Knowledge-Based Systems 246 (2022) 108709

Contents lists available at ScienceDirect

Knowledge-Based Systems

journal homepage: www.elsevier.com/locate/knosys

Contextual security awareness: A context-based approach for assessing the security awareness of users

Adir Solomon *, Michael Michaelshvili, Ron Bitton, Bracha Shapira, Lior Rokach, Rami Puzis, Asaf Shabtai

Department of Software and Information Systems Engineering, Ben-Gurion University of the Negev, P.O. Box 84105, Beer-Sheva, Israel

ARTICLE INFO

Article history:
Received 23 November 2021
Received in revised form 8 March 2022
Accepted 29 March 2022
Available online 5 April 2022

Keywords:
Information security awareness
Human factors
Mobile devices
Deep learning

ABSTRACT

Assessing the information security awareness (ISA) of users is crucial for protecting systems and organizations from social engineering attacks. Current methods do not consider the context of use when assessing users' ISA, and therefore they cannot accurately reflect users' actual behavior, which often depends on that context. In this study, we propose a novel context-based, data-driven, approach for assessing the ISA of users. In this approach, different behavioral and contextual factors, such as spatio-temporal information and browsing habits, are used to assess users' ISA. Since defining each context explicitly is impractical for a large context space, we utilize a deep neural network to represent users' contexts implicitly from contextual factors. We evaluate our approach empirically using a real-world dataset of users' activities collected from 120 smartphone users. The results show that the proposed method and context information improve ISA assessment accuracy significantly.

1. Introduction

In the context of cybersecurity, the term social engineering refers to psychologically manipulating people so they will perform actions for the benefit of an attacker [1–3]. A recent public service announcement from the U.S. Federal Bureau of Investigation (FBI) stated that the global financial loss from email scams (which are largely performed using social engineering attacks such as phishing) was $26 billion for the last three years.[1] Furthermore, businesses around the world have reported a dramatic increase in the number of social engineering attacks since the start of the COVID-19 pandemic.[2] Consequently, social engineering has been classified as one of the most serious cybersecurity threats to businesses in 2020.[3]

Information security awareness (ISA) represents the set of skills that help a user successfully mitigate social engineering attacks [3]. During a social engineering attack, the attacker exploits human behavior rather than a vulnerability in a system [4], so assessing the ISA of users and thereby identifying users who are more vulnerable to social engineering attacks is crucial for enterprise cybersecurity risk assessment. By identifying those users, security officers can implement efficient cybersecurity awareness training programs and adjust information security policies and thus improve organizational security.

Existing methods for assessing the ISA of users can be classified into three main categories based on the data source used: (1) ISA assessment using questionnaires [5,6] where the users are asked to report on their knowledge and behavior for different scenarios using surveys. Their responses are then analyzed to detect users with low ISA. (2) ISA assessment using measurements of the actual behavior [7,8], where the users' actual behavior is monitored. (3) ISA assessment using attack simulations and challenges [9,10] which simulate cybersecurity threats and are mainly conducted to record and analyze users' responses. While the measurement of the actual behavior and attack simulations can assess users' ISA objectively, they have two fundamental limitations; first, they do not consider users' context (e.g., opening an email from home compared to opening an email at work). Since human behavior often depends on a person's context, these methods are less accurate by nature. Second, because these methods typically rely on expert knowledge to integrate the raw measurements of a user's behavior into risk, they are unable to detect dynamic changes in a user's behavior, which varies based on his/her context.

In this study, we address the limitations mentioned above by proposing a context-based, data-driven approach for assessing

* Corresponding author.
E-mail addresses: adirsolo@post.bgu.ac.il (A. Solomon), michmich@post.bgu.ac.il (M. Michaelshvili), ronbit@post.bgu.ac.il (R. Bitton), bshapira@bgu.ac.il (B. Shapira), liorrk@post.bgu.ac.il (L. Rokach), puzis@bgu.ac.il (R. Puzis), shabtaia@bgu.ac.il (A. Shabtai).

1 https://www.ic3.gov/media/2019/190910.aspx.

2 https://www.jpmorgan.com/commercial-banking/insights/spotting-and-preventing-covid-19-social-engineering-attacks.

3 https://www.kaseya.com/blog/2020/04/15/top-10-cybersecurity-threats-in-2020/.

https://doi.org/10.1016/j.knosys.2022.108709
0950-7051/

users' ISA. We define the context as factors that characterize the user's state considering his/her recent behavior [11,12] and personal information [12,13], along with physical information regarding the time and space [12,13]. Namely, to assess users' ISA, we use the following contextual factors: temporal information (e.g., day of the week), points of interests (POIs), browsing habits, and user information.

The development of context-based ISA assessment faces two main challenges: (1) defining the contextual factors that affect users' ISA; and (2) measuring those contextual factors continuously and objectively. In order to deal with these main challenges, we suggest implementing our context-based, data-driven approach by utilizing a deep neural network (DNN). This will enable us to learn different latent patterns between users' different contextual factors and their ISA. Moreover, by using a DNN, we will be able to provide ISA assessments objectively and dynamically, i.e., when the user's context has changed.

We utilize a deep learning architecture to represent a user's contexts implicitly from contextual factors. That is, given a set of contextual factors (such as browsing habits, geographic location, temporal information, and personal information), the proposed deep learning architecture derives a latent context which is used to assess a user's ISA with respect to that specific context of use. In our experiments, we used a deep learning architecture with a gated recurrent unit (GRU) [14], attention layers [15], and embedding layers.

In order to assess users' ISA, we examine our method empirically on a real-world dataset of users' activities collected from 120 mobile phone users over a period of seven weeks. The dataset was collected as part of a prior study [3] in which data was obtained from different sources, such as mobile sensors (to collect information on users' actual behavior based on sensor sampling), a VPN client (to collect users' Internet traffic), and security questionnaires. Using different data sources allowed us to capture actual user behavior, as well as their self-reported behavior, thus enabling us to draw a full picture of users' routine activities and develop an effective method for assessing users' ISA.

In summary, the main contributions of this study are as follows:

- We present a novel context-based, data-driven approach for assessing the ISA of users dynamically, based on different contextual factors, such as spatio-temporal information, and users' attributes and activities.
- We implement our context-based, data-driven approach by utilizing a deep learning architecture that assess users' ISA, which is capable of learning different behavior patterns and providing ISA predictions with high accuracy.
- Relying on the dataset collected by Bitton et al. [3] we accurately evaluate users' ISA based on real dynamic data on actual users' activities and their sampled contexts, in contrast to information collected from simulations or questionnaires. We evaluated our approach on a real-world dataset of users' activities collected from 120 mobile phone users over a period of seven weeks.

2. Related work

In this section, we discuss related work on various topics related to ISA. We present studies that explore users' behavior patterns; methods, platforms, and data sources commonly used to assess users' ISA; and machine learning methods for assessing ISA.

2.1. Methods for assessing the ISA of users

Questionnaires. Security questionnaires are the most common approach for assessing the ISA of users [6,16,17]. Questionnaires can be used to measure users' activities and usage patterns regarding the users' ISA. For instance, Androulidakis et al. [18] used a questionnaire to examine the blacktooth sensor on users' ISA. Their findings showed that users' ISA is correlated to blacktooth usage. Onarlioglu et al. [19] used a security questionnaire to assess how well users deal with Internet attacks and showed that features extracted based on the URLs are highly effective for deceiving users without rich technical knowledge. Mylonas et al. [20] used a questionnaire to examine users' ISA regarding downloading applications from official repositories (e.g., Google Play). Their findings showed that most users believed that downloading applications from official repositories are risk-free, which could be indicative of a user' with low ISA.

