Your Substance Abuse Disorder is an Open Secret! Gleaning Sensitive Personal Information from Templates in an EEG-based Authentication System

Richard Matovu, Abdul Serwadda
Texas Tech University, Lubbock, TX 79409
{richard.matovu, abdul.serwadda}@ttu.edu

Abstract

Given the task of designing an authentication system that uses brain waves as input, researchers typically focus on the sole objective of maximizing authentication accuracy. In this paper we challenge this common wisdom and argue that because brain waves encode a lot of other (potentially sensitive) information about the user, this single-pronged, privacy-agnostic approach can have significant privacy implications. Based on a publicly accessible dataset, we rigorously analyze two EEG-based authentication systems built in accordance with this philosophy and show that such designs could potentially divulge more of the users sensitive personal information than that regarding the intended authentication functionality. The paper argues for privacy-aware designs for systems which take brain signals as input.

1. Introduction

Recent research has shown brain activity patterns to depict significant promise in helping solve a broad range of cyber security problems (e.g., see [12][9][16][7]). Based on a plethora of brain measurement technologies (e.g., Electroencephalogram — EEG [9], functional Near Infra Red Spectroscopy — fNIRS [16], and functional Magnetic Resonance Imaging — fMRI [12]), the research community has showcased a good number of solutions to problems ranging from the identification of malicious web sites [12], continuous behavioral biometric authentication [9][16], and the detection of insider-threats [7] to mention but a few.

The prospect of brain measurement finding its way into people's daily interactions with technology has however been met with a fair share of skepticism. With the brain being the control center for all human activities, there is a widely held phobia about the kinds of private information that practitioners operating these technologies might — against the will of the users — glean from user's brain signals. While a huge volume of research continues to explore and front innovative brain-driven technologies, there is a surprisingly limited amount of research on the dynamics of this potential privacy violation problem.

As an example, take the case of authentication based on brain signals. In the vast majority of studies on this problem, researchers primarily seek to outdo each other in terms of the system error rates — i.e., they work with the central objective of designing a system having error rates which are much lower than the state-of-the-art. A critical question that never gets much attention is that of how certain design attributes of these systems (e.g., the kinds of features used to formulate the user template) might relate to their potential to leak sensitive personal information. If for example a system with the lowest authentication error rates comes with the added baggage of leaking a significantly higher amount of private information, then such a system might in practice not be as useful as its low error rates suggest. In practice users would only accept, and get the full utility of the system if the potential privacy breaches associated with the system are well understood and appropriate mitigations undertaken. This paper takes steps towards shedding some light on this privacy breach problem.

1.1. Research Thrust

At a high level, our research hypothesis is as follows: Assume a system which uses EEG as the modality for user authentication. Typically for such a system, all variables have been optimized to maximize authentication accuracy. A selection of such variables would include: (1) The features used to build user templates, (2) The signal frequency ranges from which features are extracted, and, (3) The regions of the brain on which the electrodes are placed among other variables.

Under this assumption of a finely tuned authentication system, we empirically tackle the following questions: — If a malicious entity were to somehow access templates from this authentication-optimized system, (1) would she be able to exploit these templates to infer *non-authentication-centric* information about the users *with high accuracy*? (2) In the event that such inferences are possible – would certain template designs (e.g., in terms of the attributes listed

in the previous paragraph) be significantly more vulnerable? At the heart of this hypothesis is the fact that while an EEG system may have been built for user authentication, the EEG signal encodes a lot of other information about the users. It is this information that we refer to as *non-authentication-centric* information. Such information includes users' emotions [17], medical conditions [18], and learning ability [14], to mention but a few.

To narrow down our exposition and enable rigorous analysis on a limited problem, our notion of *non-authentication-centric* information is limited to users' substance use behavior with the focus being on alcoholism — "a problematic pattern of alcohol use leading to clinically significant impairment or distress [3]". An individual's alcohol abuse behavior is private information protected under US Federal laws (see [2]), and could, depending on the circumstances around the leakage, spark off law suits among other issues.

