Assentication: User De-authentication and Lunchtime Attack Mitigation with Seated Posture Biometric

Tyler Kaczmarek(✉), Ercan Ozturk, and Gene Tsudik

UC Irvine, Irvine, USA
{tkaczmar,ercano,gtsudik}@uci.edu

Abstract. Biometric techniques are often used as an extra security factor in authenticating human users. Numerous biometrics have been proposed and evaluated, each with its own set of benefits and pitfalls. Static biometrics (such as fingerprints) are geared for discrete operation, to identify users, which typically involves some user burden. Meanwhile, behavioral biometrics (such as keystroke dynamics) are well-suited for continuous and more unobtrusive operation. One important application domain for biometrics is *de-authentication*: a means of quickly detecting absence of a previously-authenticated user and immediately terminating that user's secure sessions. De-authentication is crucial for mitigating so-called *Lunchtime Attacks*, whereby an insider adversary takes over an authenticated state of a careless user who leaves her computer.

Motivated primarily by the need for an unobtrusive and continuous biometric to support effective de-authentication, we introduce Assentication – a new hybrid biometric based on a human user's seated posture pattern. Assentication captures a unique combination of physiological and behavioral traits. We describe a low-cost fully functioning prototype that involves an office chair instrumented with 16 tiny pressure sensors. We also explore (via user experiments) how Assentication can be used in a typical workplace to provide continuous authentication (and de-authentication) of users. We experimentally assess viability of Assentication in terms of uniqueness by collecting and evaluating posture patterns of a cohort of 30 users. Results show that Assentication yields very low false accept and false reject rates. In particular, users can be identified with 94.2% and 91.2% accuracy using 16 and 10 sensors, respectively.

1 Introduction and Motivation

Secure, correct and efficient user authentication is an integral component of any meaningful security system. Authentication schemes implemented in a typical modern workplace typically include two factors: (1) a user demonstrates knowledge of a secret password or PIN, and (2) a user proves possession of a secure device or token. However, it is becoming more popular to augment this approach with a third factor – biometrics that reflect inherent human traits or behaviors. Biometric techniques are considered as the best means of evaluating

B. Preneel and F. Vercauteren (Eds.): ACNS 2018, LNCS 10892, pp. 616–633, 2018.
https://doi.org/10.1007/978-3-319-93387-0_32

human inherence and they range widely: from a very simple (e.g., fingerprints) to rather complex, such as iris scans.

After initial authentication, users often spend long stretches of time continuously using computing devices and services. During that time, continuous presence of the originally authenticated user must be periodically re-affirmed, especially, in a shared workplace setting. Failure to do so can result in so-called *Lunchtime Attacks*. Such an attack occurs when a previously authenticated user walks away from her workplace, thus allowing the adversary to take over her login session and engage in potentially nefarious activity. This prompts the need for periodic re-authentication and/or continuous authentication. Unfortunately, the former can be quite annoying, as is the case with too-short inactivity time-outs requiring frequent password re-entry.

Meanwhile, continuous authentication (or presence verification) is challenging in its own right. For example, camera-based methods that use face recognition [1] or gaze tracking [2] might be viewed as intrusive in terms of personal privacy, since cameras can be abused (e.g., by malware) to surreptitiously record users. Furthermore, face recognition is prone to attacks, while gaze tracking requires the user to maintain line-of-sight with the camera, which can result in unnecessary de-authentication when the user turns away while remaining at the workplace. Whereas, keyboard or mouse activity profiling and monitoring, though effective in some settings, are poorly suited for cases when a user temporarily halts input activity, e.g., in order to chat with co-workers or answer the phone. Other techniques continuously measure physical distance between the user and her workplace, by requiring each user to wear an extra device, e.g., a wristband or smart badge. Such methods are: (1) potentially burdensome due to imposing an extra device, and (2) ultimately authenticate only the presence of the device and not of its owner.

Based on the above discussion, we believe that the "design space" for continuous authentication (or, equivalently, de-authentication) techniques needs to be explored further. From the outset, we acknowledge that a perfect continuous authentication method is unlikely to materialize; in fact, one might not even exist. In other words, since each previous method has a distinct set of advantages and limitations/flaws, the same will certainly hold for our current efforts.

In this paper, we propose and evaluate a new biometric called Assentication. It is based on a user's seated posture patterns in an average office chair over the course of a typical workday. We examine the applicability of Assentication for continuous user authentication, i.e., ensuring that – after the initial successful login – the person currently using a particular computer is the same as the one who initially logged in. One of Assentication's key advantages over many other de-authentication methods is its ability to operate in an unobtrusive manner, with no effort on the part of the user.[1] To evaluate its viability and effectiveness, we built a low-cost Assentication prototype by instrumenting a commodity office

[1] This is in contrast with, for example, fingerprint-based continuous authentication, which would prompt the user to periodically swipe her finger(s) on the fingerprint reader; which is obtrusive and disrupts the typical workflow.

chair with ultra-thin flexible sensors that gather user posture data. Its purpose was to assess whether users are correctly authenticated, based on their own training data. The same platform was used to test the uniqueness of Assentication within a sample population of measured users. Our results demonstrate that the prototype unobtrusively captures the necessary data for continuous authentication and identification while the user engages in a typical use of a desktop or laptop computer.

