Prediction of Human Error Using Eye Movements Patterns for Unintentional Insider Threat Detection

Hassan Takabi, Yessir Hashem, Ram Dantu
Department of Computer Science and Engineering, University of North Texas
Denton, Texas, USA
Takabi@unt.edu, YassirHashem@my.unt.edu, Ram.Dantu@unt.edu

Abstract

Threats from the inside of an organization's perimeters are a significant problem since it is difficult to distinguish them from benign activities. Recent reports indicate that the accidental/unintentional incidents account for the majority of all insider security incidents. Human error is a major factor in unintentional insider threat. In this paper, we propose a novel approach for unintentional insider threat (UIT) detection and mitigation based on eye movement patterns. We perform experiments to capture unique characteristics of a user's eye movements as they perform several computer-based activities in different scenarios. The goal is to evaluate the effectiveness of using eye movement patterns in determining a user's subjective mental workload which is one of the main contributing factors to human error. We extract eye movement and pupil features which allow us to reliably achieve this goal. We evaluate our proposed approach using several classifiers and examine how different subsets of features affect the performance. The results show about 82% accuracy on average for users wearing eye glasses and an average accuracy of 84.5% for users without eye glasses. Our results demonstrate that users' eye movement patterns and pupil behaviors can reveal valuable clues about their subjective mental workload and could be used in developing effective tools for unintentional insider threat detection and mitigation in real-world environments.

1. Introduction

The "insider threat" has received significant attention and is considered one of the most serious security problems [20]. It is also considered the most difficult problem to deal with because insiders often have information and capabilities not known to external attackers, and as a consequence can cause serious harm. However, little is known about the insider threat and the threat of insider activities continues to be of paramount concern in both the public and private sectors [18].

Recent reports show that a significant concern of security professionals is employees accidentally jeopardizing security through data leaks or similar errors. We use unintentional insider threat (UIT) to refer to this type of threat which is also referred as an accidental insider threat, inadvertent insider threat, etc. Verizon 2015 Data Breach Investigation Report found that insiders are responsible for 90% of security incidents and 71% of insider threats are unintentional [26]. Similar trends are echoed in a recent report by Ponemon Institute found that the accidental/unintentional incidents account for 70% of all insider security incidents [5]. The same report shows that the cost of time wasted responding to security incidents caused by human error to be as much as 1.5 million for a U.S. company and 1.6 million for a German company. It also shows that if negligence was reduced by as much as 50%, an average of 31% (U.S.) or 28% (Germany) of IT security spending could be saved.

1.1. Unintentional Insider Threat and Human Error

The CERT defines an unintentional insider threat as "(1) a current or former employee, contractor, or business partner (2) who has or had authorized access to an organization's network, system, or data and who, (3) through action or inaction without malicious intent, (4) causes harm or substantially increases the probability of future serious harm to the confidentiality, integrity, or availability of the organization's information or information systems." [23]. In addition to threats such as phishing, spear phishing, and social engineering in general, examples of unintentional insider threats include the inadvertent disclosure of confidential information, accidentally exposing company networks

2018 IEEE 4^{th} International Conference on Identity, Security, and Behavior Analysis (ISBA)
978-1-5386-2248-3/18/$31.00

to hacking due to misconfiguration for instance, accidentally leaking sensitive information on social networks or mishandling sensitive information such as sending to the wrong email.

Although there has been a significant amount of research on insider threat, most of the efforts focus on intentional or malicious insider threats and little research has explicitly focused on UIT and UIT as a research topic has largely been unrecognized until only recently [24]. It is very difficult if not impossible to map common patterns and contributing factors from intentional insider cases to UIT cases. Many of the contributing factors to malicious insider threats especially motivational factors do not play a role in UIT cases. However, it is possible to study the literature in several domains such as human factors, cognitive science, safety, and human error to suggest possible contributing factors to UIT.