Notable methods in this domain are the HAIS-Q and SeBIS security questionnaires [5,21], which were respectively designed for assessing the ISA of PC users and smartphone users. Moreover, HAIS-Q is shown to be a correlative and reliable measure with the users' ISA [6]. Moreover, organizations can employ the HAIS-Q to define the impact of security they should employ [6].

The main advantage of security questionnaires over alternative approaches for assessing users' ISA is in their ability to explore the knowledge and attitudes of users with respect to various aspects of security. However, security questionnaires have three main limitations: First, since questionnaires are based on self-reporting, they are very biased and therefore cannot be used to accurately evaluate users' ISA [3]. Second, since questionnaires require users' active involvement and cooperation, they therefore cannot be used to evaluate users' ISA frequently. Third, security questionnaires cannot consider users' context which is dynamic and varies over time.

Passive measurements. Recent studies have evaluated users' ISA based on passive measurements, i.e., assessing users' ISA by sampling sensors in real time [3,7,8]. The authors of these studies developed different frameworks (for both PC and mobile platforms) that can be used to monitor users' actual behavior (e.g., monitoring application usage, blacktooth and Wi-Fi usage, and browsing habits, etc.) The main advantage of these methods is their ability to evaluate users' ISA objectively and continuously. However, these methods suffer from two main limitations: First, they do not consider the context of use during the assessment, which could improve the accuracy of ISA assessment. Second, they rely heavily on expert knowledge. For example, the method presented in [3] assists security experts in integrating sensor readings into a single measure that can be used to reflect users' ISA.

Attack simulations. Other studies used attack simulations in order to evaluate users' ISA [9,10,22]. The vast majority of these methods only focus on simulated phishing attacks. There are two main limitations when employing assessment methods that are based solely on attack simulations. First, attack simulations measure the momentary behavior of subjects during specific events, and therefore they are very sensitive to environmental and contextual factors. Second, attack simulations require interaction with the user (e.g., sending an email to the user), and therefore such simulations cannot be used very often.

In the proposed method we manage to capture real users' behavior by extracting different contextual factors. Moreover, we assess users' ISA based on their activities and on real events.

2.2. The impact of contextual factors on ISA

Recent studies [23,24] examined the relationship between ISA and demographic attributes, such as age and gender. These studies measured the ISA using the widely adopted HAIS-Q security questionnaire [5]. Their results show a significant difference in the average ISA scores of different age groups. However, with respect to gender, a significant difference was only observed in one of the studies [23]. Bitton et al. [3] also examined the relationship between ISA and demographic attributes. In contrast to [23,24] which measured the ISA using a security questionnaire, the authors calculated the ISA based on the user's actual behavior. Their results showed a difference in the average ISA scores of users from different age groups and academic backgrounds; those findings however, were not statistically significant.

Other methods [1,3,25,26] examined the relationship between ISA and contextual factors. Duncan et al. [1] developed a persona-centered ISA methodology to reduce a company's security risks. In their work, the authors identified contextual factors based on a user's persona (e.g., attitudes, motivations), which was derived during multiple interviews, in order to assess ISA. Then, they analyzed the characteristics of each persona and incorporate their findings into their awareness design methods (e.g., quizzes, short video clips, or short topic briefs). Their methodology was shown to improve business security. Karyda et al. [27] demonstrated that contextual factors derived from the organizational culture, organizational structure (e.g., code of ethics), and management support improved information system security policies. Similar findings were observed by Ifinedo et al. [28] whose work showed a correlation between information security concerns in the financial services industry and various contextual factors, such as transparency levels, information, and communication technology use regulations.

Dang et al. [26] explored the relationship between points of interest (POIs) and users' ISA (measured using a security questionnaire). Their findings show that the user's ISA varies for different points of interest. In other research, Dang et al. [25] demonstrated that contextual factors that users reveal in their social network profile can also be used to assess their ISA. Bitton et al. [3] showed that an ISA measure that is based on the actual behavior of users (as measured using a mobile agent and network probe) correlates with the user's ability to mitigate social engineering attacks. Canali et al. [29] explored users' visits to malicious URLs based on their browsing history. They observed that users are more likely to visit such URLs on weekends and at night (between 9:00 pm and 2:00 am).

These works support our main assumption that in order to assess a user's ISA accurately, the context of use must be considered.

2.3. Machine learning approaches for assessing user's ISA

Several studies [30,31] showed that phishing attacks can be detected using different machine learning classifiers. Tjostheim et al. [32] trained a logistic regression classifier on users' information (e.g., demographic attributes, education) to predict an individual's susceptibility to phishing attacks. Sharif et al. [33] proposed a system that predicts whether a user could be exposed to malicious sites; the authors used a random forest classifier based on several features: the presence of an antivirus application on the user's device, previous downloads from unofficial app stores, etc. Their results indicate that users that have previously been exposed to malicious sites have a higher probability of being exposed again. In this study we implement our approach with a deep learning architecture, thus, assess the ISA of users dynamically. Using deep learning we manage to capture different behavior patterns and providing ISA predictions with high accuracy.

In the work of Foroughi et al. [34] the authors proposed a machine learning multi-agent model to profile users based on the activities performed with their computers at home. Thus, they were able to identify users' profiles that have low security awareness and suggest that the users change the way they perform these activities. This study covered only users' activities performed within a limited context at the users' homes with their personal computers. In contrast to their work, in this study, we examine the users' context based on their mobile devices; thus, we are able to model different contexts and all of the users' routine activities. Another recent study [35] used a machine learning-based method to assess users' ISA based on an ISA score, i.e., the awareness was encoded into three awareness levels: low, average, and high. The authors extracted users' information and used traditional machine learning classifiers (e.g., KNN, SVM, decision tree) to assess users' ISA. Unlike our study, which used dynamic data derived from users' actual activities, their method is static, and their dataset is only derived from an online HAIS-Q security questionnaire. Furthermore, the authors did not examine different contexts (such as spatio-temporal) to assess the users' ISA, although using contextual factors has been shown to be useful for advancing solutions for different tasks in several domains, such recommendation systems [36,37], crime prediction [38], and user modeling [39,40]. To the best of our knowledge, in the domain of cybersecurity only limited users' context (e.g., user information) has been explored to assess users' ISA. Therefore, in this study, we propose assessing users' ISA dynamically, based on rich contextual information, by employing a deep learning framework.

3. Method

In this section, we present an implementation of our context-based, data-driven approach utilizing a deep learning architecture. We begin by formulate the problem in Section 3.1. Next, in Section 3.2 we describe the rational for deriving users' ISA score dynamically for short timeframes. In Section 3.3, we describe the dataset used in this study. Next, in Section 3.4, we specify the various contextual factors considered in the assessment and describe the techniques used to derive latent context from contextual factors. Finally, in Section 3.5, we outline how to use that context to evaluate users' ISA.

3.1. Problem formulation

To formulate our problem we define the following:

- T_i^u – The ith timeframe of user u. The whole data period is split into non-overlapping consecutive five-minute timeframes.
- $C(T_i^u)$ – The context derived for user u in timeframe T_i^u.
- $I(T_i^u)$ – The set of ISA indicators derived for user u in timeframe T_i^u.
- $I_j(T_i^u) \in I(T_i^u)$ – The jth ISA indicator derived for user u in timeframe T_i^u, where $I_j(T_i^u) = 0$ denotes that u behaves securely with respect to the jth ISA indicator in timeframe T_i^u. In contrast to $I_j(T_i^u) = 1$ which indicates that u behaves unsecurely with respect to the jth ISA indicator in timeframe T_i^u.
- $ISA(T_i^u)$ – The observed ISA of user u in timeframe T_i^u, where

$$ISA(T_i^u) = \begin{cases} 0, & \exists j \text{ such that } I_j(T_i^u) = 1 \\ 1, & \text{otherwise} \end{cases}$$

Definition 1 (*Context-based ISA Prediction*)**.** Given a time-series of (1) *historical* user's contexts $C(T_0^u)...C(T_n^u)$ and (2) *historical* user's ISA indicators $I(T_0^u)...I(T_n^u)$, our goal is to predict $ISA(T_{n+1}^u)$.