The rest of this paper is organized as follows: Sect. 2 overviews related work. Next, Sect. 3 provides the background on continuous authentication and de-authentication. Section 4 describes the Assentication biometric. Then, Sect. 5 outlines the adversarial model. Section 6 describes the Assentication prototype and methodology used for data collection, followed by results in Sect. 7, a detailed discussion of Assentication is provided in Sect. 8. The paper concludes with directions for future work in Sect. 9 and a summary in Sect. 10.

2 Related Work

Biometric traits have been extensively explored in the context of authentication. Jain et al. [3] provides an authoritative overview of many well-known techniques, including: fingerprint, face, iris, palm-print and keystroke dynamics. However, since our focus is on biometric-based continuous authentication which can be used to achieve effective de-authentication, we do not discuss methods that are not amenable for the intended application.

Eberz et al. [4] provide an overview of the state of the art of the evaluation of biometric techniques for authentication. Additionally they provide recommendations for the evaluation of future methods. Our evaluation strategy is informed by this framework.

Rasmussen et al. [5] use human body's response to electric signals as a biometric. In the proposed system, a weak pulse signal is applied to the palm of one hand and measured on the palm of the other hand. Pulse-response biometric can be used as a second or third factor in user authentication and/or as a continuous authentication mechanism. The system achieves 100% accuracy over a static set, and 88% accuracy on permanence tests performed over several weeks.

Eberz et al. [2] investigate eye movement patterns as a biometric. Based on gazing data gathered from 30 participants, pupil, temporal and spatial features are defined. Reported equal error rate is 3.98% in a single session and 92.2% of attacks are detected within 40 s. Measurements done two weeks apart show that this biometric is stable over time.

Mare et al. [6] propose wearing a bracelet that has a gyroscope and an accelerometer for continuous authentication. When the user interacts with the computer (e.g., typing or scrolling), the bracelet transfers collected sensor data to the computer, which evaluates whether user actions match the sensor data. The proposed system, ZEBRA, achieves continuous authentication with 85% accuracy and detects attacks within 11 s. However, a recent study by Huhta et al. [7] presents a set of credible attacks on ZEBRA.

Keystroke dynamics are another means of continuous authentication. Ahmed and Traore use 1,500 digraphs from each user as a base profile and applies neural networks to guess missing digraphs [8]. In a 53-user experiment, a false reject rate of 0.0152%, a false accept rate of 4.82% and an equal error rate of 2.46% are achieved.

Finally, Conti et al. [9] describe FADEWICH, a continuous authentication system that uses attenuation of wireless signals when a human body is on the signal's path. FADEWICH is **not** based on any biometrics. It tracks the user by placing 9 sensors in a 6m-by-3m office environment. Once detected as having left the environment, the user is logged out. FADEWICH successfully de-authenticates users with 90% accuracy within 4 s, and 100% accuracy within 6 s.

There have been prior attempts to use posture and seated pressure for both identification and continuous authentication. Gia et al. [10] use data gathered from: (1) four pressure sensors placed on the seat bottom, (2) an accelerometer, and (3) light sensors placed on the seat-back, to identify the user. Pressure sensors are used to differentiate among users, while weight and accelerometer readings determine chair movements when someone sits down. Light sensors help determine how much space is covered by the sitting user. In an experiment involving only 10 people, a rather low accuracy of 72% is achieved.

Furthermore, Yamada et al. [11] describe a hip-print authentication method which uses pressure data from 32 sensors placed along the seat bottom. However, in experiments that use only first 1.5 s of measurement, quite low accuracies of 74.3% and 59.9% are reported for 10 and 25 subjects, respectively.

Among prior work, the one closest to this paper is [12]. It proposes a continuous driver identification system for automobiles that uses pressure data from two mats, each containing 32×32 sensors, placed on the seat cushion and backrest of the driver seat. Features used for classification are based on one's pelvic bone signature, mid- to high-pressure distribution and weight. In a study involving 34 participants, a fairly low uniqueness rate is reported. The authors assert that this is because a car setting is not appropriate for detecting pressure distribution changes. Most drivers adopt a single, constant posture adjusted to their preferred driving position.

Finally, Mutlu et al. [13] investigate how to use fewer sensors to detect posture. To determine optimal sensor placement, a classifier is constructed that learns the probabilistic model between the chosen subset of sensor values and feature vectors used for posture classification. With 19 sensors, classification accuracy of 87% is reported. As discussed later, this study guides our sensor placement strategy.

3 Background

This section sets the stage for the rest of the paper by overviewing user authentication, de-authentication, attack scenarios and continuous authentication requirements.

3.1 User Authentication

User authentication can involve one or more of the following factors:

F1: What one knows, or what one recognizes.
The former corresponds to knowledge of: passwords, PINs, drawing patterns and free-text answers to security questions. The latter corresponds to recognition of: correct answers to multiple-choice questions, faces or other types of images.

F2: What one has in their possession.
This generally means some form of a personal (even passive) device, such as a badge, bracelet, key-fob, token or smartphone.

F3: What one is, or how one behaves.
The former type is referred to as a *static* and includes biometrics based on: fingerprints, irises, palms, wrists, faces, ears and pulse-response. The latter type is called *behavioral* and includes biometrics based on: gait, keystroke dynamics, head movements, hand gestures and gaze tracking.