Human errors are the cause of about 80 percent of accidents in different contexts ranging from air transport operations to nuclear power plants. Human error analysis techniques have long been a mainstay of effective safety programs, with each implementation being tailored for a particular context. Researchers in human error have adapted human information processing models to illustrate how individuals process hazards and how errors may occur in different stages of processing. In the communication-human information processing (C-HIP) model by Conzola et al. [9], stimuli from the environment enter channels in the human sensory system. Errors and limitations in the cognitive processes of perception, attention, comprehension, and decision making may produce bottlenecks and decrease performance.

It is widely accepted and documented that many mishaps of various types are inadvertent or unintentional, and that human errors underlying these mishaps often form patterns of recurrence when examined over time [19]. A detailed analysis of real-world UIT cases by CERT found that a large number of UIT incidents are caused due to mental fatigue, mind wandering, mood, etc. that contribute to human error [23]. Fortunately, evidence suggests that many precursors that have been identified as relevant to safety are also relevant to security-related errors [19].

One of the most important challenges in predicting and stopping unintentional insider threats is the difficulty of predicting human errors and failure in human performance which are the primary factor in UIT. Although human errors cannot be entirely eliminated, being able to predict them allows the organizations to take proper actions to mitigate human error and consequently reduce the potential effects. Human error has been investigated in other domains, however, most of the previous studies traditionally rely on self-reporting mechanisms or physiological measures such as functional magnetic resonance imaging (fMRI), electroencephalogram (EEG), which are either inaccurate or invasive, not practical for day-to-day use, and not directly applicable to UIT.

1.2. Contributions

In this study, we take the first important steps towards gaining a better understanding of UIT and ultimately developing tools to mitigate it. Our goal is to utilize recent advances in eye tracking technologies and devices in investigating their application in UIT mitigation and developing a non-invasive approach to predict human error which is a major cause of UIT incidents. The neuroscience community has made significant advances in understanding the connections between the retina and the human brain regions that generate eye movements [22]. The eye tracking devices can capture spontaneous responses that are unfiltered by the conscious mind and generate a rich and distinctive feature space of voluntary, involuntary, and reflexive eye movements. Eye tracking devices have become cheaper, and a variety of inexpensive devices including open source hardware and software are available [25].

The current non-invasive video-based eye tracking technologies do not require any physical contact with the users and can easily be used in a conventional workplace.

In cybersecurity domain, the eye movement dynamics have been studied mainly for authentication and identification [3][4][14]. In this study, we investigate the effectiveness of using eye tracking technology and eye movement dynamics in predicting human error for help unintentional insider threat incidents. Our study is informed by the fact that conscious and unconscious mental processes generate different patterns of eye movements. Our goal is to capture these changes using eye tracking technology and use it for unintentional insider threat detection and mitigation frameworks. We design experiments to capture the unique characteristics of eye movements patterns using an eye tracker. We perform experiments with 25 participants and capture their eye movements while performing different tasks. We extract unique features of each participant's eye movements and to ensure the time stability of our extracted features. The participants perform the experiment tasks during two sessions on different days and at different times of day as explained in section 3.3.

We use several different classifiers to evaluate our experiment results. Our proposed framework includes the following components:

- Data acquisition: to record the participant's eye movements during the experiment, we used a Tobii Pro X2-60 device, a screen-based eye tracker capturing gaze data at 60 Hz [25]. The device delivers accurate gaze position data within the entire experiment screen.

- Feature extraction: we use the output of the device to extract useful features including pupil, temporal, and

dimensional features.

- Activity classification: We use four different classification algorithms to detect the malicious activities.

The rest of the paper is organized as follows. In Section 2, we provide a brief background about the human visual system. Section 3 describes our experimental design including the experiments setup and experiment tasks. Section 4 presents the experimental results, and in Section 5 we discuss the related work. Section 6 discusses the limitations of the proposed approach. Finally, section 7 discusses the future work and concludes the paper.

2. Background: Human Eye and Eye Movements

Human eyes are the visual organs of the human body that help him to visualize the world around him and eye movements direct the optical axes to a new position in order to observe. In the human visual system, the eyes move to scan the visual scenes, the eyes' pupils control the light and pass it to the retina. The retina consists of light-sensitive cells located at the back of the eye that converts the light into electrochemical signals, and these signals travel along the optic nerve to the visual cortex part of the brain. As a result of this process, human eyes frequently move to observe the objects. Six muscles control the voluntary or involuntary movements of the eyes. These muscles give the eyes the flexibility of movement in 6 degrees. There are two main types of eye movements: the saccades and the fixations.