As a vehicle to investigate the leakage of information pertaining to alcoholism, we build two EEG-based authentication systems, one similar to Chuang et al.'s system in [4] and the other loosely modeled similarly to that in [6][1]. With both systems fully tuned for user authentication, we play the devil's advocate and apply a range of pattern recognition techniques to perform a rigorous evaluation of the extent to which users' alcohol usage disorders could be inferred as a side-effect of the system's primary (authentication) purpose. Our specific findings and contributions are summarized next.

1.2. Our Findings and Contributions

1. For almost a quarter of the population in our study, we found that the authentication template divulged *significantly* more information about their alcohol use behavior than it did for the primary authentication purpose for which the system was designed.

2. By making changes in variables such as the numbers and locations of EEG electrodes and the features used for template building, we found that a template's propensity to leak alcohol usage behavior can be largely reduced while only causing a slight reduction in the mean authentication accuracy. This observation raises an interesting question of whether one could find an optimal mix of electrode locations which maximize authentication performance while minimizing information leakages about specific conditions that one might want to keep secret.

The rest of the paper is organized as follows. We present related work in Section 2 and describe our threat model and authentication design process in Section 3. Our findings on the privacy violation inferences are then presented in Section 4 before we make our conclusions in Section 5.

[1] We use the same electrodes and statistical measures as in [6], however, we do not transform our data to the wavelets space used in [6].

2. Related Work

Several studies have shown how EEG signals can be exploited as a side-channel to leak sensitive personal information. For example in [10], EEG data was shown to reveal information about details such as the month when one was born, the area in which one lives, the user's bank and the user's preferred bank card, to mention but a few. This information was revealed by monitoring user's EEG signals at the instants when certain carefully chosen stimuli were exposed to the users. In a related paper [11], the EEG signal was shown to delineate between people who believe in God and those who don't. This information was revealed while 28 subjects from mixed religious backgrounds undertook a stroop task and completed the Religious Zeal scale. Two key differences between these works and our research are that we: (1) simulate an attacker who only exploits the authentication templates and has no access to the raw (and much more informative) EEG stream, and (2) the sensitive private information we seek to exploit is embedded in the continuous EEG signal and does not depend on the analysis of the EEG stream registered during the brief spell when a user reacts to a specific stimuli (e.g., the data segment collected during the short time span when a user answers a question on where she lives).

The other stream of research that closely relates to our research is in the medical field, where a good number of papers have examined the connection between alcohol consumption and brain signals. For example in [5], a neural network was used to distinguish between the EEG patterns of subjects who were drunk from the patterns of those whose were not. Meanwhile in [13], investigations on the connection between the EEG beta power and alcohol use behavior were made. A key variation of these works from our work is that we focus on the connection between an EEG *user authentication* template and the leakage of a user's health-related information, explicitly evaluating the amount of information leaked and how different authentication template designs impact this leakage. To our knowledge no past work has addressed these research questions.

3. Threat Model and Authentication System Details

3.1. Threat Model

We assume an attacker who somehow gains access to a database of EEG authentication templates. Such access could be gained in several ways, e.g., by hacking into the template database, or by an unscrupulous insider such as a database administrator who misuses her privileges and decides to maliciously exploit the templates. To more clearly

highlight the magnitude of the potential threat, our presentation in the rest of the paper will assume the insider attacker scenario. Such an attacker would have fine details about things such as the template formulation approach, which would in turn enable her to exploit the system to the fullest.

Having made the alcohol usage investigations, the inferred information, if shared with management, could, for example, form the basis for further observations on the performance of employees in question, which might eventually lead to them being victimized. Note than an insider attacker could easily determine which template belongs to whom.

3.2. Description of Dataset used in the Study

Our experiments are based on a dataset collected by the Neurodynamics Laboratory, SUNY Downstate Medical Center. The dataset was part of the data used in a large study to examine the EEG correlates of genetic predisposition to alcoholism (see dataset source [1] and a recent detailed description of the dataset [8]). The dataset comprises of 77 alcoholic subjects and 45 control subjects. The alcoholic subjects were diagnosed as such after formal medical evaluations and had a history of heavy drinking for at least 15 years. For a minimum of 30 days before the data collection, the alcoholic subjects were fully detoxified.