Though widely used, F1-type authentication alone is widely considered to be insufficient, mainly due to the low entropy of secrets involved. By itself, F2 is also inadequate, since a personal device is not guaranteed to always be in possession of its intended owner. Finally, F3 can be subverted, at least for some static methods, e.g., via cloned fingerprint moulds [14], fake irises using contact lenses [15], and face masks[2]. It also usually requires a non-trivial training or enrollment phase. Meanwhile, some behavioral biometrics are unstable or fragile, e.g., gait, breathing patterns, blinking and head movements. Consequently, multi-factor user authentication is usually recommended in order to achieve better security.

3.2 De-authentication and Lunchtime Attacks

As part of everyday office or workplace activity, an average user might engage in one or more of the following activities (not an exhaustive list):

[A1]: Work by continuously utilizing one or more traditional input devices, such as a keyboard, touchscreen or mouse.
[A2]: Take a quick seated nap or meditation break.
[A3]: Read some printed matter, e.g., a paper or book.
[A4]: Use another personal device, e.g., a smartphone.
[A5]: Turn away from one's desk to talk to other people directly, or on the phone.
[A6]: Watch videos and/or listen to music without using any input devices.
[A7]: Take part in an audio or video conference.
[A8]: Get up momentarily to fetch something from the immediate vicinity (or simply to stretch) and return.

[2] "Biometric Update", "spoofing iris recognition technology with pictures", http://www.biometricupdate.com/201503/spoofing-iris-recognition-technology-with-pictures, 2015, accessed: 2017-05-19.

[A9]: Walk away from the workplace for a short (e.g., bathroom), longer (e.g., lunch), or long (e.g., done for the day) time, before returning.

In a security-conscious setting, these activities might require periodic reassurance that the same user (who initially authenticated and/or logged in) is still present. Ideally, when the original user remains present [A1–A7], no reassurance should be needed. However, [A9] results in leaving the workplace unattended, while [A8] might. (Also, [A2] could be viewed as the user not really being there.) An important challenge is to distinguish among these types of activities. The term *de-authentication* denotes the process of deciding whether the original user is no longer present and, if so, terminating active secure sessions.

In a perfect world, each user would always log out or otherwise terminate all active sessions before stepping away. Unfortunately, this is far from reality, which triggers the threat of *Lunchtime Attacks*. As the name suggests, attack of this type occurs when the adversary takes over the secure session(s) of a legitimate user who has left, even for a short time. Such attacks are quite common, as noted in the recent work of Marques et al. [16].

3.3 Default Approach: Inactivity Timeouts

The most common current means of dealing with Lunchtime Attacks and reassuring original user presence is inactivity timeouts. Most users of personal and workplace computing devices are familiar with them: whenever keyboard and/or mouse inactivity exceeds a certain threshold, de-authentication takes place, i.e., log-in and other (previously authenticated) sessions are terminated. Various operating systems, apps and websites set their own timeout rules and policies. In some cases (e.g., macOS or Windows) users can select their own timeouts. At a typical workplace, mandatory timeouts are often imposed.

Inactivity timeouts are almost universally disliked. As noted in [6], most users find too-short timeouts annoying, while too-long timeouts are insecure, since they defeat the purpose of Lunchtime Attack mitigation (by extending the attack time window). Even more importantly, timeouts achieve their desired effect only in case [A1] and fail in several other ways:

- They operate under the assumption that keyboard/mouse inactivity (i.e., "NOT [A1]") indicates user absence. This is often not true, e.g., in cases [A2]–[A5] and [A8]. De-authenticating the user in these cases is both unnecessary and annoying.
- Conversely, timeouts naïvely suppose that resumption of activity (within the timeout threshold) indicates presence of *the same* user. This is clearly wrong in situations where the original walks away [A9] and the adversary quickly starts typing.
- In case [A6], if timeouts are activated, the user is also unnecessarily burdened. Otherwise, if timeouts are automatically disabled while music and/or videos are playing, the user can walk away for a potentially long time, thus leaving the computing device(s) open to Lunchtime Attacks.

- The same holds for case [A7], except that user's voice and/or camera movements might be used to infer continuous presence. However, this would require additional voice or visual authentication.
- In case [A9], timeouts only work correctly (by de-authenticating the original user) if no attack occurs. Knowing the timeout threshold, which is usually not secret, allows the adversary to easily succeed in a Lunchtime Attack.

3.4 Continuous Authentication

Given the inadequacy of inactivity timeouts, one appealing alternative is continuous authentication. Methods of this variety are generally unobtrusive, i.e., require none or very little user burden. As discussed in Sect. 2, these include: keystroke dynamics [8], wrist movement [6], pulse response [5], gaze tracking [2], and wireless signal monitoring [9]. (Note that only the first four are biometric-based methods, while the last is purely a de-authentication technique.)

3.5 Design Goals

Since our main goal is the design of a biometric-based de-authentication method, we first consider general design goals for biometrics. A popular survey of biometric techniques by Jain et al. [3] provides a comprehensive overview of many popular methods, and discusses design criteria, which include the following:

Universality: The biometric must be (ideally) universally applicable. For example, an iris scanner is not useful for users who are missing an eye or have cataracts, while fingerprint readers are similarly useless for people with severe eczema.

Uniqueness: The biometric must be unique within the target population. It must be possible to distinguish users using the biometric.

Unobtrusiveness: The biometric should be maximally transparent. Ideally, it should be used in a passive manner, without any extra requirements or interference with users' normal behavior.