Saccades are the movements of both eyes in the same direction to receive visual information and used to shift the line of vision from one position to another. Saccadic duration increases from about 20ms for the smallest movements to about 100ms for the largest ones [1]. On the other hand, the fixations are related to maintaining of the visual gaze on the objects, and they have very low-velocity movements compared to the Saccades with around 100 to over 500ms duration.

3. Methodology

Our primary goal in this study is to design a framework to predict human error using the human eye movement patterns and pupils characteristics and behaviors. We aim to analyze users' eye movements while they perform several activities and determine whether participants are more likely to make an error. In general, there are two types of experiments that can be conducted for the purpose of eye tracking. The first is the control experiment where the participant reacts to visual stimuli displayed on the screen, and the device records his reaction. The second type is the free experiment where there are no specific stimuli and the participant freely moves his eyes on the screen. In order to closely imitate the real-world unintentional insider threat scenarios we aim to address in this study, we have chosen the second type of experiments. We do not use any stimuli in the experiments, and instead we focus on developing tasks are close to realistic situations.

Our experimental design includes four different tasks that are a mix of regular computer-based activities such as data entry, browsing the Internet, using applications, etc. Each task emulates a real-world scenario very close to a typical work environment as described in section 3.4.

Our participants performed the tasks in two sessions, two tasks per each session. In order to ensure the stability of the eye movements features during different times, eye conditions, and levels of fatigue, the sessions were conducted on two separate days and at different times of the day. The main reason is to test participants in different situations w.r.t mental workload which may affect eye movements. The first session was conducted at the early morning hours while the participant is still fresh and the second session was conducted at the end of the day when the participant is probably tired as a result of performing many daily activities.

The experiments were conducted with the approval of Institutional Review Board (IRB) from the University of North Texas, and the participants were compensated $30 for one hour of their time. In the followings, we describe the study and the experiments in more detail.

3.1. Study Participants

We recruited a total of 30 participants to participate in our experiments, but we use data of 25 participants in our data analysis; the data records of the other participants were incomplete and were removed from the analysis. Out of 25 participants, 15 were male, and 10 were female. All participants were between the age of 18 and 34 years old and were a graduate or undergraduate students at the University of North Texas with different levels of programming skills and cybersecurity knowledge. There were a total of 10 participants with prescription eye glasses to emulate real-world scenarios for normal work environment where some employees may wear eye glasses. Also, we wanted to ensure that our approach can be applied to any users regardless of the condition of their eyes.

3.2. Experiment Environment and Tracking Device

For our experiments, we used the Tobii Pro X2-60 eye tracking device, a screen-based eye tracker that captures the gaze data at 60 Hz [25]. The device delivers the gaze position data within 24-inch screen with a 1920x1200 resolution experiment screen. We used a regular personal computer with Intel Core i5-4210U Processor, 6GB PC3 DDR3L SDRAM, and Windows 10 Home operating system.

The experiments were conducted in a room which was set up to keep the same environmental conditions for all

tasks and all participants. For example, the lighting and the screen brightness and resolution were the same during the experiments. Any change in the environment can affect the eye features gathered from the participants, so we made sure to stabilize the environmental factors.

3.3. Experiment Procedure

The total time for the experiments was one hour. The experiment was divided into four different tasks, each task being about 10 minutes long. There were two regular activity tasks and two tasks that were performed under a high mental workload. The experiments were performed on two separate days, one at the early morning and one at the afternoon. The participants were given a pre-screening form. The form includes questions about user's health, education level, programming skills, cybersecurity knowledge, etc. After filling the form, the participants were briefed on the objectives of the study and given a written informed consent form to read and sign. The consent form includes a precise information about experiment procedure and participant's right to participation.