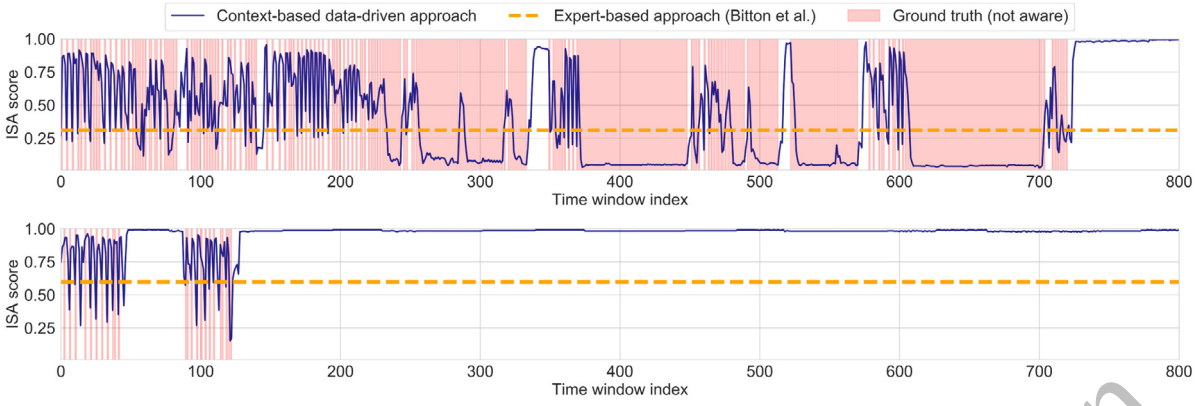

Fig. 1. Comparison of the ISA scores derived for two users by the context-based, data-driven approach (black) and the method of Bitton et al. (yellow).

We denote this prediction as $ISA_{score}(T^u_{n+1}) \in [0, 1]$ such that 1 indicates that we predict that u will be completely aware of cyber threats in timeframe T^u_{n+1}.

3.2. Rationale for calculating users' ISA for short timeframes

We argue that users' behavior is too complex to be represented using a single ISA measure. That is, it is quite possible that a user with high security awareness will behave in an unsecure manner in a certain context and vice versa. For this reason, we believe that a user's ISA score should be calculated continuously for short timeframes.

To better emphasize the benefits of calculating a user's ISA score for short timeframes, we select two users from our dataset. The users were selected based on their ground-truth behavior (*ISA*), where one of the users performs multiple unsafe activities, and the other has only a few unsafe activities. We then calculate their ISA scores (ISA_{score}) based on two different approaches: (1) the proposed context-based, data-driven approach for assessing user's ISA, which calculates users' ISA scores for short timeframes, and (2) a recently published, static, expert-based method for assessing users' ISA, which calculates a single ISA score for a user [3] based on four different social engineering attack vectors: application-based attacks (application score, e.g., Trojan application, permission abuse), man-in-the-middle attacks (e.g., certificate manipulation), phishing attacks, and attacks that are based on weak authentication (e.g., password cracking).

The results are presented in Fig. 1. The x-axis represents the timeframe index, which is ordered chronologically, and the y-axis represents the ISA score. The ground truth, *ISA*, (indicated by the red background) denotes timeframes in which the user behaved unsafely. The yellow line indicates the user's ISA score derived from the expert-based method, and the black line indicates the score derived from the proposed data-driven approach. As can be seen, the user classified with a low ISA score by the expert-based method (the upper figure) has many more instances of unsafe behavior compared to the user that was classified with a high ISA score (the lower figure). However, it can also be seen that a user with a high degree of security awareness can sometimes behave unsafely and vice versa.

Unlike the expert-based method, our data-driven approach provides ISA scores that are correlated with the users' actual behavior. For example, as seen in the lower figure, when the user did not perform any unsafe activities, our score was in fact high, reflecting the user's actual activities, in contrast to the ISA score obtained using the expert-based method which was fixed at 0.6 and did not reflect changes in the user's actual behavior. These types of changes cannot be modeled by the expert-based method. On the other hand, the proposed data-driven approach (which considers the context of use) is able to capture these changes. Thus, this example supports our assumption that contextual factors can be used to better assess the ISA of mobile phone users.

Table 1
Dataset collection — subject statistics.

User's attribute	Categories	Distribution
Gender	Male	0.66
	Female	0.33
Age group	18–24	0.45
	25–30	0.55
Academic status	Alumnus/Alumna	0.17
	Graduate student	0.51
	Undergraduate student	0.32
Faculty	Engineering	0.37
	Humanities	0.32
	Natural sciences	0.22
	Other	0.09
Programming skills	Low	0.50
	Medium	0.16
	High	0.33
Build website	Low	0.66
	Medium	0.26
	High	0.07

3.3. Dataset

In this study, we use the dataset collected in the work of Bitton et al. [3], which was collected from 120 subjects over a period of seven to eight weeks. The subjects are 18 to 30 year-old undergraduate students, graduate students, and alumni from various disciplines. Descriptive statistics of the subject population are presented in Table 1. As part of this work we are publishing the dataset for academic and research purposes.[4]

The dataset was derived from three data sources: sensor readings from a mobile device agent installed on the subjects' smartphones; Internet traffic collected using a VPN client installed on the subjects' smartphones; and demographic attributes and self-reported behavior collected using a security questionnaire.

Mobile device agent. The mobile device agent installed on the subjects' smartphones monitored their behavior continuously and objectively thus providing us the ability to measure users' ISA in different situations. The mobile device agent's complete list of sensors is as follows:

[4] https://bit.ly/3HTCj5i.

- Wi-Fi – Records the connected Wi-Fi access point (SSID, BSSID) and its security capabilities.
- blacktooth – Detects connected blacktooth devices.
- Traffic – Collects statistical information about the volume of ingoing/outgoing network traffic for each package and process on the mobile device.
- Installed apps – (1) Samples the installed applications and their permissions, and (2) scans the installed packages in VirusTotal and issues a notification when a package is updated or removed.
- Running apps – Samples the running applications and processes.
- Application changes – Monitors application updates, application deletions, and application installations.
- Browser search – Monitors the browser searches and URLs, scans the URLs in VirusTotal and Web of Trust.
- Emails – Provides statistics about the email account (email sent, email received, spam received, etc.).
- SMS – Monitors links sent in SMSs.
- Hardware – Collects model and brand information.
- Software – Collects information about the OS, build number, and firmware.
- Root checker – Checks whether the device is rooted.
- Screen lock – Checks which type of screen lock is used (PIN, pattern, none, etc.).

Internet traffic. The users' Internet traffic was collected by using a VPN client installed on the subjects' smartphones. Monitoring the users' Internet traffic allowed us to extract users' behavioral attributes passively, i.e., it did not require any collaboration from the users. We focus on four behavioral attributes:

- Domain names — We inspect the domains exchanged within application layers' protocols (e.g., HTTP host, SSL server name). We identify two domain categories that could be indicative of users' high or low security awareness level: information security-based categories (such as spam, ads, malware, etc.) and content based categories that focus on the content of the session (such as social networks or email).
- Application level protocols — We focus on two application level protocols: HTTP and TLS. Inspecting the header of an HTTP packet could reveal general information about a user's smartphone. Specifically, we can extract the device model and operating system from the user-agent field. This information can be used to better understand whether the user is conscientious about updating his/her operating system. The application level protocol of the TLS can be used to detect when a user accepts an untrusted SSL/TLS certificate.
- Transport layer protocols — We detect unencrypted protocols in which the user transmits private information (e.g., email addresses, passwords, credit card numbers, GPS coordinates) by employing a deep packet inspection (DPI) process.
- Contextual-based data — We examine the user's behavior within a specific context of use, e.g., clicking malicious links while reading emails. This type of data could not be evaluated directly from the raw traffic due to the encryption of the raw network traffic. However, it can be learned by using indirect indicators, e.g., analyzing users' communication with email services.