Circumvention Difficulty: To be meaningful in any security context, the biometric must be difficult to circumvent. That is, false accept (fraud) rate (FAR) should be minimal, i.e., it should be hard to impersonate a genuine user.

Low Error Rate: The biometric must have a low false reject (insult) rate (FRR), i.e., should very rarely fail to recognize an enrolled user.

Collectability: The biometric should be measurable in a fast, easy and meaningful quantitative way.

Cost Effectiveness: The biometric's distinguishing power as related to the cost of deployment and maintenance. In our design, this is a key goal.

Easy enrollment: The biometric's initial (training) phase should be as short and burden-free as possible.

Acceptability: The ideal biometric is one which (most) users are comfortable to use.

We now present design goals for an ideal de-authentication method, not necessarily based on biometrics.[3]

- Minimal extra components (particularly, physical or hardware) and monetary cost
- Quick and correct detection of activities requiring de-authentication, i.e., [A9] or a circumvention attempt, e.g., another user sits down
- Minimal False Reject Rate (FRR), i.e., probability of mistaking [A1]–[A8] for [A9]
- Minimal False Accept Rate (FAR), i.e., probability of mistaking [A9] for [A1]–[A8]
- Maximal user transparency, i.e., unobtrusiveness

We recognize that the last goal might be ethically dubious. De-authentication methods with user transparency can be abused, e.g., by unscrupulous employers, to surreptitiously spy on unsuspecting users. We acknowledge that it is very difficult to reconcile positive and negative connotations.

4 Assentication Biometric

Assentication works by monitoring, over time, changing pressure patterns exerted by a person seated in a typical office chair. This pattern is influenced by both behavioral and physical characteristics. The former stem from one's seating preferences. For example, some people cross their legs, which leads to an asymmetric pressure distribution, while others keep both feet firmly on the ground which results in nearly symmetric pressure distribution. Other contributing factors include height, hip width and weight.

4.1 Strengths and Weaknesses

Since exact distribution of seated pressure depends on the user's physical dimensions as well as on adopted postures, Assentication is a hybrid biometric blending physiological and behavioral factors. This allows it to benefit from some strengths of both. In particular, one's posture pattern can be captured in a strictly passive manner. Even though this property is shared by other biometrics, such as facial recognition or pulse response, posture pattern is not easily circumventable (unlike, e.g., facial recognition), and does not alter normal user behavior, unlike, e.g., pulse-response. We believe that this combination of unobtrusiveness, difficulty of circumvention, and behavior agnosticism make Assentication an attractive biometric.

Additionally, Assentication requires very little in terms of specialized hardware to capture the physiological biometric it uses. As discussed later in Sect. 6, we constructed a Assentication prototype of an instrumented office chair.

[3] We do this while keeping in mind that all of them are unlikely to be achievable.

4.2 Liveness and Replay

In any biometric system used for continuous authentication, liveness detection is a serious concern. For example, a face recognition system needs to detect blinking, breathing, and/or some other artifact of a user being alive and present. Otherwise, as has been demonstrated in the past, it can be subverted by a photo or a mask (face-cast). Traditionally, liveness is attained via some form of a challenge by the system that requires the user to act. In case of facial recognition, the system might prompt the user to turn her head or look in a particular direction. While this helps achieve liveness and protect against subversion, it also sacrifices transparency and increases user burden.

Some modern de-authentication systems, such as gaze tracking or keystroke patterns, can passively check for liveness by relying on dynamic user behavior instead of constant physical characteristics. However, they require the user to act in a particular (not necessarily free or natural) manner. For example, gaze tracking requires the user to face in the general direction of the gaze tracking apparatus, which may not always be in the user's typical workflow. Furthermore, gaze tracking requires the user's eyes to be open. In the same vein, keystroke analysis requires the user to type on the keyboard. For its part, the pulse-response biometric needs the user to complete an electrical circuit by touching conductive implements with both hands. With all these systems, if the user fails to behave in the required manner, the likely outcome is a false accept.

In contrast, Assentication is more forgiving in such cases. It does not rely on specific user actions. Instead, Assentication is based merely on user's physical presence. It monitors seated pressure distribution regardless of whether the user faces the workstation, touches the keyboard with both hands, is currently typing, or keeps their eyes open. The only requirement for collection of posture pattern data is that the user must be seated in the chair. We believe that this makes our system a good candidate for both continuous authentication and de-authentication.

5 Adversarial Model and Attacks

The Assentication biometric focuses on protecting against insider threats. We are particularly concerned with aforementioned Lunchtime Attacks whereby the adversary steps in to access a co-worker's computer after the latter walks away. Insider threats are not limited to such attacks, and might include scenarios ranging from a disgruntled employee staying after hours to sabotage a colleague, to the trivial case of a user deliberately giving access to a co-worker. In all scenarios, the adversary "wins" by gaining access to secure log-in or application sessions.

We assume that the original user provides authentic log-in credentials at session initiation time. However, the same user neglects to log-out before physically leaving the workplace. Once the original user leaves, the adversary approaches the computer, accesses secure log-in sessions and performs some actions, e.g.,

copy or erase sensitive files, read or send private email. Such attacks are particularly dangerous since they originate from valid and logged-in user accounts. Also, it might be very difficult for the victim to repudiate the adversary's actions.

Insider attacks are unfortunately quite commonplace. In fact, they account for about 28% of all electronic crimes in industry [17]. This includes some high-publicity attacks, such as the infamous 2014 Sony hack [18].