Data was collected while subjects were exposed to three different kinds of stimuli. The stimuli were pictures chosen from the 1980 Snodgrass and Vanderwart picture-set. In one experiment a single stimuli (S_1) was show to each subject. In another experiment, two stimuli (S_1 and S_2) were successively shown to a subject, with the second stimulus shown to the subject about 1.6 seconds after the first. The two stimuli were in some cases selected such that they matched while in other cases they did not match. Subjects pressed a different mouse key depending on whether they determined that the stimuli matched or not. The EEG device used for data collection had 64 electrodes. Measurements from each electrode were sampled at 256Hz for 1 second. More details about the full dataset can be found in [8][1]. The data used in this paper is only from stimulus S_1 and is from 50 subjects; 25 from each of the alcoholic and control groups. We refer to each 1-second segment from a user as a sample. The 50 users selected for this work had between 30 and 80 samples each.

3.3. EEG Authentication Systems Under Investigation

3.3.1 EEG Feature-sets/Templates

Here we describe the structure of the templates used by our two described experimental systems previously introduced in Section 1.

Template #1: This template uses data collected from the 6 blue electrodes (see Figure 1). After computing the power spectral density of the signal from each electrode, only data from the alpha (8-12 Hz) and beta (12-30 Hz) ranges is retained. From each of the alpha and beta ranges, 3 features are computed; namely the entropy, mean and standard deviation. Each electrode gives a total of 6 features.

Template #2: Like was done in [4], this template uses features derived from only one electrode; the Front Polar (FP-1) electrode colored yellow in Figure 1. Again, only data from the alpha (8-12 Hz) and beta (12-30Hz) ranges is retained after computing the power spectral density of the signal from this electrode. For each frequency component in this range, the median magnitude of the psd is recorded as a feature.

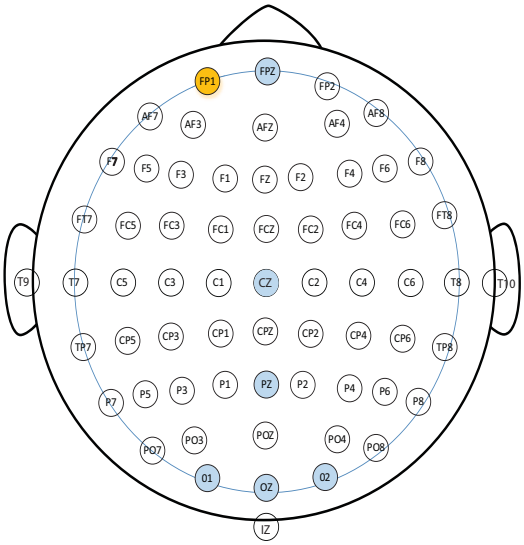

Figure 1. Locations of EEG electrodes used to collect the data used in this study. One of the templates studied in this work is built based on data collected from only the yellow electrode (as was done in [4]), while the second template is based on data collected from the 6 blue electrodes.

3.3.2 Classification Performance of Authentication Systems

Classifier	Template #1		Template #2	
	Mean HTER	Std HTER	Mean HTER	Std HTER
Random Forest	0.155	0.093	0.238	0.117
Naive Bayes	0.218	0.094	0.264	0.139

Table 1. Baseline authentication results. Std HTER stands for the standard deviation of the Half Total Error Rate.