Once the consent form was obtained, the participants have seated in a comfortable chair about 2 feet away from the computer screen. Then, the examiner performed the device calibration process. The calibration was performed before each task to ensure the device was capturing the eye movements data with the best accuracy. The examiner then explained the task and what the participant was supposed to do step by step using a task script document located on the computer desktop. The participants were given five minutes to read the script before each task and to feel comfortable with the test environment.

3.4. Data Collection

In the followings, we describe the tasks participants performed during the experiment.

3.4.1 Task 1: users perform regular daily activities

In this task, the participating subject received through email an excel sheet containing names of students who participated in a previous survey and their associated information. The participant needed to use the browser to log in to the database system and find out the students' names and update their missing data using the excel sheet. Participants were not required to finish all the students' records, and there was no pressure introduced during the task.

3.4.2 Task 2: users perform high mental workload activities

The participants were asked to complete a short coding project (designing a calculator) which demands more mental workload compared to regular daily activities in Task 1. The participants were allowed to choose any programing language they felt comfortable with (C++, Java, Python). They were also allowed to browse the Internet for help if needed. However, the participants were told that copying the code from the Internet was not allowed and they had to write their own code. The participants were encouraged to finish the code project. However, the task did not require them to complete the project and there was no pressure introduced to the experiment.

3.4.3 Task 3: users perform regular daily activities under pressure

In this task, we repeat Task 1 with some changes to introduce stress to the participants in order to emulate the scenario where employees work under pressure or emotional stress.

To do this, we conducted the experiment at the end of the working day, so the participants came to the lab after attending classes, exams or labs during the day causing them to have more mental workload compared to Task 1. The participants were also given twice the number of the records compared to the first task and asked to make sure to finish all the records during the 10 minutes experiment. In addition, we removed some of the students' names from the database to add more pressure as the participants were not able to find those names. The participants were told that there will be a special prize for the participant who finished his/her report faster than the others adding more stress and time control to the experiment.

3.4.4 Task 4: users perform high mental workload activities under pressure

In this task, we repeated Task 2 with some changes to introduce stress to the participants in order to emulate situations where employees work under pressure or emotional stress. To do this, we conducted the experiment at the end of the working day, so the participants came to the lab after attending classes, exams or labs during the day causing them to have more mental workload compared to Task 3. The participants were asked to repeat the same project but with another programing language. Participants were still able to browse the Internet for help. However, the participants had to complete the code and test it in 10 minutes. The participants were told that there will be a special prize for the participant who completed his/her code faster than the others.

3.5. Data Analysis

The eye tracking device records the eye position of our participants on the screen (x,y coordinates) by generating raw data which include the saccade and fixation positions

on the screen as well as the current pupil diameter. The device sampling frequency is 60Hz (60 samples per second). We use these recorded raw data to extract and calculate useful features set that can reveal the behavior of the eye in a particular time window. We use 10-second time window to sample our recordings. In other words, we segment our recorded time to 10-second time windows where each one represents one sample and extract the features from that sample and label it with its current task activity (regular or high subjective mental workload). In this work we extract two type of features:

- The pupil's features: we extract the pupil diameter and calculate the statistic measures such as the minimal, maximal, mean and standard deviation of the pupil diameter values.
- The eye movements features: we extract a total of 38 eye movements features including the timing and velocity features such as the saccade and fixation duration, pairwise movement speed and acceleration. And the spatial features such as the pairwise distance, the distance from the center of the screen, and the direction of saccades.

Also, we consider other useful features such as the eye blinking, fixations, and saccades frequency. After extracting the features and labeling our features vector, we use four different classifiers to evaluate our approach: the Support Vector Machine (SVM) classifiers, k-nearest-neighbors (k-NN), Random Forests classifier and Bagging predictors .

4. Experimental Results

Once the experiments are conducted for all participants, we extract the eye movements and pupil features per every 10 seconds recording time window and label the extracted feature vectors to the activities scenario the participant performed as the following:

- We assign the feature vectors extracted from each participant and obtained from the regular activities tasks a class label = "0".
- The feature vectors extracted from each participant and obtained from the tasks under high subjective mental workload are given a class label ="1".