We consider two types of insider adversaries: *casual* and *determined*. In both cases, the adversary is aware of Assentication's use and presence. The adversary is considered successful if it manages to circumvent the system, either by physically imitating the victim's pressure patterns, or by constructing an accurate model (replica) that does the same. We assume that the adversary cannot disable the system, or interfere with its correct operation through physical sabotage, since such manipulation would leave traces.

The *casual* adversary aims to subvert Assentication through behavioral imitation of the victim's posture patterns. We assume that this adversary is familiar with the habits and schedule of the victim, and has physical access to the victim's workplace. Success of the *casual* adversary relies on the discriminating power of the system. In our prototype design (discussed later), posture pattern data is aggregated and evaluated against the previously constructed profile every 10 s. Even in the unrealistically ideal scenario where the *casual* adversary instantly appears in the victim's chair immediately after the victim walks away, only 10 s would remain to perform any attack. However, in our experimental office setting, this attack time window is substantially shorter, ≈2-to-4 s, since it takes 3–4 s for the victim to leave and about as long for the adversary to enter and sit down. After that, posture data is flagged as incorrect, the victim is de-authenticated and all active secure sessions are terminated.

The *determined* adversary seeks to defeat the system by fabricating a physical model of the victim user. We assume that this adversary has access to the exact sensor data of the victim, as well as precise measurements of the victim's posterior and lower back. This data might be obtained if the adversary manages to previously trick the victim to sit (for a sufficiently long period) in a staged chair instrumented the same way as the victim's.

A perfect mold or cast of the victim with the correct pressure distribution would circumvent Assentication. However, creation and deployment of such a mold is not trivial. The determined adversary would have to create a bulky and heavy object that accurately replicates the victim's posterior as well as lower back and weighs enough to exert the necessary pressure upon the instrumented chair, in the right places. Physically and logistically, deploying the mold onto the victim's chair is burdensome and likely to be detected by extraneous means.

However, we recognize that a mold is not the only way to subvert Assentication. We conjecture that a more effective and discrete approach is to use a set of strategically placed hydraulic or pneumatic contraptions, each calibrated to exert an accurate amount of pressure on each sensor on the victim's chair. This kind of precision is difficult to achieve and, unlike a monolithic mold, placing the entire set of contraptions onto the chair at the same time is also quite hard.

Fig. 1. Assentication prototype chair: (a) as seen by the user, and (b) uncovered seat-bottom sensor placements.

All in all, we consider this attack to be quite improbable and close to the realm of "Mission: Impossible".

6 Methodology

This section describes our Assentication prototype design, experimental setup, procedures, subject parameters as well as classifiers used for data analysis.

6.1 Prototype Design

To demonstrate viability and facilitate ease of experimentation, we built the Assentication prototype by modifying a standard inexpensive office chair with commodity (off-the-shelf) sensor components. Figure 1(a) shows the prototype chair, and Fig. 1(b) focuses on the placement of sensors across the seat and back of the chair. Our sensor placement was guided by the experience in Mutlu et al. [13].

The prototype consists of three components:

1. One 2003/2004 Hon Mid-Back Task Chair.[4]
2. Sixteen (16) Tekscan Flexiforce A401 Large Force Sensing Resistors.
3. Two Arduino 101 modules[5], one of which is connected to 6 A401 resistors. The other module is connected to the remaining 10 sensors in a similar configuration, augmented with an analog multiplexer in its 6-th analog port, in order to support the use of 10 sensor inputs.

Acquired measurements are sent from the Arduino to a commodity desktop PC for collection and evaluation. Arduinos are connected to the desktop via USB cables. Obviously, in a real office setting, having wires running between the chair

[4] See https://www.hon.com.
[5] See https://www.arduino.cc/en/Main/ArduinoBoard101.

and the computer would be highly undesirable. We expect that either Bluetooth or WiFi would be used instead.

Total instrumentation cost of $275 was incurred for the initial single-chair prototype, for feasibility and testing purposes. For a medium-size office with 50 chairs, the per-chair cost can be cut significantly, to approximately $150 due to volume pricing of Tekscan A401.

6.2 Data Collection Procedure

To collect data in a realistic setting, rather than bringing subjects to an unfamiliar office and encountering complications cited by Yamada et al. [11] in collecting posture data in a lab setting, we brought the prototype instrumented chair to the subjects' workplace. Each subject was briefed on the nature of the experiment, and was asked to sit naturally. Subjects allowed us to swap out their office chair with the prototype, and continued their normal work activities while sitting on the latter. We collected posture data in rounds of 10 min per subject. 17 subjects participated in two collection sessions over the course of several days and 13 subjects participated in a single session only. We sampled subjects in order to accommodate typical day-to-day fluctuations in mood and posture, e.g., one session in the morning, and the other – shortly after lunch, on a different day.

A total of 30 subjects were recruited primarily from the student population of a large public university. Because of this, overwhelming majority (27 out of 30) were between the ages of 22 and 30, while the remaining 3 were somewhat older faculty and staff. The gender break-down was: 10 female and 20 male.

Finally, despite its somewhat ungainly appearance (as shown in Fig. 1), the prototype chair is rather comfortable for sitting and none of the subjects expressed any unease or discomfort during the data collection phase.