Before proceeding with our privacy violation investigations, we first evaluate the performance of the two templates (or systems) for user authentication. For our findings to be easily put in the context of the state-of-the-art, it is

crucial that our system's authentication performance should not drastically differ from the average results seen in past literature. Table 1 summarizes the authentication performance of our two systems. The mean Half Total Error Rates (HTER)[2] reported in the table were computed based on data from 50 users; with each user providing 30 samples. 70% of this data was used for training while 30% was used for testing. The table shows that the two templates give mean HTERs of between 0.155 and 0.264. While these error rates are far from those of a system that is ready for deployment, they are not so different from those reported previously in the literature (e.g., see error rates for various activities in [9][4]). With our systems performing similarly to findings in past research, we proceeded to study the privacy violation problem.

4. Examining the Privacy Violation Attacks

4.1. Do EEG Authentication Templates Leak a *Significant* Amount of Information About Alcoholism?

4.1.1 Metric For Analyzing Information Leakage

In this sub-section we address the question of whether *a statistically significant* dependence exists between a user's EEG authentication template and their alcohol use behavior. If such a connection exists, then the adversary could have a good chance of inferring a user's alcohol usage behavior from their biometric authentication template. It is noteworthy that the investigations in this sub-section are meant to provide a preliminary exploration into the feasibility of the privacy leakage attack and do not simulate the attack itself. Building on the intuition developed in this section, an account of what the attacker would find is given in Section 4.2.

To understand the dependency between a user template and alcohol use behavior, we use the mutual information metric. Unlike traditional correlation measures such as Pearson's, Spearman's or Kendall's which only detect linear dependencies, mutual information detects non linear relationships as well, and hence enables us to get a rigorous account of how an EEG biometric template might leak information about the user's alcohol use behavior. The mutual information $I(X;Y)$ between two random variables X and Y is defined by Equation 1. P_X and P_Y respectively represent the probability mass functions (*pmf*) of X and Y while P_{XY} is the *pmf* of the joint distribution between X and Y.

$$I(X;Y) = \sum_{x \in X} \sum_{y \in Y} P_{XY}(x,y) log_2 \frac{P_{XY}(x,y)}{P_X(x) P_Y(y)} \quad (1)$$

[2]The HTER is the average of the False Accept Rate and False Reject Rate. It has been used before in several EEG authentication studies (e.g., see [9][4])

4.1.2 Approach and Rationale Behind Information Leakage Analysis

We tackle the mutual information question from two angles. From one angle, we study the mutual information between the biometric features and the biometric class labels (which are the user IDs). Results from this evaluation give us insights into how strongly the features reduce uncertainty about the users' identities. To conveniently refer to this setting later in the discussion, we assign it a short acronym and refer to it as the MI_{auth} setting. From the second angle, we study the mutual information between the (same) biometric features and the alcohol usage class labels (which are labels on whether users are alcoholic or not). Results from this evaluation give us insights into how strongly the features reduce uncertainty about the users' alcohol usage behavior. For back reference during the rest of the discussion, we will refer to this second angle of analysis as the MI_{alco} setting.

In an idealized EEG authentication *system which attains high authentication accuracy while leaking minimal or no information about users' alcohol abuse behavior*, results from the MI_{auth} setting should exhibit high amounts of mutual information between the features and the user IDs while results from the MI_{alco} setting should exhibit much less mutual information (or a much lower extent of dependency) between the (same) features and alcohol usage behavior. If the dependencies seen from either setting are comparable to each other, or if the MI_{alco} setting exhibits higher mutual information (i.e., stronger dependences), then it is likely that the system's high authentication accuracy comes with the baggage of a comparably high, or even higher accuracy with regard to the inference of users' alcohol usage behavior.

Supported by a series of statistical tests of significance, the logic expressed in the previous paragraph will form the backbone of our mutual information investigations. A detailed account of how we undertake the mutual information computations follows in Sections 4.1.3, 4.1.4 and 4.1.5.

4.1.3 Dependency Between Features and User IDs

We compute the mutual information between the features and user IDs as follows (i.e., the MI_{auth} case): *For each user and each template type*, we create a feature matrix F_{auth} for which each row is either: (1) a feature vector extracted from a sample belonging to the user, or, (2) a feature vector extracted from a sample belonging to a different user (i.e., an impostor). The matrix in total contains 30 such feature vectors (or rows) from the user and 30 such feature vectors (or rows) from the impostors. The impostor samples are randomly selected from all the others users. Each column in matrix F_{auth} represents a single feature.