First, we train a separate classifier for each participant (private model setup) and build a model for each participant. In other words, the training and testing process performed on the same participant data. We divide the data into 70% tuning and training dataset, and 30% testing dataset. In more details, we split the 10 minutes task time to seven minutes of training and three minutes for testing. We disregard the last ten seconds of the seven minutes and the first ten seconds of the three minutes to ensure a good separation between the training and testing sets. In order to evaluate our results, we use four evaluation metrics. The classification accuracy shows the percentage of correctly classified activities from all the activities. The true positives (TP) rate to measure how many activities under high subjective mental workload are classified correctly, the false positives (FP) rate to show how many regular activities are misclassified as being performed under a high subjective mental workload. Finally, we use the F-measure to present the geometric mean between the precision and recall of the classifiers' results.

We run the classifiers for each participants' dataset and calculate the average results among all the participants. Table 1 shows the average results of our 25 participants using four classifiers. The results indicate that the Random Forest classifier outperformed the other three classifiers and were able to achieve up to 84.52% average classification accuracy. Also, the Random Forest classifier provides better average true positives rate, average false positives rate, and average F-measure. Bagging classifier achieved the second best results with 81.36% average classification accuracy. The lowest accuracy at 70.30% was obtained by the k-NN classifier which also shows lower performance with the other evaluation parameters.

Moreover, we extend our evaluation to measure the performance change between participants who have healthy eyes and participants who wear eye glasses. As in real life scenarios, users have different eyes health conditions, and it is likely that some of the employees wear eye glasses or contact lenses.

Table 2 presents the average results for 10 participants who were wearing glasses during the experiments. The results show that the Random Forest classifier still outperforms the other classifiers and has the best average accuracy at 81.93%. These results are slightly below the average results from the entire participants set and indicate that the lower results are mostly related to the participants who wore eye glasses as glasses can reflect the light and impact the recording accuracy of the eye tracking device. In contrast, the 15 participants who had healthy eyes (didn't use eye glasses during the experiments) show an average detection accuracy up to 86.24% using the Random Forest classifier which is higher than the average for all the participants as expected.

All the previous results were achieved using seven minutes of activities to train the classifiers. However, we might not have access to that much data for training in all situations. In order to address this issue, we train classifiers with much fewer data and measure the performance of the approach; then, we gradually increase the amount of training data to measure the impact of increased training data on the accuracy of the proposed approach. We start by using only

Classifier	Avg. Acc	Avg. TP Rate	Avg. FP Rate	Avg. F-Measure
SVM	76.79	0.77	0.24	0.77
Random Forest	**84.52**	0.85	0.17	0.84
k-NN	70.30	0.70	0.30	0.70
Bagging	81.36	0.81	0.20	0.81

Table 1: Results for distinguishing regular activities from the ones performed under high subjective mental workload for all the participants.

Classifier	Eyes condition	of Participants	Avg. Acc	Avg. TP Rate	Avg. FP Rate	Avg. F-Measure
SVM	Without Eyes Glasses	15	78.37	0.78	0.32	0.78
Random Forest	Without Eyes Glasses	15	**86.24**	0.86	0.15	0.86
k-NN	Without Eyes Glasses	15	70.58	0.71	0.29	0.71
Bagging	Without Eyes Glasses	15	81.73	0.86	0.15	0.86
SVM	With Eyes Glasses	10	74.43	0.74	0.26	0.74
Random Forest	With Eyes Glasses	10	**81.93**	0.82	0.19	0.82
k-NN	With Eyes Glasses	10	69.86	0.70	0.31	0.70
Bagging	With Eyes Glasses	10	80.82	0.81	0.20	0.80

Table 2: Results for distinguishing regular activities from the ones performed under high subjective mental workload for two groups of participants based on their eye health condition.