6.3 Features

We collected data in the form of $1,200$ sample time-series reflecting the force exerted on each of the 16 pressure sensors captured each 0.5 s over a 10-min session, for a total of $19,200$ samples per subject, per session.

For continuous authentication, we treat the first 5 min of each session as a training phase, and evaluate the subject on the next 3 frames of sensor data, representing 1.5 s of measurement. If this data is consistent with the training set, it is accepted as valid and included in the training set for future evaluations. If the data is deemed inconsistent with the training set, it is marked as adversarial and de-authentication takes place.

6.4 Feature Selection and Quality

Riener and Ferscha [12] use the highest pressure points on a car seat instrumented with a uniform array of $1,024$ pressure sensors to determine a subject's gender. Based on these results, we constructed a Random Forest (RF) classifier

Table 1. RF-based gender classification results.

Gender	TPR	FPR
Female	96.5%	1.4%
Male	98.6%	3.5%

Table 2. Pressure sensors ranked according to information gain.

Sensor	Inf. gain	Sensor	Inf. gain
15	2.4054	1	1.3519
4	1.9285	5	1.2846
12	1.8822	0	1.063
3	1.7893	8	1.034
13	1.7332	14	1.021
2	1.6246	6	0.5634
7	1.4498	10	0.3578
11	1.4407	9	0.0128

to identify subjects' gender in order to assess the quality of gender information of our features. We obtained a 97.9% True Positive Rate (TPR), and a 2.7% False Positive Rate (FPR), as shown in Table 1. This demonstrates that our strategy of utilizing 16 well-placed sensors preserves gender information reflected by pressure data. However, preliminary stratification of subjects by gender before classification did not yield a higher true acceptance rate, and we do not utilize such an approach.

Table 2 shows the ranking of our extracted features based on their information gain for all 16 sensors. The Information gain is measured with respect to the class as follows:

$$InformationGain(Class, Feature) = H(Class) - H(Class|Feature)$$

where H is entropy function and $H(A|B)$ is entropy of B conditioned on A.

6.5 Classification Algorithm

Since we are dealing with a fairly commonplace time series clustering problem, there are many well-known candidate techniques. We compared three popular classification algorithms to determine the one that yields the best results.

Random Forest (RF): we found that it consistently yields the best results. It produces precise, accurate results, closely clustered for all subjects. Both FNR and FPR are acceptably low in cross-validation of user data, as discussed in more detail below.

K-Nearest Neighbors (KNN): we tested the KNN classifier for $k = 1,\ 3$ and 5 using Euclidean Distance. KNN is a simple lazy classifier that is quite effective in many settings. However, for our classification needs, it did not perform as well as RF.

Support Vector Machine (SVM): For each subject, we trained a single binary classifier in a one-against-one case. The final prediction was determined by voting. While SVM provided extremely consistent and highly accurate results for

some users, it did not perform as well as RF, on average. It also had a few outliers with unacceptably high FPRs. We tried different kernel functions, degrees, cost and γ values.

7 Results

We present results for two classifiers: one for identification and the other – for continuous authentication. The former is based on RF and provides verification of a one-to-n match of a sample of a known user against every sample in a database. The continuous authentication classifier is based on anomaly detection in training data's inter-quartile range and provides verification of a one-to-one match of a sample of unknown origin against that of a single known user.

7.1 Identification

Identification is a classification problem across many classes. Our RF-based classifier is ideal for this – it achieves, on average, 94.2% TAR, as shown in Fig. 2. We also achieve an average FAR of 0.2%, as shown in Fig. 3 using 16 sensors. Moreover, Assentication with only 10 sensors achieves a TAR of 91.2% and FAR of 0.3%, as shown in Figs. 4 and 5, respectively.

The low FAR indicates that the *casual adversary* (as described in Sect. 5) can not successfully impersonate another enrolled user in a Lunchtime Attack with a reasonably high probability. Furthermore, the insult rate of 0.2% suggests that users experience minimal annoyance through mis-identification.

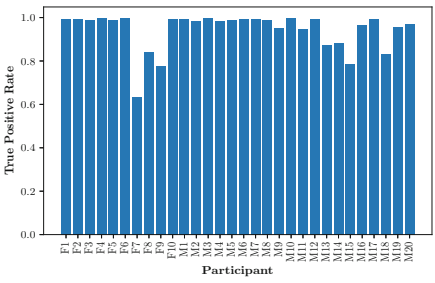

Fig. 2. Identification classifier (16 sensors)

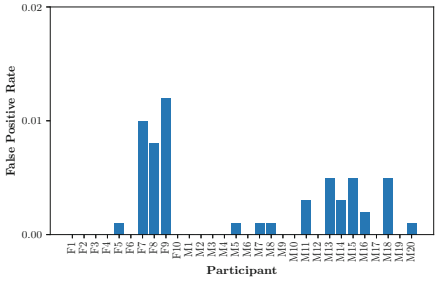

Fig. 3. FPR for RF cross-validation on all subject data (16 sensors).

7.2 Continuous Authentication

Our classifier for continuous authentication is focused on identifying outliers and extreme values in fresh incoming data. After a 5-min training window, subject training data is compared to the next 3 data slices collected by the sensors.

