconscious recognition of faces. *Neuroscience Research*, 46(4):435–442, 2003.
- [22] Johan C. Karremans, Wolfgang Stroebe, and Jasper Claus. Beyond vicary’s fantasies: The impact of subliminal priming and brand choice. *Journal of Experimental Social Psychology*, 42(6):792–798, 2006.
- [23] Pieter-Jan Kindermans, Hannes Verschore, David Verstraeten, and Benjamin Schrauwen. A P300 BCI for the masses: Prior information enables instant unsupervised spelling. In *NIPS*, pages 719–727, 2012.
- [24] R. T. Knight. Contribution of human hippocampal region to novelty detection. *Nature*, 383:256–9, 1996.
- [25] Belinda J. Liddell, Leanne M. Williams, Jennifer Rathjen, Howard Shevrin, and Evian Gordon. A temporal dissociation of subliminal versus supraliminal fear perception: An event-related potential study. *J. Cognitive Neuroscience*, 16(3):479–486, April 2004.
- [26] 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 Association, Aug 2012.
- [27] Neurobehavioral Systems, Inc. <http://www.neurobs.com/>.
- [28] Neurooptimal Personal Trainer. <http://www.zengar.com/>.
- [29] NeuroSky Inc. www.neurosky.com.
- [30] NeuroSky’s App Store. <http://store.neurosky.com/>.
- [31] Nick Epley. Laboratory manual: Science or science fiction? investigating the possibility (and plausibility) of subliminal persuasion.
- [32] Josef Parvizi, Corentin Jacques, Brett L Foster, Nathan Withoft, Vinitha Rangarajan, Kevin S Weiner, and Kalanit Grill-Spector. Electrical stimulation of human fusiform face-selective regions distorts face perception. *The Journal of Neuroscience*, 32(43):14915–14920, 2012.
- [33] SmartBrain Technologies. <http://www.smartbraintech.com/store/pc/all-attention-brain-exercisers-c9.htm>.