Security questionnaire. Today, the most common method for assessing users' ISA is to use a security questionnaire [6,16,17]. The questions in the questionnaire utilized in our study focus on three areas:

- Likelihood for performing an action — In Appendix, we present the security questionnaire used to measure the likelihood that the user will perform a certain action. All the questions were measured according to a five-point Likert scale with the following answers: "Never", "Unlikely", "Medium Likelihood", "Very Likely", and "Always".
- Device connectivity — An additional set of questions was used to measure the frequency with which device connectivity was turned off. The question was "How often do you turn off the following components in your device?" for the mobile components of GPS, Wi-Fi, NFC, and blacktooth. The answers were: "Always off, "Unless needed", "Turned on when needed, but often forget to turn off", "Turned on if the battery allows", "Always on", and "I didn't know it could be turned off".
- Updating a service or application — We used a set of security questions to measure the frequency of updating a service/application. The question was "How frequently do you update the following services/applications on your smartphone?" regarding the following: operating system, system application (e.g., web browser)s, security applications (e.g., anti-virus), and other applications (e.g., Facebook). The answers were: "I use auto-update", "When I receive a notification", "Once a month", "Once in six months", and "Only when I must".

The dataset includes undergraduate students, graduate students, and alumni from two academic institutions located in two different regions. When we analyzed the dataset more closely, we saw that (1) users' geographical locations were spread across the country; and (2) users had different applications installed on their mobile devices. More specifically, popular applications, such as WhatsApp, Chrome, Waze, Facebook, Shazam, Skype, Gett, and Instagram were installed on more than 50% of the devices. On the other hand, most of the applications were installed on less than 20% of the devices, and more than 40% of the applications were installed on only one device. Therefore, the population examined in this research is quite diverse.

In order to label the dataset, we use the taxonomy presented in [41], which defines the criteria for a security aware smartphone user. Specifically, we define eight behaviors which indicate that the user is unsafe in a given context of use. The user behaviors were defined based on the top mobile security threat categories (mentioned in the 2020 Kaspersky report[5]):

1. **ISA Indicator 1 — Visiting a malicious website.** We use VirusTotal[6] to classify visits to malicious websites. Visiting a malicious website can expose mobile users to different dangerous cyber threats, e.g., phishing, account hijacking, malicious QR bar codes. Moreover, visiting a malicious website could result in exposing the user's browser to gain system-level privileges or even stealing the user's sensitive information from the mobile phone's embedded sensors (e.g., camera, GPS, fingerprint, accelerometer) [42,43].
2. **ISA Indicator 2 — Accepting an untrusted SSL/TLS certificate.** Recent papers [44,45] and attack case studies [46,47] indicate that accepting an untrusted certificate is a dangerous behavior.[7] To measure this behavior, we counted network sessions that included untrusted certificates (the classification as untrusted certificate was performed using

[5] https://www.kaspersky.com/resource-center/threats/top-seven-mobile-security-threats-smart-phones-tablets-and-mobile-internet-devices-what-the-future-has-in-store.

[6] http://www.virustotal.com/.

[7] https://www.computerweekly.com/news/2240214897/Apple-users-at-risk-of-SSL-man-in-the-middle-attack.

OpenSSL) but were not terminated immediately after the handshake, which indicates that the user accepted the untrusted SSL/TLS certificates. To eliminate 'false positives' we filtered out localhost connections, which in many cases can indicate a security aware person who intentionally accepts untrusted certificates in internal networks.

3. **ISA Indicator 3 — Clicking on pop-ups or advertisements.** This indicator was extracted from the user's 'network traffic.' Specifically, we counted instances where the user
visited domains that are classified by Web of Trust and VirusTotal as advertisement sites and analytical services.
4. **ISA Indicator 4 – Transmission of private information over an unencrypted network protocol.** We identify the transmission of private information (e.g., the transmission of private content, such as email addresses, phone numbers, and credit card numbers) within the user's unencrypted network traffic. This could be caused by (1) the existence of untrusted applications, which do not use the standard encryption AIP to send data to the server; (2) dangerous permissions given by the users to an untrusted application. In both cases, this indicator points to low security awareness.
5. **ISA Indicator 5 — Downloading files in an unsecure manner.** Downloading files in an unsecure manner (e.g., over HTTP), exposes the user to man-in-the-middle attacks. Specifically, a man-in-the-middle attacker can replace/change files downloaded via HTTP, which compromise the file's integrity.
6. **ISA Indicator 6 — Connecting to an unsecure Wi-Fi network.** We identify users that connected to unsecured Wi-Fi networks without using a VPN. Connecting to untrusted Wi-Fi networks increases the chances of exposing the users to eavesdropping, SSL strip, and man-in-the-middle attacks [42,48,49].[8].
7. **ISA Indicator 7 — Installing unsafe applications.** In the Android operating system third-party applications need to request the user to grant sensitive permissions through the access control mechanism at run-time or during an update [50]. Additionally, updating applications is indicative of the users' ISA as it is being studied through questionnaires [16,20]. An attacker can hack the access control mechanism with social engineering attacks, or due to the user's limited understanding [42,51]. Therefore, we identify unsafe applications that the user installed on his/her mobile phone. We define an unsafe application based on the VirusTotal scan or if the application used more than 14 sensitive permissions, such as recording audio, using the camera, sending and reading SMS messages, etc.
8. **ISA Indicator 8 — Clicking on URLs from an SMS.** Examine different mobile phone infections, previous works indicate that malicious websites are mostly delivered by blacktooth and SMS messages [42,52]. Therefore, we identify when the user clicked on URLs that he/she received in an SMS message. For this indicator, we focus only on messages that the user received from mobile phones that are not part of the user's contacts.

If we identify one of the above behaviors within the timeframe of a given context of use, we label the context as 'the user is unsafe', $ISA(T_i^u) = 0$. Otherwise, we label the context as 'the user is safe', $ISA(T_i^u) = 1$.

The distributions of the ISA indicators based on a timeframe of five minutes are presented in Table 2. We also present the data source for each ISA indicator in Table 3. The sampling rate for all indicators was event-based, i.e., each indicator was recorded based on the user's actions. The probability of observing one of the eight behaviors indicating users' security unawareness in a five-minute timeframe is 0.2451. As can be observed, in most cases (0.8818), the unaware behavior is exhibited by *transmission of private information over unencrypted network protocol*. In order to detect personal information that was sent in an unsecure manner, we implement a dedicated deep packet inspection (DPI) process [53,54]. Specifically, we reconstruct the TCP session, and for any non-TLS transmission we also decode GZIP/XML/JSON encoded traffic and extract the content in a readable format. This content is then analyzed using regular expressions and string matching techniques, to identify personal information, such as email addresses, username and password pairs, phone numbers, and credit card numbers. The information most commonly seen consists of username and password pairs.

Focusing more closely on the personal information, we note that the user must specifically grant permission allowing a given application to send his/her private information. In order to discern between active user actions vs actions made by a background process, we implement a simple procedure that filters actions made by a background process.

Ethical considerations. The data collection process includes sensitive and private information of users' identities (such as age and gender) and their routine activities (such as browsing habits). Therefore, the data collection process has been fully approved by the institutional review board (IRB).

Table 2
Distribution of ISA indicators in the dataset (based on a five-minute timeframe).