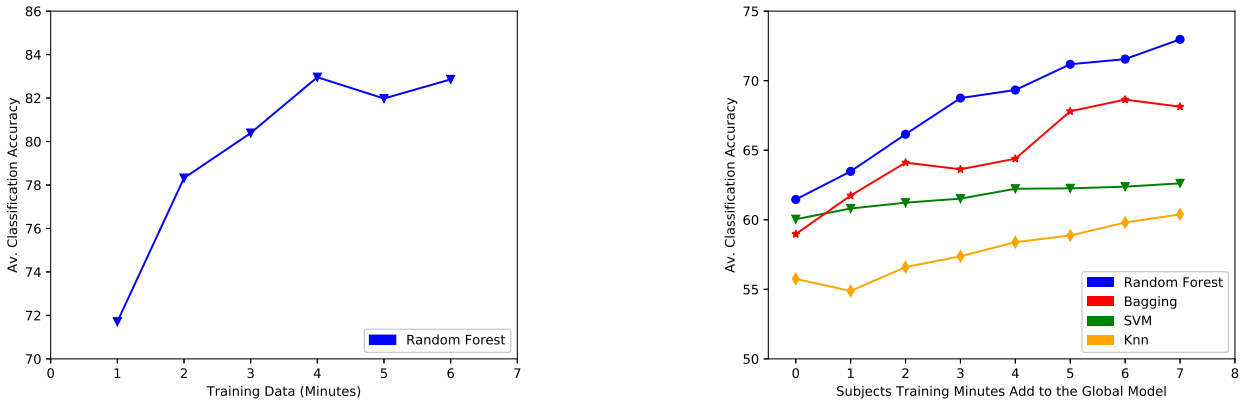

(a) (b)

Figure 1: (a) Results of using different amount of training data. (b) Classification results for using the global model.

one minute of activities for training, and then we add another minute to the training data and repeat the process and record the detection accuracy along the way. The goal is to see how much training data we need in order to achieve a reasonable detection accuracy.

Figure 1a shows the results of the Random Forest classifier using different amounts of training data. As we can see, the average detection accuracy increases when we increase the amount of training data. However, when we reach to 4 minutes of training data, the results remain steady. These results suggest that we can reduce the amount of time the participants required to perform for training the classifier to four minutes without incurring any impact on the detection performance.

Finally, building different models for each employee might not be an ideal solution and some organizations might prefer to build one global model for all their employees rather than generating a private model per each employee. For this purpose, we extend our proposed approach using transfer learning technique [17] to test the performance of building a global model on a particular set of users and testing on a different set of users. Since some eye features are unique to the individual and as previous research has shown the use of eye movements on the field of biometric identification [21], we expected that the global model (group-based model) would result in lower performance than the private model (user-based model) setting. We repeat our evaluation by training a classifier using data from a group of partici-

pants and test on the data from another participant without using any of his/her data in the training phase. Then, we slightly add some of the participant's data to the training data and re-train the model.

Figure 1b shows the results of building a global model of our four classifiers. As the results show, our four classifiers perform very low when none of the participant's data is used in the training phase; the best average accuracy is about 61.47% and achieved by the Random Forest classifier. However, as we add some of the participant's data to the training process, the accuracy improves. The average accuracy increases to 63.48% when we add one minute of the participant's data to the training data. The performance consistently increases for the four classifiers as we increase the training data of the tested participant until we reach to the average accuracy of 72.97% as our best detection accuracy using the Random Forest classifier. These results still show lower performance compared to the private model, but this was expected as some eye movements features are unique to an individual and negatively impact the training process and performance of the classifiers.

5. Related Work

There has recently been some work on different aspects of unintentional insider threats. The CERT Insider Threat Team has released three reports on unintentional insider threats studying contributing factors, social engineering, and phishing and malware incidents [23][24]. They have summarized the relevant research to identify possible contributing factors [23]. Greitzer et al. have developed an operational definition of UIT [7] and examined UIT cases that derive from social engineering exploits to identify possible behavioral and technical patterns [8].

The eye tracking and gaze tracking technologies have been mainly used for authentication and identification. For example, eye movements have been used to propose a gaze-based password entry [3]. Cantoni et al. proposed a task-dependent approach to implement a biometric framework based on the density and duration of fixations for the user's eye while looking at different pictures [2].