For comparison with F_{auth}, we create a vector C_{auth} for which each element is a class label mapping to the cor-

responding row in F_{auth}. For example, if the first row in F_{auth} is a feature vector belonging to $User$ #1, then the first element in C_{auth} is the user ID for *User* #1. If on the other hand a given row in F_{auth} is a feature vector belonging to a user other than *User* #1 (i.e., an impostor), then the corresponding element in C_{auth} has the label of an impostor. All impostors, irrespective of their User IDs have the same class label, which means that all 60 elements in C_{auth} are either of two labels; — the user's ID, and the impostor label. For each one of the 50 users in our study, we create the matrix F_{auth} and the corresponding vector C_{auth} for each of Templates #1 and #2.

Given a matrix F_{auth} and vector C_{auth} generated for a given user and template, the mutual information between the i_{th} feature and class labels is the mutual information between the i_{th} column in F_{auth} and the vector C_{auth}. To perform the mutual information computations, we discretize F_{auth} into 20 bins and C_{auth} into 2 bins (C_{auth} inherently has 2 bins since it has 2 discrete labels). After computing the mutual information for all the features, we generate a vector I_{auth} which contains mutual information values for the user and template in question. We undertake this mutual information computation for all 50 users and obtain 50 vectors similar to the vector I_{auth}. Each value in each one of the 50 I_{auth} vectors gives us some measure of how knowledge of the corresponding feature reduces our uncertainty about the class labels in C_{auth}.

4.1.4 Dependency Between Features and Alcoholic Behavior

For the mutual information computations in this case, we create a matrix F_{alco} which contains: (1) 30 rows where each row is a feature vector representing a sample randomly drawn from the sub-set of users who are alcoholic, and, (2) 30 rows where each row is a feature vector representing a sample randomly drawn from the sub-set of users who belong to the control group (i.e., not alcoholic). Like was the case with F_{auth}, each column in F_{alco} represents a single feature. Further, we create a vector C_{alco} in which each element represents the class label of a corresponding row in F_{alco}. The class labels in this case are one of two values: *alcoholic* or *not alcoholic*. Similarly to what we did in Section 4.1.3, we compute the mutual information between each feature (i.e., each column in F_{alco}) and the vector C_{alco} to create the vector I_{alco}. Each element in I_{alco} gives us some measure of how knowledge of the corresponding feature reduces our uncertainty about the class labels in C_{alco}.

4.1.5 Comparing the Feature Dependences

The elements in I_{auth} and I_{alco} can be compared pair-wise because elements at a given index in either vector represent the same feature. The only difference between the elements of the 2 vectors is that the elements in I_{auth} represent mutual information from the perspective of user authentication, while those in I_{alco} represent mutual information from the perspective of separability between an alcoholic user and a user who is not alcoholic. If for a certain user the elements in I_{alco} are significantly larger than their corresponding elements in I_{auth}, this means that the feature in question is much better at reducing uncertainty about alcoholic behavior than about a user's identity.

To determine if the elements in I_{alco} are significantly larger than those in I_{auth}, one can use the paired t-test if the two vectors follow a Gaussian distribution or can choose between the Wilcoxon Signed Rank test and the Sign test if the vectors do not follow a Gaussian distribution. The Wilcoxon Signed Rank test requires that the vectors are symmetric while the Sign test is generally considered not to be so powerful because it only uses the signs and not the sizes of the differences between the elements of the vectors being compared [15]. In this work we zeroed on the Wilcoxon Signed Rank test after finding that the majority of our difference vectors were not Gaussian (ruling out the

