Continuous and Transparent Authentication of Haptic Users

Fatimah Elsayed[1], Kiran Balagani[1], Paolo Gasti[1], Chung Hyuk Park[2], and Anand Santhanakrishnan[1]

***Abstract*— Telerobotic systems are used to perform critical tasks in sensitive environments. The security of these systems is of paramount importance, because compromising them can result in significant harm. In this paper, we attempt to address threats leading to illegitimate access to telerobotic devices. We conducted an experiment in which users explored a scene using a GeoMagic Phantom Omni haptic device. The scene provided only limited visual feedback, and required users to interact with it by primarily relying on haptic feedback. We recorded how 32 users interacted with the haptic device over a total of 180 sessions. Our results show that haptic signals collected during a session can be successfully used to distinguish between users. As a result, telerobotic operators can be authenticated transparently throughout a session (i.e., *continuously*) by relying on haptic measurements alone.**

I. INTRODUCTION

In recent years, telerobotic systems have been increasingly used to perform critical tasks in sensitive environments. These systems enable greater accuracy and/or safety while users perform critical and time-sensitive tasks, such as surgeries [4] and search-and-rescue operations [10]. Security of telerobotic systems is, therefore, very important: compromising any of the components of these systems can result in significant harm, including loss of human lives.

Current research on telerobotic security has focused primarily on system- and network-level threats [2], [8]. While these are certainly important to telerobotic security, they do not address threats leading to illegitimate access through stolen or spoofed user credentials, impersonators, and insiders [7].

In this paper we address this problem by evaluating user behavioral signatures under realistic operational constraints. Our system leverages unique user behavioral signals (e.g., certain aspects of velocity, acceleration, and pressure), collected while the user operates the haptic device, to perform authentication *continuously* and *transparently*. Because our system authenticates users based on their behavior, rather than using something they know (e.g., a password), it is substantially more resilient to impersonators, insiders, and adversaries with access to spoofed credentials.

To evaluate our system, we collected data from 32 subjects using a custom haptic setup. We performed experiments in which users perceived and recognized several static and dynamic object geometries, while the system recorded users' haptic inputs. Our results show that haptic signals carry enough information to distinguish between subjects. These signals can therefore be used for continuous authentication.

The research presented in this paper is unique because we aim to determine the biometric individuality of several classes of user behavioral signatures under realistic operational constraints. This allows us to determine the inherent individuality of the signatures, devoid of the artificial discriminability caused by operational constraints such as reduced visual or haptic feedback. Though operational constraints are commonly encountered by telerobotic users, to our knowledge their impact on biometric signatures has not been investigated.

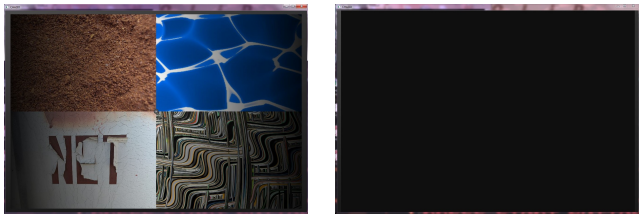

(a) (b)

Fig. 1. Figure (a) shows a visual representation of the texture presented to the user. Figure (b) depicts the user's obstructed view of the scene.

II. RELATED WORK

There is very limited research on behavioral authentication of telerobotic users. Current studies performed evaluation on a small number of subjects (12 in [1], 14 in [9], which was extended to 22 in [6]). The resulting error rates are therefore of questionable statistical value. Moreover, results from [6], [9] were obtained through experiments that may have been influenced by contextual factors. Specifically, because each subject was asked to sign a virtual check, it is difficult to tell whether the reported discriminability was due contextual artifacts (e.g., differences in the subjects signatures), or due to individual user characteristics.

III. EXPERIMENT SETUP

We performed experiments using a data collection application that we developed using the Chai3D open-source simulator [3] and the Open Haptics library [5]. In our experiments, users were presented with a scene composed of four textures of various roughness and complexity. Each texture covered an equal portion of a plane, divided in four quadrants. Users were asked to rank the textures from the roughest to the smoothest by relying exclusively on haptic feedback from a Geomagic Phantom Touch Omni. Visual feedback was concealed by covering the texture with a black veil with no haptic feedback. Figure 1 shows a sample scene used in our experiments. Figure 1(a) depicts a visual representation of the scene, where darker colors represent deeper parts of the plane. The users' view of the scene is represented in Figure 1(b). With each experiment, the position of the textures was randomly assigned.

[1] School of Eng. and Computing Science, New York Inst. of Technology
[2] Dept. of Biomedical Engineering, The George Washington University

We collected data from 32 users (9 female, 23 male), primarily from the NYIT student and faculty population, over of 180 sessions (5.6 sessions per user, on average; up to 12 sessions per user).[1] Five subjects were left-handed, and 27 were right-handed.

While the user interacted with the scene, we extracted two high-level events: *probes*, and *strokes*. A user generates a probe event when she touches an object briefly using the haptic pen. In our experiments, we defined probes as events in which the tip of the haptic pen moved 1 cm or less while touching the surface of the object. Strokes are events in which users move along the surface of an object while touching it. For instance, we consider a single uninterrupted line drawn on a surface as a stroke event. In our experiments, we defined strokes as events in which the tip of the haptic pen moved by at least 1 cm while continuously touching the surface of an object within the boundaries of one of the four textures. Specifically, if the user started a stroke on one of the texture, and then completed it on another texture, we considered the signals generated on each texture as a separate stroke.