On the other hand, Hashem et al. introduced a multi-modal neuro-physiological study based on the user's psychophysiological and computer-based behaviors that uses the Electroencephalogram (EEG), Electrocardiogram (ECG), eye movement dynamics and the mouse movement and keystroke dynamics to reveal valuable knowledge about users malicious intent and utilize them to build an insider threat monitoring and detection framework [10][11][12][13].

Two recent studies utilize the eye movements for insider threat mitigation. Neupane et al. conducted a three-dimensional study of phishing detection and malware warnings, focusing on what users' task performance and how users process these tasks based on (1) neural activity captured using Electroencephalogram (EEG) cognitive metrics, and (2) eye gaze patterns captured using an eye-tracker [16]. In this work, the eye movements are used only to identify whether users look at the security indicators.

Another recent study used the eye tracking to propose a biometric-based on distinctive eye movement patterns with the goal of mitigating the insider threat in so-called lunchtime attack scenarios where a person temporarily gains physical access to a workstation that he is not supposed to use (e.g., using a coworkers workstation while he is at lunch) [4]. The approach they propose is essentially a biometric-based authentication mechanism to detect if someone else other than the authenticated user has physical access to a device.

6. Limitations and Discussion

One potential drawback of the study is the eye tracking system used in experiments and its applicability in the real world. Although we used a specialized hardware in this study, there have been significant advances in eye tracking technologies in recent years, and several inexpensive eye tracking systems have been developed that can achieve good accuracy [6][15]. Due to the advances in the video-based tracking devices, these technologies can be implemented using a simple webcam or a built-in laptop webcam with a simple software. It would be useful to perform a field study using built-in laptop/computer webcam which addresses both issues of time constraint and the applicability of the hardware; this is left to future work. It is worth noting that since our goal is developing a UIT prediction method using eye movements and based on the behavioral and psychological perspectives of the user, this new approach can be combined with other technical and behavioral approaches to increase the accuracy and the performance. As mentioned earlier, UIT is mostly unexplored and this work is a first step towards addressing this important issue. More in-depth research should be conducted on factors that contribute to UIT and various contributing factors to human error. Based on the results, a framework could be developed for automatic modeling and measurement of contributing factors of the UIT threat vectors, and proactively responding to the threats. This could be an unobtrusive, non-invasive multi-modal measurement framework which is able to assess users' cognitive and affective state using only inexpensive standard computer devices (e.g. keyboard, mouse, touchpad).

7. Conclusion and Future Work

Although accidental/unintentional incidents account for the majority of all insider threats, unintentional insider threat (UIT) as a research topic is largely unexplored. We

presented a novel approach based on eye movement patterns for predicting human error which is a major factor in unintentional insider threat. We conducted experiments using an eye tracking device to capture users' eye movements as they performed several computer-based activities in different scenarios. Our experimental results show that our approach achieved an average accuracy up to 82% average accuracy for users wearing eyeglasses and 86.24% for users without eye glasses. In conclusion, our results demonstrate that eye movements and pupil behaviors can reveal valuable knowledge about the user and can be used as an effective indicator in designing unintentional insider threat detection and mitigation frameworks. For future work, we plan to integrate standard computer devices (e.g. keyboard, mouse) into our framework and examine if a multi-modal framework improves the accuracy of human error prediction.

References

[1] R. Baloh, A. Sills, W. Kumley, and V. Honrubia. Quantitative measurement of saccade amplitude, duration, and velocity. *Neurology*, 25(11):1065–1065, 1975.

[2] V. Cantoni, C. Galdi, M. Nappi, M. Porta, and D. Riccio. Gant: Gaze analysis technique for human identification. *Pattern Recognition*, 48(4):1027–1038, 2015.

[3] A. De Luca, M. Denzel, and H. Hussmann. Look into my eyes!: Can you guess my password? In *Proceedings of the 5th Symposium on Usable Privacy and Security*, page 7. ACM, 2009.

[4] S. Eberz, K. B. Rasmussen, V. Lenders, and I. Martinovic. Preventing lunchtime attacks: Fighting insider threats with eye movement biometrics. In *Proceedings 2015 Network and Distributed System Security Symposium (NDSS)*, 2015.