ALGORITHM 1: Evaluating Information Leakage

Input: ω_{eeg} , U , $C_{template}$
// ω_{eeg}: EEG data
// U: List of user IDs
// $C_{template}$: Authentication template
Output: U_{count}
// Count of users for whom we fail to reject the null hypothesis
$I_{auth}, I_{alco} \leftarrow \emptyset$
for *each* $u_j \in U$ **do**
  $F_{auth} \leftarrow getAuthFeatures(\omega_{eeg}, u_j)$
  $C_{auth} \leftarrow getAuthClassVector(\omega_{eeg}, u_j)$
  for $i \leftarrow 1$ ***to*** n **do**
    $I_{auth}[j] \leftarrow computeMI(F_{auth}[i], C_{auth})$

$F_{alco} \leftarrow getAlcoFeatures(\omega_{eeg})$
$C_{alco} \leftarrow getAlcoClassVector(\omega_{eeg})$
for $i \leftarrow 1$ ***to*** n **do**
  $I_{alco} \leftarrow computeMI(F_{alco}[i], C_{alco})$

$U_{count} = 0$
for *each* $u_j \in U$ **do**
  $p_{value} \leftarrow wilcox.test(I_{auth}[j], I_{alco}, 0.05)$
  if $p_{value} > 0.05$ **then**
    $U_{count} = U_{count} + 1$
return U_{count}

t-test) yet were either approximately skewed or moderately skewed. See [15] for a past study which applied this test to a moderately skewed dataset.

Formally, for each user the null hypothesis is that the differences vector $I_{alco} - I_{auth}$ follows a continuous symmetric distribution with zero median. The alternative hypothesis is that the differences vector $I_{alco} - I_{auth}$ follows a symmetric continuous distribution with median less than zero. Informally, rejection of the null hypothesis in favor of the alternative hypothesis implies that the elements in I_{alco} are on average less than those in I_{auth}, which means that for this particular user, the template (or features) in question provide stronger discriminative power for the user authentication problem than for the problem of determining users' alcoholic behavior. *Here we are interested in counting the number of users for whom we are unable to reject the null hypothesis*. For this subset of users, the implication of the test result is that there is no evidence to indicate that the authentication template in question performs worse at giving away their alcoholic behavior than it does at performing authentication. An attacker who accesses such templates from the authentication server could *potentially* exploit them to infer alcohol usage behavior with higher accuracy than the accuracy which the system attains during its primary user authentication function. Note our usage of the word *potentially* — the test result does not offer guarantees on what the attacker would find for this subset of users. It however gives us a lead that warrants further exploration (done in Section 4.2).

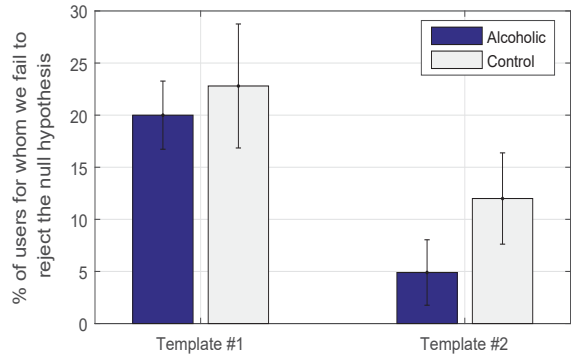

Figure 2. % of users for whom we are unable to reject H_o in favor of H_1. H_o specifies that a template (or feature-set) reduces our uncertainty for the user authentication problem in the same way it does for the alcoholism inference problem. H_1 specifies that a template (or feature-set) reduces our uncertainty for the user authentication problem to a lesser extent than in does for the alcoholism inference problem. α for the tests is 0.05.

Algorithm 1 summarizes this process; it inputs the EEG data as ω_{eeg}, a list of user IDs as U, $C_{template}$ as authentication template in question and outputs U_{count}, a count of the users for whom we fail to reject the null hypothesis. The function $computeMI()$ obtains the dependency (mutual information) between a feature matrix and class vector. I_{auth} and I_{alco} are generated in lines #7 - #11 and #13 - #16 respectively. A wilcoxon signed rank test is then computed between I_{alco} and each user's I_{auth} on line #20 to obtain a p_{value}. When the p_{value} is greater than 0.05, the user's count for that template, U_{count}, is incremented by one (1) implying we fail to reject the null hypothesis for that user.