In our experiments, we collected 2684 strokes and 520 probes. This discrepancy is due to the fact that strokes provide significantly more information to the user about the texture being explored, and as such they were typically preferred by users. Due to large quantity of strokes collected, and their richness in terms of user-spec characteristics, our analysis focused exclusively on this type of event. To analyze the signals resulting from strokes, we extracted the following features: average stroke velocity, average stroke angular velocity, average stroke pressure, stroke length, stroke duration, stroke start position, and stroke end position. Velocities, pressure, and positions were recorded on three axes (x, y, and z), while lengths and durations were represented as scalar values. This resulted in 17 features. Each feature was further characterized by the quadrant in which it was collected, thus resulting in a total of $17 \cdot 4 = 68$ features.

To test the quality of these features, we followed standard authentication procedures. We first built user behavioral profiles (*training vectors*) by computing feature-wise averages from multiple strokes. We calculated *testing vectors* by computing feature-wise averages from multiple strokes collected in sessions other than the ones used for training. We then computed the Scaled Manhattan distance, defined as $D(X,Y) = 1/n \sum_{i=1}^{n} (|x_i - y_i|)/\sigma_i$, between all training vectors and all testing vectors (x_i represents the i-th component of vector X, while σ_i is the standard deviation of feature i).

We calculated false rejects by computing, for various thresholds τ, whether $D(X,Y) < \tau$ with X and Y obtained from different sessions for the same users. Similarly, we calculated false accepts by computing whether $D(X,Y) > \tau$ with X and Y from different users. Finally, we computed Equal Error Rates (EERs) by identifying the rate of false accepts for the value of τ for which the rate of false accepts was the same as the rate of false rejects.

[1]IRB approval was obtained prior to performing the experiments under NYIT IRB protocol BHS-1003.

IV. RESULTS

None of the users in our experiments had previous experience with haptic devices. For this reason, as user were able to practice with the haptic device used in our experiments, their usage patterns changed visibly. To evaluate this phenomenon, we collected data during up to 12 sessions from each user. (For each user, different session were performed on different days.) Our results, summarized in Table I, show that training the classifier using later sessions (sessions 9-10) led to better results compared to using earlier sessions (e.g., sessions 3-6).

TABLE I

EQUAL ERROR RATES OBTAINED USING THE SCALED MANHATTAN VERIFIER (LOWER VALUES INDICATE BETTER RESULTS).

Exp. #	Training sessions	Testing sessions	EER
1	Sessions 3, 4	Sessions 5-12	46.22%
2	Sessions 5, 6	Sessions 3, 4, 7-12	46.49%
3	Sessions 7, 8	Sessions 3-6, 9-12	48.33%
4	Sessions 9, 10	Sessions 3-8	**43.08%**

Our results show that haptic signals collected during probes and stroke events can be successfully used to distinguish between users. We consider this work a first step towards a comprehensive understanding of the dynamics of continuous authentication via haptic devices. As such, future research directions include the exploration of feature selection and feature-level fusion mechanisms in order to maximize the discriminative power of the features being captured. Finally, we consider the exploration of additional verification techniques, such as SVM, Random Forest, and Gaussian Mixture Models, as an avenue for further reducing authentication error rates.

Acknowledgements. This work was supported in part by an NYIT Institutional Support of Research and Creativity (ISRC) grant. Gasti and Balagani were supported by NSF grant CNS-1619023.

REFERENCES

[1] A. Bianchi. Haptic and audio authentication: Empirically exploring usability, security and feaibility of non-visual passwords. *PhD Dissertation, KAIST*, 2012.

[2] T. Bonaci and H. J. Chizeck. Surgical telerobotics meets information security. In *Robotics, Science and Systems (RSS) Workshop on Algorithmic Frontiers in Medical Robotics: Manipulation in Uncertain, Deformable, Heterogeneous Environments*, July 2012.

[3] Chai3d - opensource haptic simulation toolkit. `http://www.chai3d.org`. Accessed: 2018-01-24.

[4] da vinci surgical system. `https://www.cancercenter.com/treatments/da-vinci/`. Accessed: 2018-01-24.

[5] GeoMagic. Open haptics toolkit programmers guide, 2014.

[6] M. Orozco, M. Graydon, S. Shirmohammadi, and A. E. Saddik. Experiments in haptic-based authentication of humans. *Springer Science Jl. on Multimedia Tools and Applications*, 37, Apr. 2008.

[7] M. R. Randazzo, M. Keeney, E. Kowalski, D. M. Cappelli, and A. P. Moore. Insider threat study: Illicit cyber activity in the banking and finance sector. *CERT Coordination Center, Software Engineering Institute, Carnegie Mellon University (PA, USA)*, 2004.

[8] Researchers hijack teleoperated surgical robot: Remote surgery hacking threats. `http://www.computerworld.com/article/2914741`. Accessed: 2018-01-24.

[9] A. E. Saddik, M. Orozco, Y. Afshaw, S. Shirmohammadi, and A. Adler. A novel biometric system for identification and verification of haptic users. *IEEE Trans. on Instrumentation and Measurements*, 56, 2007.

[10] Toshiba readies scorpion-like robot for fukushima nuclear plant. `http://www.computerworld.com/article/2942037/`. Accessed: 2018-01-24.