[5] Forcepoint LLC. The cost of an unintentional insider threat. http://www.accudatasystems.com/,last view Feb 2017, 2016.

[6] J. Gómez-Poveda and E. Gaudioso. Evaluation of temporal stability of eye tracking algorithms using webcams. *Expert Systems with Applications*, 64:69–83, 2016.

[7] F. L. Greitzer, J. Strozer, S. Cohen, J. Bergey, J. Cowley, A. Moore, and D. Mundie. Unintentional insider threat: contributing factors, observables, and mitigation strategies. In *System Sciences (HICSS), 2014 47th Hawaii International Conference on*, pages 2025–2034. IEEE, 2014.

[8] F. L. Greitzer, J. R. Strozer, S. Cohen, A. P. Moore, D. Mundie, and J. Cowley. Analysis of unintentional insider threats deriving from social engineering exploits. In *Security and Privacy Workshops (SPW), 2014 IEEE*, pages 236–250. IEEE, 2014.

[9] S. G. Hart and C. D. Wickens. Workload assessment and prediction. In *Manprint*, pages 257–296. Springer, 1990.

[10] Y. Hashem, H. Takabi, and R. Dantu. Insider threat detection based on users' mouse movements and keystrokes behavior. In *Proceedings of the Secure Knowledge Management Conference 2017 (SKM 2017), Tampa, FL, USA*, 2017.

[11] Y. Hashem, H. Takabi, R. Dantu, and R. Nielsen. A multi-modal neuro-physiological study of malicious insider threats. In *Proceedings of the 9th ACM CCS International Workshop on Managing Insider Security Threats*, MIST '17. ACM, 2017.

[12] 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*, pages 71–74. ACM, 2015.

[13] Y. Hashem, H. Takabi, M. GhasemiGol, and R. Dantu. Inside the mind of the insider: Towards insider threat detection using psychophysiological signals. *Journal of Internet Services and Information Security (JISIS)*, 6(1):20–36, 2016.

[14] T. Kinnunen, F. Sedlak, and R. Bednarik. Towards task-independent person authentication using eye movement signals. In *Proceedings of the 2010 Symposium on Eye-Tracking Research & Applications*, pages 187–190. ACM, 2010.

[15] O. Mazhar, T. A. Shah, M. A. Khan, and S. Tehami. A real-time webcam based eye ball tracking system using matlab. In *Design and Technology in Electronic Packaging (SIITME), 2015 IEEE 21st International Symposium for*, pages 139–142. IEEE, 2015.

[16] 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*, pages 479–491. ACM, 2015.

[17] S. J. Pan and Q. Yang. A survey on transfer learning. *IEEE Transactions on knowledge and data engineering*, 22(10):1345–1359, 2010.

[18] C. P. Pfleeger. Reflections on the insider threat. In *Insider attack and cyber security*, pages 5–16. Springer, 2008.

[19] D. J. Pond and K. R. Leifheit. End of an error. *Security Management*, 47(5):113–117, 2003.

[20] R. Richardson and C. Director. Csi computer crime and security survey. *Computer Security Institute*, 1:1–30, 2008.

[21] I. Rigas, G. Economou, and S. Fotopoulos. Human eye movements as a trait for biometrical identification. In *Biometrics: Theory, Applications and Systems (BTAS), 2012 IEEE Fifth International Conference on*, pages 217–222. IEEE, 2012.

[22] V. Schöpf, T. Schlegl, A. Jakab, G. Kasprian, R. Woitek, D. Prayer, and G. Langs. The relationship between eye movement and vision develops before birth. *Frontiers in human neuroscience*, 8, 2014.

[23] C. I. T. Team. Unintentional insider threats: A foundational study. *Software Engineering Institute Technical Report*, 2013.

[24] C. I. T. Team. Unintentional insider threats: social engineering. *Software Engineering Institute*, 2014.

[25] Tobii pro x2-60 eye tracker, 2016.

[26] Verizon Enterprise. Verizon data breach investigations reports. http://www.verizonenterprise.com/verizon-insights-lab/dbir/,last view Feb 2017, 2016.