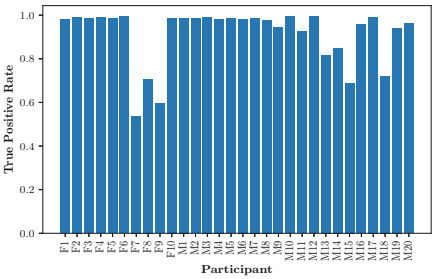

Fig. 4. RF identification TPR by subject (10 sensors).

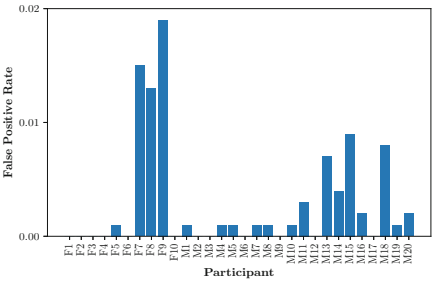

Fig. 5. RF identification FPRs by subject (10 sensors).

Each data slice is then classified as an "extreme" value (i.e. an outlier) or not. If all three measurement sets are "extreme", data is considered invalid, and de-authentication results. Otherwise, data slices are added to the training set. This comparison occurs every 1.5 s as new data is collected.

This approach results in no (0% rate) false de-authentications. Hence, a valid user who sits down on an instrumented chair will not – with very high probability – be erroneously de-authenticated while remaining seated. Additionally we find that 91% of imposters are de-authenticated in the first measurement frame (after 1.5 s) and all imposters are de-authenticated by the end of the first 45 s (30 measurement cycles). This demonstrates further resistance to *casual* adversary attacks in the continuous authentication case.

8 Discussion

We now assess Assentication in the context of design goals for an ideal de-authentication system outlined in Sect. 3.

Assentication was designed with the emphasis on *minimal extra components and monetary cost*. The use of an instrumented chair does require specialized hardware. However, it does not impose any behavioral requirements on the user. Furthermore, the per-unit cost of $150 (at scale of about 50) is reasonable in the context of other posture-based techniques, which can easily cost thousands of dollars, as noted by Mutlu et al. [13].

We claim that *maximal user transparency* is achieved by Assentication because of the ubiquitous nature of sitting in office or workplace settings. In fact, over 70% of the workforce in a traditional office setting spend upwards of 6 hours a day seated [19]. Enrollment, authentication and de-authentication phases of Assentication all occur transparently while the user is seated and engaged in normal workflow activities. Because of this, there are no behavioral modifications required from the user to participate in Assentication and no need for modifying everyday activities.

Quick detection of activities requiring de-authentication is trivial in Assentication. A user who engages in any activity covered by [A9] is de-authenticated

as soon as a single collection window passes. Though in the initial prototype implementation this window was set to 1.5 s, it can be adjusted up or down.

As evidenced by the average 94.2% accuracy of user identification and 100% accuracy for continuous authentication, false rejections would only occur in exceptional circumstances, which satisfies the *minimal insult rate* design goal. This holds during most typical office activities [A1–A7] that are typically performed while the users is seated. However, if a user leaves the chair to grab something nearby [A8] and spends over 10 s away, potentially erroneous de-authentication can occur.

8.1 Deployment Scenario

The physical and behavioral features measured as part of Assentication are somewhat ephemeral in nature. An individual's weight can fluctuate over a kilogram (2.2 pounds) day-to-day [20]. Also, one's posture is influenced by the emotive state [21]. This makes permanence of user posture rather doubtful. Therefore, we believe that an Assentication-like de-authentication system should be operated as follows:

- At the start of a session, a user sits down in an instrumented chair, and authenticates to the system normally, e.g., via username and password.
- Upon successful authentication, Assentication system collects posture data and forms a temporary profile.
- After enrollment period of 10 min, the system evaluates new posture data against the profile throughout the day.
- Once the user leaves at the end of the session (e.g., for the day), the profile is deleted.

This provides several benefits. First, no additional configuration is needed to deploy Assentication in an enviroment where several users are authorized to use the same physical terminal or workstation under different authorized accounts. Second, unlike more permanent biometrics, there is no costly institution-wide enrollment or re-enrollment required to initialize the system. Finally, the use of temporary profiles avoids the need for permanent secure storage of sensitive biometric data, which is currently an open problem for long-lived biometrics.

8.2 Ethical Considerations

All experiments described in this paper were duly authorized by the Institutional Review Board (IRB) of the authors' employer, well ahead of the commencement of the study. The level of review was: Exempt, Category II. No sensitive data was collected during the experiments and minimal identifying information was retained. In particular, no subject names, phone numbers or other personally identifying information (PII) was collected. All data was stored pseudonymously.

9 Future Work

There are several directions for future work:

First, we plan to conduct a larger-scale, longer-term (longitudinal) study, obtaining multiple measurement sessions from each subject over the course of several weeks. This would lead to a better understanding of the posture pattern biometric as a whole.

Second, we intend to evaluate accuracy of Assentication in its typical deployment scenario, as described in Sect. 8.1. For this, we intend to have the subjects replace their office chair with our prototype for an entire workday. We would use this data to obtain the rate of both false rejects and accepts, throughout the day, as well as measure associated user burden.

Next, we plan to evaluate attack vectors outlined in Sect. 5, starting with a *casual* adversary. This will entail recruiting pairs of subjects with similar physical characteristics, and training them to impersonate each other's posture patterns. Finally, we explore the attacks by a *determined* adversary. For this, we need to construct a contraption that imitates the victim's posture pattern.