Category	Dist.
The user is safe	0.7548
The user is unsafe	0.2451
- ISA Ind. 1 – Visiting a malicious website	0.0577
- ISA Ind. 2 – Accepting an untrusted SSL/TLS certificate	0.0113
- ISA Ind. 3 – Clicking on pop-ups or advertisements	0.0127
- ISA Ind. 4 – Transmission of private information	0.8818
- ISA Ind. 5 – Downloading files in an unsecure manner	0.0013
- ISA Ind. 6 – Connecting to an unsecure Wi-Fi network	0.0321
- ISA Ind. 7 – Installing unsafe applications	0.0025
- ISA Ind. 8 – Clicking on URLs from SMSs messages	0.0003

Table 3
Data source and sampling rate for all ISA indicators.

ISA indicator	Data source
ISA Indicator 1	Agent or Traffic
ISA Indicator 2	Traffic
ISA Indicator 3	Agent or Traffic
ISA Indicator 4	Traffic
ISA Indicator 5	Traffic
ISA Indicator 6	Agent
ISA Indicator 7	Agent
ISA Indicator 8	Agent

3.4. The contextual factors used for assessing users' ISA

In order to assess the ISA of mobile phone users, we use previously observed ISA indicators and the following contextual factors: temporal information, points of interest (POIs), browsing habits, and the user information. The contextual factors are extracted over time, using a sliding window feature extraction technique (except for the user information which is static). We focus only on timeframes that contain an occurrence of any contextual factors or ISA indicators. For each user, we order all of the data chronologically and use time threshold of five minutes

[8] https://www.cybsafe.com/research/security-behaviour-database/behaviours/disabling-automatically-connect-to-wi-fi-on-mobile-devices/.

Table 4
Data source and sampling rate for all contextual factors.

	Data	Data source	Sampling rate
Temp.	Day of week	Agent or Traffic	Continuously
	Part of day	Agent or Traffic	Continuously
POI	GPS coordinates	Agent	Every 15 min
Browsing habits	External services – URL	Traffic	Event-based
	External services – DNS	Traffic	Event-based
	HTML components – iframe	Traffic	Event-based
	HTML components – applet plugin	Traffic	Event-based
	URL's features – URL address	Traffic	Event-based
	URL's features – Length of URL	Traffic	Event-based
	URL's features – Symbols	Traffic	Event-based
	URL's features – Browsing behavior	Traffic	Event-based
User Info.	Gender	Questionnaire	Static
	Age group	Questionnaire	Static
	Academic status	Questionnaire	Static
	Academic faculty	Questionnaire	Static
	Ability to build websites	Questionnaire	Static
	Programming skills	Questionnaire	Static

as the timeframe. It should be noted that we use non-overlapping sliding windows. The contextual factors and their data sources and sampling rates are presented in Table 4.

Temporal information. Previous works [55,56] proposed robust methods by discretizing users' temporal information to different bins, thus managing to achieve accurate recommendations for various tasks of recommendation systems. Motivated by these studies, we extract the following features to model users' temporal contexts: the day of the week (1–7) and the part of the day (night 00:00–06:00, morning 06:00–12:00, afternoon 12:00–18:00, and evening 18:00–00:00). The temporal information is represented using two ordinal variables for each timeframe.

Points of interest (POIs). We extract users' POIs based on sampled GPS coordinates (i.e., latitude and longitude). We define a POI as the semantic meaning of a physical location sampled by the GPS coordinates for a given timeframe. We extract the semantic meaning (e.g., residential, commercial, restaurant) for each pair of GPS coordinates using OpenStreetMap [57], and treat this feature as categorical. Using the POI we are able to capture the user's location and to learn the POIs in which the user is less aware of cyber threatens [26]. Moreover, by using generic location categories derived from OpenStreetMap, and not specific geolocation coordinates, we managed to create a robust model that could be applied on different areas.

Browsing habits. We extract the visited URLs from the users' Internet traffic. The visited URLs include the entire browsing history, as well as indications of application usage. Recent works [33,58,59] showed that URLs' metadata could be beneficial for detecting phishing attempts and malicious websites. Motivated by these studies, we also use URLs' metadata to extract features that could be indicative of users' ISA assessments. In particular, we extract the following features:

- External services — we check whether the visited URL is in the top 100,000 visited websites, using the Alexa service,[9] similar to the work of [33]; we also examine whether the domain name server (DNS) of the URL exists, using the WHOIS service, similar to work of [58].
- HTML components — we check the number of times that the iframe tag and applet plugin appeared in the HTML of the visited URL, since they can be indicative of malicious websites [58].
- URL's features — we treat the URLs as categorical feature and extract additional features based on the URL's textual address similar to previous works [19,59]: the length of the URL, the number of times that the URL contains a dot, dash, double backslash, or the at symbol (@).
- Users' previous browsing behavior — based on the work of [33], we extract features that can be used to imply the users' behavior and features that are indicative of the number of times that the user was exposed to malicious sites in the past, and compute the ratio between the number of previously visited malware websites and all of the user's visited websites.

User information. Motivated by the findings presented in previous studies [23,24,41], we utilize the following user demographic attributes to assess users' ISA: gender, age groups (two groups: 18–24 years-old and 25–30 years-old), academic status (undergraduate student, graduate student, alumnus/alumna), and academic background/faculty (engineering, humanities, or natural sciences). Note that in contrast to other contextual factors which vary over time, users information is static, and therefore it is fixed across the different timeframes. In addition, we consider users' technical experience, such as the their ability to build websites, and their programming skills (divided into three groups: low, medium, high). All user information is represented using categorical variables.

Using all of the above-mentioned contextual factors, we are able to collect information regarding the user identity, drawing a full picture for the user's physical state, and answer on four important questions: "when the user performed a recent activity?" with the context of the temporal information context, "where the user has been recently active?" with the POI context, "what a recent activity the user performed?" with the browsing habit context, and "who is the user?" with the user information context.

Data processing and integration. We represent the POIs with a binary vector (each POI has an entry in the vector), and for each visited POI entry in each timeframe we marked '1'. The browsing habit features were processed as follows: the external services were mapped to binary values, the HTML components (e.g., the number of times that the iframe tag), and URL's features (e.g., the number of times the URL contains a dot) were normalized by the number of URLs in each timeframe. We also represent the URLs with a binary vector (similar to the POI representation). Each feature values from the temporal information (e.g., day of week) and user's information (e.g., gender, age group) contexts

[9] https://www.alexa.com/.

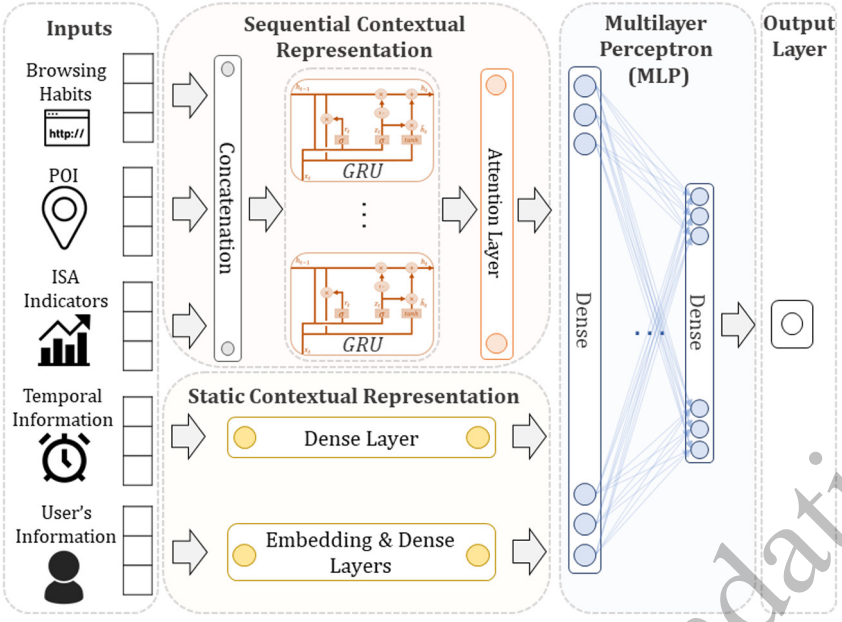

Fig. 2. The deep learning architecture used to implement our context-based, data-driven approach.

were mapped to numeric values. We also use a binary vector to represent ISA indicators, and for each detected ISA indicator in each timeframe we marked '1'. The integration was performed based on the actual observed user behavior and sampled sensors in each timeframe. For instance, if the user did not opened any URL, the URL binary vector values were consist only of zeros.