Figure 2 summarizes our findings from the statistical tests. For each of Template # 1 and Template #2, the figure shows the percentage of alcoholic users for whom we fail to reject the null hypothesis, and the percentage of users in the control group (non alcoholic users) for whom we fail to reject the null hypothesis. For Template # 1 we were unable to reject H_o for between 20 to 25% of the users while for Template # 2 this number was much less – specifically between 5 and 15% depending on whether the users were alcoholic or not. Looking at the sizable difference seen between the two templates (20 to 25% for one template and 5 and 15% for the other), these results suggest that the sheer careful design of the templates (e.g., identities of electrodes used, metrics computed, etc.) could largely reduce the proportion of users whose templates are vulnerable to attack.

4.2. Accuracy of Inference of Alcoholic Behavior

The statistical tests performed in the previous section provide a solid foundation to one's understanding of aspects such as the proportion of users who could be susceptible to attack. These tests do a good job at giving us a classifier-independent perspective of the potential privacy problem, however, to get a more concrete picture of what the attacker could infer in practice, it is interesting to carry out an explicit supervised classification process. In practice the attacker can find training data from various sources — e.g., the dataset used in this study is freely available online. Other kinds of training datasets can be similarly found if the attacker is interested in a variable other than alcoholism. After training the classifier with this data, she can then make inferences about users' templates. Table 2, shows findings from this supervised classification setting. We use data from 25 of our users for training, and the other 25 for testing with each user providing 30 samples. The class labels in this classification are whether a given sample comes from

Classifier	Template #1		Template #2	
	Mean HTER	Std HTER	Mean HTER	Std HTER
Random Forest	0.245	0.011	0.421	0.017
Naive Bayes	0.318	0.001	0.549	0.043

Table 2. Accuracy of classification when the class labels mapped to the alcohol use behavior (i.e., one class is "alcoholic user"; the other class is "not alcoholic user").

an alcoholic user or not. The HTERs in the table (Table 2) are computed as an average of the individual HTERs of the 25 users in the test set. The results indicate that Template # 1 enables a higher classification accuracy for alcohol usage behavior than Template # 2. This agrees with the findings in Section 4.1.5 where Template # 1 had a higher proportion of users who we found to have a high possibility of being vulnerable to the privacy violation attacks. The table further shows that Template # 2 has very poor classification accuracies (approximately equivalent to random guessing). This observation suggests that the sheer change of template could largely minimize the potential of success of the attacks. This observation is even much more interesting when one notes that Template # 2 did not perform much worse than Template # 1 at user authentication (recall Table 1, Section 3.3.2). The result suggests that by compromising a little on authentication accuracy, one could make significant gains in terms of reducing the kinds of information that could be inferred from the authentication template.

5. Conclusions

In this paper we have studied the problem of a biometric authentication template leaking a user's non authentication-centric sensitive personal information. We have found that by tweaking the design parameters of the system, one could largely reduce the information leakages without significantly compromising authentication accuracy. Further, we have found that there exist certain users who are significantly more vulnerable than others to these kinds of attacks. The paper calls for brain waves authentication implementations that optimize authentication performance while keeping track of potential privacy violations.

References

[1] Eeg database. `http://kdd.ics.uci.edu/databases/eeg/eeg.data.html`. Last accessed in March, 2016.

[2] Health information privacy. `http://www.hhs.gov/hipaa/for-professionals/special-topics/related-links/index.html`. Last accessed in March, 2016.

[3] Internet mental health. `http://www.mentalhealth.com/home/dx/alcoholdependence.html`. Last accessed in March, 2016.

[4] J. Chuang, H. Nguyen, C. Wang, and B. Johnson. I think, therefore i am: usability and security of authentication using brainwaves. In A. A. Adams, M. Brenner, and M. Smith, editors, *Financial Cryptography and Data Security*, volume 7862 of *Lecture Notes in Computer Science*, pages 1–16, Berlin Heidelberg, 2013. Springer.