10 Conclusions

In summary, this paper proposed and described a new Assentication biometric based on seated posture patterns. We built and experimented with a prototype implementation of Assentication. Furthermore, experimental results show that posture pattern biometric captures a unique combination of physiological and behavioral traits. We found that users can be identified with, on average, 94.2% accuracy from a population of 30. We also believe that it is infeasible for a *casual* adversary to circumvent Assentication by impersonation of the victim's posture patterns. We also argue that physical and logistical burdens of fabricating and deploying an accurate mold (replica) of the victim's relevant body parts make circumvention very challenging even for the *determined* adversary. Finally, we provided a thorough comparison of several prominent modern biometric-based techniques for continuous authentication.

References

1. Chang, K., Bowyer, K.W., Sarkar, S., Victor, B.: Comparison and combination of ear and face images in appearance-based biometrics. IEEE Trans. Pattern Anal. Mach. Intell. **25**(9), 1160–1165 (2003)
2. Eberz, S., Rasmussen, K.B., Lenders, V., Martinovic, I.: Preventing lunchtime attacks: fighting insider threats with eye movement biometrics. In: NDSS (2015)
3. Jain, A.K., Ross, A., Pankanti, S.: Biometrics: a tool for information security. IEEE Trans. Inf. Forensics Secur. **1**(2), 125–143 (2006)
4. Eberz, S., Rasmussen, K.B., Lenders, V., Martinovic, I.: Evaluating behavioral biometrics for continuous authentication: challenges and metrics. In: Proceedings of the 2017 ACM on Asia Conference on Computer and Communications Security, pp. 386–399. ACM (2017)

5. Rasmussen, K.B., Roeschlin, M., Martinovic, I., Tsudik, G.: Authentication using pulse-response biometrics. In: NDSS (2014)
6. Mare, S., Markham, A.M., Cornelius, C., Peterson, R., Kotz, D.: Zebra: zero-effort bilateral recurring authentication. In: 2014 IEEE Symposium on Security and Privacy (SP), pp. 705–720. IEEE (2014)
7. Huhta, O., Shrestha, P., Udar, S., Juuti, M., Saxena, N., Asokan, N.: Pitfalls in designing zero-effort deauthentication: opportunistic human observation attacks. arXiv preprint arXiv:1505.05779 (2015)
8. Ahmed, A.A., Traore, I.: Biometric recognition based onfree-text keystroke dynamics. IEEE Trans. Cybern. **44**(4), 458–472 (2014)
9. Conti, M., Lovisotto, G., Martinovic, I., Tsudik, G.: Fadewich: fast deauthentication over the wireless channel. In: 2017 IEEE 37th International Conference on Distributed Computing Systems (ICDCS), pp. 2294–2301. IEEE (2017)
10. Gia, N., Takimoto, T., Giang, N.D.M., Nakazawa, J., Takashio, K., Tokuda, H.: People identification based on sitting patterns. In: Workshop on Ubiquitous Data Mining, p. 33 (2012)
11. Yamada, M., Kamiya, K., Kudo, M., Nonaka, H., Toyama, J.: Soft authentication and behavior analysis using a chair with sensors attached: hipprint authentication. Pattern Anal. Appl. **12**(3), 251–260 (2009)
12. Riener, A., Ferscha, A.: Supporting implicit human-to-vehicle interaction: driver identification from sitting postures. In: The First Annual International Symposium on Vehicular Computing Systems (ISVCS 2008), p. 10 (2008)
13. Mutlu, B., Krause, A., Forlizzi, J., Guestrin, C., Hodgins, J.: Robust, low-cost, non-intrusive sensing and recognition of seated postures. In: Proceedings of the 20th Annual ACM Symposium on User Interface Software and Technology, pp. 149–158. ACM (2007)
14. Uludag, U., Jain, A.K.: Attacks on biometric systems: a case study in fingerprints. Proc. SPIE **5306**, 622–633 (2004)
15. Bowyer, K.W., Doyle, J.S.: Cosmetic contact lenses and iris recognition spoofing. Computer **47**(5), 96–98 (2014)
16. Marques, D., Muslukhov, I., Guerreiro, T.J., Carriço, L., Beznosov, K.: Snooping on mobile phones: prevalence and trends. In: Twelfth Symposium on Usable Privacy and Security, SOUPS 2016, Denver, CO, USA, 22–24 June 2016, pp. 159–174. USENIX (2016)
17. Mickelberg, K., Pollard, N., Schive, L.: US cybercrime: rising risks, reduced readiness key findings from the 2014 US state of cybercrime survey. US Secret Service. National Threat Assessment Center, Pricewaterhousecoopers (2014)
18. Robb, D.: Sony hack: a timeline (2014). http://deadline.com/2014/12/sony-hack-timeline-any-pascal-the-interview-north-korea-1201325501/
19. Ryan, C.G., Dall, P.M., Granat, M.H., Grant, P.M.: Sitting patterns at work: objective measurement of adherence to current recommendations. Ergonomics **54**(6), 531–538 (2011)
20. Jéquier, E., Tappy, L.: Regulation of body weight in humans. Physiol. Rev. **79**(2), 451–480 (1999)
21. Jaimes, A.: Sit straight (and tell me what i did today): a human posture alarm and activity summarization system. In: Proceedings of the 2nd ACM Workshop on Continuous Archival and Retrieval of Personal Experiences, pp. 23–34. ACM (2005)