3.5. Network architecture

Each contextual factor changes with a different frequency; therefore we use different approaches for representing the various contextual factors. Specifically, in order to represent the browsing habits, POIs, and ISA indicators, which change with high frequency, we utilize a gated recurrent unit (GRU) and attention layers (the orange area in Fig. 2), based on the concatenation of these sequential contextual factors. We use embedding and dense layers (the yellow area in Fig. 2) to represent user and temporal information (which are fixed for many timeframes).

Gated recurrent unit. A GRU is a variant of a recurrent neural network (RNN) that is able to capture lengthy time series information. The GRU concept, introduced by Cho et al. [60], can selectively remember and update previous information.

Attention layer. Attention is a vector of weights that measure the contribution of each part in the input vector [61]. Attention layers have been shown to be effective at capturing complex patterns by focusing on the most relevant information in the input sequences [15]. The GRU units are able to capture long-term dependencies, however a GRU cannot model the contribution of each hidden state to the latent representation of the context. In order to address this limitation, we add an attention layer to the output of the GRU.

First, we scale the output of the GRU hidden states, and then we use the *softmax* function in order to calculate the weights of each part of the scaled hidden units. Finally, we multiply the weights of the output of the GRU hidden states. Formally, we describe the self-attention mechanism with the following equations:

$$\begin{aligned} s_t &= tanh(W_s h_t + b_s) \\ a_t &= \frac{exp(s_t)}{\sum_{i=1}^{n} exp(s_i)} \\ \hat{a} &= \sum_{i=1}^{n} h_i a_i \end{aligned} \tag{1}$$

such that in s_t, we apply a *tanh* on the vector h_t, which is the GRU's hidden units' output at timestamp t, with the weights W_s and bias b_s. In a_t, we apply the *softmax* function on s_t and use $\hat{a}$ to represent the dot product on h_t with a_t.

Using this calculation allows us to focus on the most important parts of the GRU outputs, thus improving the sequential contextual factors' representation.

Embedding and dense layers. Embedding layers are hidden layers that represent high-dimensional data with a compressed representation. The embedding representation was shown to be useful for various applications, such as representing words [62], users' stay points [63], graph representations [64], and even facial recognition [65]. In this study, we propose using embedding layers to represent the user information which is fixed across all of the timeframes. Therefore, we used a simple embedding layer to integrate the categorical variables into the neural network. By using the embedding representation within the deep learning architecture, we are able to capture hidden patterns within the user information and other contextual factors. We leverage these patterns to improve ISA assessment accuracy.

Dense layers are hidden layers that represent high-dimensional data with a compressed representation by applying a mathematical activation function (e.g., *ReLU*). In our network we use dense layers to represent temporal contexts and the following user information: gender, academic status, programming skills, and the ability to build websites.

Multilayer Perceptron (MLP). We use fully connected dense layers in order to provide ISA assessments (the black area in Fig. 2).

4. Evaluation

4.1. Evaluation settings and methods

Hardware. We implement our deep learning framework on a machine with the following settings: a GPU card of RTX 2080, CPU of Intel(R) Xeon(R) Silver 4214 CPU @ 2.20 GHz and 72G of RAM, Samsung DDR4 2666 MHz.

Optimization. To find the optimal network parameters, we employ the Bayesian optimization framework [66] and report on the optimal parameters and the examined parameter ranges. For selecting the optimal dense and embedding size we examine different values from the range of [1, 5, 10, 15, 20, 25, 30]. Additionally, to select the optimal number of GRU units we examine the range of [10, 15, 20, 25, 30]. We also use the range of [10, 20, 30] for selecting the optimal number of epochs, and the range of [32, 64, 128, 256] for the selecting the optimal batch size. We examine the value range of [0.0001, 0.001, 0.005, 0.01] to select the optimal learning rate. We train the models using backpropagation through time, with the Adam optimizer. The optimized selected parameters are as follows:

- GRU — In our network implementation we use 30 GRU units to represent the sequential contextual factors. All GRU units analyze six timestamps (each timestamp represents a timeframe) simultaneously. That is, when analyzing contextual factors that change with high frequency, our model considers six timeframes to derive the latent context. For instance, if the time threshold (for the sliding window) is set at five minutes, then each latent context is constructed based on the previous 30 min.
- Embedding and dense layers — We set 10 units for each of the embedding layers that represent the user's age and academic faculty, 15 units for the embedding layer that represents the user's ID, and 15 neurons with *ReLU* activation for all of the neurons in all of the dense layers.
- MLP — we use four dense layers (with 50, 20, 5, and 1 neuron(s) respectively) to represent the latent context and predict a user's ISA. We use *ReLU* as the activation function for all of the neurons in the first four layers, and *sigmoid* in the last layer.

We also use 10 epochs (and a batch size of 256), where the learning rate is set at 0.01. In addition, to cope with the imbalanced dataset, we use the weighted cross-entropy loss function, with the proportion of 1:5 for aware and unaware ISA classes.

Evaluation methods. We use the following two evaluation methods:

- Time-based estimation — In this evaluation method, the data of each user is first put into chronological order and then split so that the first 60% of the records are used for training the model, the next 20% used to validate the model, and the remaining 20% are used for testing the model. Unless we specify otherwise, this was the evaluation method used.
- User-based cross-validation — In this evaluation method, the users are divided into 10 groups, each of which is composed of 12 users. Then, we train the models with 10-fold cross-validation over the 10 groups so that each time a different group is used as the test set. This ensures that new users that were not observed in the training phase are assessed in the testing phase. We specify when this evaluation method is used.

Table 5
AUC and F1 score for the ISA assessments for all timeframes with the contextual-based approach and the expert-based method.

Measure	Expert-based method	Contextual-based approach
AUC	0.4920	**0.8675**
F1 score	0.6161	**0.8206**

4.2. Evaluation measures

In all of our experimenters we use applicable and realistic timeframes of five minuets. In order to measure our model's performance, we use the following measurements:

- **AUC** — The area under the ROC curve (AUC) is one of the most common measures used for evaluating classification tasks [67,68]. The ROC curve is created by plotting the true positive rate (TPR) against the false positive rate (FPR) for different classification thresholds. Thus, the AUC reflects the average performance of the classifier with different classification thresholds. Note that a higher AUC indicates better performance.
- **F1 Score** — The F1 score measures the harmonic mean of the precision and recall [69]. It is also a very common measure used to evaluate classification tasks, especially in unbalanced problems.

4.3. Results — assessing the ISA of mobile phone users

General performance. In order to demonstrate that the proposed method can be used to assess users' ISA with high accuracy, we apply the time-based estimation evaluation method. The results are reported in Table 5. As can be seen, the proposed method yields high AUC values. The practical implication of the reported performance is that the proposed approach can accurately model changes in the user's context and thus identify, in advance, that a user is more likely to be less aware of cybersecurity threats.

Baseline comparison. To better emphasize the benefits of a context-based, data-driven approach, which considers contextual factors and derives a user's ISA score dynamically for short timeframes, we compare the proposed approach with the expert-based method [3], which derives a *single* ISA score for each user (as described in Section 3.2).