[5] Z. Ekhi, A. Akgul, and M. R. Bozkurt. Article: The classificaton of eeg signals recorded in drunk and non-drunk people. *International Journal of Computer Applications*, 68(10):40–44, April 2013.

[6] Q. Gui, Z. Jin, and W. Xu. Exploring eeg-based biometrics for user identification and authentication. In *Signal Processing in Medicine and Biology Symposium (SPMB), 2014 IEEE*, pages 1–6, Dec 2014.

[7] Y. Hashem, H. Takabi, M. GhasemiGol, and R. Dantu. Towards insider threat detection using psychophysiological signals. In *Proceedings of the 7th ACM CCS International Workshop on Managing Insider Security Threats*, MIST '15, pages 71–74, New York, NY, USA, 2015. ACM.

[8] N. Karamzadeh, Y. Ardeshirpour, M. Kellman, F. Chowdhry, A. Anderson, D. Chorlian, E. Wegman, and A. Gandjbakhche. Relative brain signature: a population-based feature extraction procedure to identify functional biomarkers in the brain of alcoholics. *Brain and Behavior*, 5(7), 2015.

[9] S. Marcel and J. d. R. Millan. Person authentication using brainwaves (eeg) and maximum a posteriori model adaptation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 29(4):743–752, Apr. 2007.

[10] I. Martinovic, D. Davies, M. Frank, D. Perito, T. Ros, and D. Song. On the feasibility of side-channel attacks with brain-computer interfaces. In *Proceedings of the 21st USENIX Conference on Security Symposium*, Security'12, pages 34–34, Berkeley, CA, USA, 2012.

[11] I. Michael, M. Ian, H. Jacob, and N. Kyle. Neural markers of religious conviction. *The Association of Psychological Science*, 20(3), 2009.

[12] A. Neupane, M. L. Rahman, N. Saxena, and L. Hirshfield. A multi-modal neuro-physiological study of phishing detection and malware warnings. In *Proceedings of the 22Nd ACM SIGSAC Conference on Computer and Communications Security*, CCS '15, pages 479–491, New York, NY, USA, 2015.

[13] M. Rangaswamy, B. Porjesz, D. B. Chorlian, K. Wang, K. A. Jones, L. O. Bauer, J. Rohrbaugh, S. J. OConnor, S. Kuperman, T. Reich, and H. Begleiter. Beta power in the {EEG} of alcoholics. *Biological Psychiatry*, 52(8):831 – 842, 2002.

[14] N. A. Rashid, M. N. Taib, S. Lias, N. Sulaiman, Z. H. Murat, and R. S. S. A. Kadir. Learners learning style classification related to {IQ} and stress based on {EEG}. *Procedia - Social and Behavioral Sciences*, 29:1061 – 1070, 2011. The 2nd International Conference on Education and Educational Psychology 2011.

[15] A. Serwadda and V. V. Phoha. Examining a large keystroke biometrics dataset for statistical-attack openings. *ACM Trans. Inf. Syst. Secur.*, 16(2):8:1–8:30, Sept. 2013.

[16] A. Serwadda, V. V. Phoha, S. Poudel, L. M. Hirshfield, D. Bandara, S. E. Bratt, and M. R. Costa. fnirs: A new modality for brain activity-based biometric authentication. In *Biometrics Theory, Applications and Systems (BTAS), 2015 IEEE 7th International Conference on*, pages 1–7, Sept 2015.

[17] M. Soleymani, S. Asghari-Esfeden, M. Pantic, and Y. Fu. Continuous emotion detection using eeg signals and facial expressions. In *Multimedia and Expo (ICME), 2014 IEEE International Conference on*, pages 1–6, July 2014.

[18] A. T. Tzallas, M. G. Tsipouras, and D. I. Fotiadis. Epileptic seizure detection in eegs using time-frequency analysis. *IEEE Transactions on Information Technology in Biomedicine*, 13(5):703–710, Sept 2009.