To put the two methods on the same scale, we used the single ISA score derived by the expert-based method for each user as the ISA score of all of the user's timeframes. The results are reported in Table 5. As can be seen, the proposed method dramatically outperforms the expert-based method. These results emphasize the main gap between the current ISA assessment methods and the proposed approach. Namely, users' behavior is too complex to be represented using a single ISA measure. That is, it is quite possible that a user with high security awareness will behave in an unsecure manner in a certain context and vice versa.

Performance relative to different ISA indicators. We are also interested in understanding how different ISA indicators (which are used to calculate the ground truth) affect the performance of the proposed method. Therefore, we calculate the accuracy of the model with respect to the most frequent ISA indicators of unsafe behavior: sending private information in an unsecure manner (private information), visiting a malicious website (malicious website), connecting to an unsecure Wi-Fi network (unsecure Wi-Fi), and clicking on pop-ups or advertisements. The results are presented in Fig. 3.

As can be seen, the proposed method was better at classifying aware timeframes, obtaining a true positive rate (TPR) of 86%.

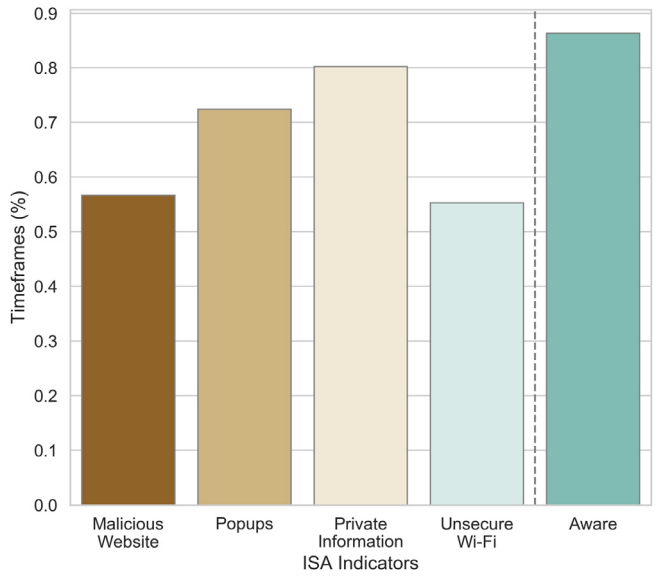

Fig. 3. Classification accuracy with respect to different ISA indicators.

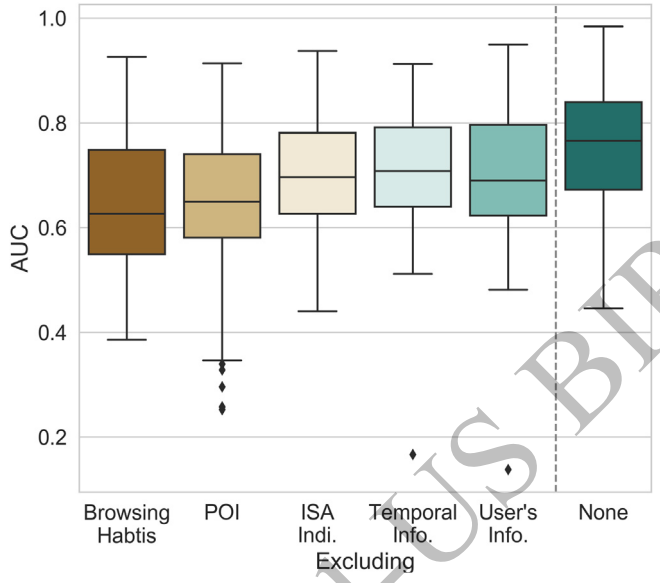

Fig. 4. Boxplot describing the AUC score of all users when excluding different contextual factors.

Exploring the unaware timeframes, the data-driven approach's best performance was observed when classifying timeframes in which the user sent private information in an unsecure manner, obtaining a TPR of 80% in this case, and in the ISA indicator of clicking on pop-ups or advertisements, with a TPR of 72%. The weakest performance was observed for the ISA indicators in which the user connected to an unsecure Wi-Fi network and visited a malicious website, which obtained a TPR of 55% and 56% respectively. We attribute this to the fact that these unsecure behaviors are less common in our dataset (see Table 2). Thus, with less data for these ISA indicators, the data-driven approach was unable to learn the user's patterns, resulting in poor performance. In contrast, the data-driven approach achieved the best performance for the ISA indicator in which the user sent private information in an unsecure manner, which is the most common ISA indicator in the dataset.

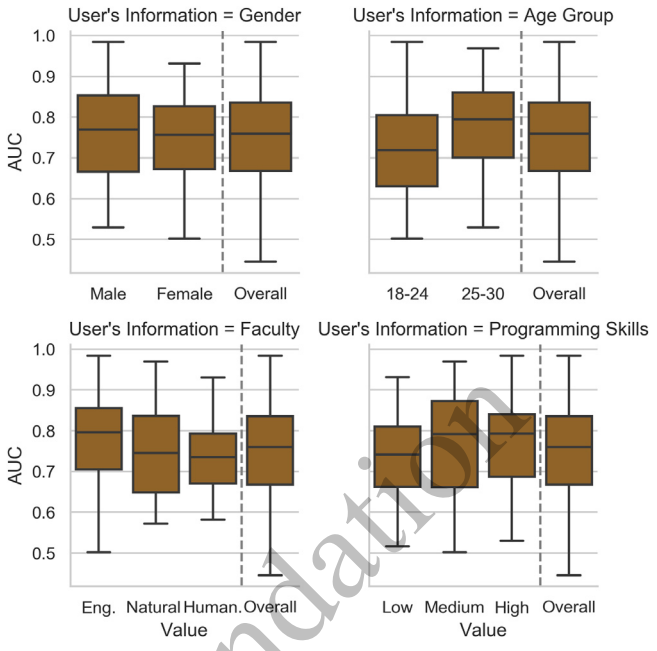

Fig. 5. The AUC performance based on different user attributes.

Contribution of the context. We perform the following experiment to better understand the contribution of the different contextual factors to the ISA assessment: First, we only use the complete set of contextual factors to assess users' ISA. Then, we create a subset of contextual factors by excluding a single contextual factor from the complete set and use that subset to assess users' ISA. The results of this experiment are reported in Fig. 4. This figure presents a boxplot for all users for each excluded contextual factor. We refer to the network's performance based on all contextual factors and ISA indicators as 'None.'

Exploring the results, we can observe that the largest deviation in the AUC performance is found by excluding two contextual factors: 'POI' and 'Browsing Habits.' These factors are able to capture the user's dynamic behavior in the physical space with short timeframes. To verify the significance of these results, we perform the Mann–Whitney U test. In this case, the input for the Mann–Whitney U test is a set of users, where each user is represented by his/her AUC with and without the excluded factor. All results were found to be significant ($p < 0.001$). This emphasizes the contribution of all selected contextual factors for providing accurate ISA assessments.

Ablation study. To better understand the contribution of different components of the suggested deep learning architecture, we assess users' ISA without each of the examined components. More specifically, we examine the necessity of the MLP and the attention layer. When the MLP was not employed, the AUC and F1 score were 0.8602 and 0.8103 respectively, while when the attention layer was not employed, the AUC and F1 score were 0.8596 and 0.8170 respectively. In both cases, we can see that the AUC and F1 score decreased, highlighting the contribution of the MLP and attention layer in obtaining more accurate results for the ISA assessment task.

Performance subject to different user attributes. We examine the performance of the data-driven approach by measuring the AUC based on different user attributes and present the results in Fig. 5. We observe that the smallest AUC differences between user attributes were found for the demographic attribute of gender.

In contrast, the largest differences were found for the age group attribute. To verify the significance of these results, we perform the Mann–Whitney U test, which is a nonparametric test that is not sensitive to the input data distribution type [70]. In our case, the input for the Mann–Whitney U test is a set of users, where each user is represented by his/her AUC based on his/her demographic attribute. The results are significant for the demographic attribute of age group, with a p-value and test statistic of 0.005 and 814 respectively. This analysis could indicate that the data-driven approach was better at learning older users' (between the ages of 25–30) contextual factors to provide more accurate ISA assessments than learning the contextual factors of younger users (between the ages of 18–24).

Additionally, we observe that the AUC of the students with an engineering background is higher than users with a background in the natural sciences or humanities. We also see that the AUC is lower for users with a low level of programming skills than that of more highly-skilled users. However, these differences were not shown to be significant. Focusing on the user's gender, we observe the lowest AUC differences between male users and female users. This result is consistent with a prior study [3] that demonstrated that the difference between the users' gender is not significant for assessing users' ISA.

Assessing the ISA of new users. We also explore the capability of our approach to assess the ISA of new users that were not observed in the training set. In this experiment, we used *user-based cross-validation*, as described in Section 4.1. The results show that the proposed method can be used to assess the ISA of new users, obtaining an average AUC and F1 score of 0.814 and 0.772 respectively for this task. These results indicate that we can employ the data-driven approach to assess new users' ISA with high confidence. Additionally, we report a standard deviation of 0.06 and 0.07 respectively for the AUC and F1 score across different folds when performing user-based cross-validation. This indicates that our method is robust for different users across different folds.

4.4. Comparison with different machine learning models

We perform a comprehensive evaluation by comparing the suggested deep learning framework implementing our approach to (1) an expert system-based method, which is the state-of-the-art and most commonly used solution for assessing users' ISA (the results are presented in Section 4.3), (2) different implementation methods for our approach using traditional machine learning classifiers: decision tree, random forest, logistic regression, AdaBoost, and CatBoost, and (3) three advanced deep learning architectures that were found to be effective for performing related cybersecurity tasks [71,72]:

- **CNN with GRU** [73]. This architecture, which was presented by [73], was shown to be effective in classifying different types of malware (e.g., virus, worm, Trojan) with high accuracy. The CNN with GRU architecture consists of four components: in the first component the authors used a convolutional neural network (CNN). In the second component a gated recurrent unit (GRU) with 20 units was used. In the third component they used a multilayer perceptron. The multilayer perceptron consists of two layers; the first and second layers are composed of ten and five units respectively. In the final component a single neuron with *sigmoid* was used. This architecture was used to classify malware to their respective different families.
- **Deep autoencoder** [74]. The deep autoencoder was used to predict intrusion detection system attacks based on the denial-of-service, user to root, and remote to local and probing categories. We implemented the model presented by [74], which used a deep autoencoder (DAE) and proved to be effective for building an intrusion detection system. The autoencoder consists of seven layers: the first layer is composed of 20 units, the second and the third layers are composed of ten units, the fourth layer is composed of five units, and the fifth and sixth layers are composed of ten units. The final layer is composed of a single neuron with *sigmoid*.
- **GRU with SVM** [75]. We used the architecture introduced by [75] for the classification of the intrusion detection task, based on network traffic data from the honeypot systems of Kyoto University. This architecture consists of a GRU that composed of 20 units and a support vector machine (SVM) on top of it.

Table 6
ML methods and the proposed deep learning architecture for assessing users' ISA.

Model	AUC (SD)	F1 score (SD)
Decision tree	0.6242 (0.06)	0.6854 (0.14)
Random forest	0.6876 (0.09)	0.7741 (0.13)
Logistic regression	0.7406 (0.11)	0.7017 (0.17)
CatBoost	0.7674 (0.11)	0.7550 (0.17)
AdaBoost	0.6215 (0.06)	0.6897 (0.14)
CNN with GRU	0.8284 (0.11)	0.8045 (0.11)
Deep autoencoder	0.8349 (0.11)	0.7713 (0.15)
GRU with SVM	0.8245 (0.11)	0.8058 (0.13)
Proposed architecture	**0.8675** (0.10)	**0.8206** (0.10)

In order to set a fair comparison, we optimize all deep learning models with the same optimization framework [66] used for optimizing the suggested deep learning architecture. Furthermore, to find the optimal parameters we use the same range of values as examined for the suggested deep learning architecture (mentioned in Section 4.1). We also use the same contextual factors and previously observed ISA indicators as input, i.e., the input for all of the ML methods and deep learning architectures described above was a vector, which is the result of a concatenation of all of the contextual factors and ISA indicators mentioned in Section 3.4. We present the AUC and F1 score achieved with the various deep learning architectures in Table 6. In the table we also present the standard deviation (SD) of each measurement based on the performance of all users.

In the table, we can observe that the best performance was achieved by the DNN models: CNN with GRU, deep autoencoder, GRU with SVM and the proposed architecture. These results indicate that users' ISA can be predicted with high accuracy using deep learning models based on contextual information.

In order to validate the significance of the results, we perform statistical tests based on the AUC obtained for each user in the test set. First, we perform the Friedman test, which is a nonparametric statistical test. The null hypothesis of this test is that multiple paired samples have the same distribution, i.e., all models will obtain similar AUC performance. We reject the null hypothesis for the test with a p-value that is less than 0.001 and a test statistic of 649. Next, we use the Bonferroni–Dunn test to determine whether our framework performed significantly better than all of the other baselines examined. The results of this test were found to be significant with a p-value of 0.01, except with the CNN with GRU architecture where the results were not found to be significant.

Similarly, we perform the Friedman test based on the users' F1 score and reject the null hypothesis with a p-value that is less than 0.001 and a test statistic of 439. Then, we perform the Bonferroni–Dunn test; the results of the test were found

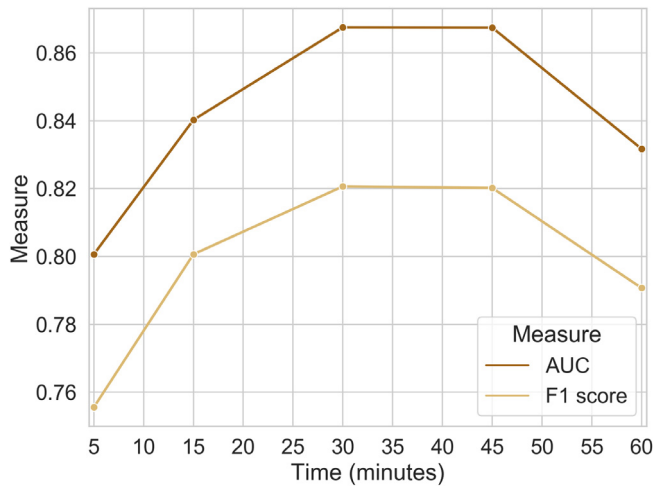

Fig. 6. AUC and F1 score for different learning periods before providing an ISA assessment.

to be significant with a *p*-value of 0.05, except with the CNN with GRU architecture where the results were not found to be significant. Based on the statistical tests' results, we can conclude the our method statistically outperformed all baselines, except for the CNN with GRU architecture. This could be explained by the fact that in both our suggested architecture and the CNN with GRU architecture we used the same GRU component to model the time series based contextual information. Nevertheless, our architecture performed better than the CNN with GRU architecture. These results highlight the ability of the proposed architecture to capture all of the contextual factors based on their type: sequential or static. Thus, the architecture can be used to capture various hidden patterns and connections between the contextual factors and the user's ISA.

4.5. Evaluating number of learning periods

We examine different learning periods for the task of providing an ISA assessment and present the results in Fig. 6. For example, for a given user, learning for a 15 min period (three timeframes of five minutes each) before providing his/her ISA assessment for the next five minutes results in AUC and F1 score of 0.80 and 0.75 respectively.

When examining the results further, we observe that the optimal learning time is 30 min. Using less learning time (five or 15 min) results in the lowest AUC and F1 score. Therefore, we conclude that for a short learning period, the user's behavior with respect to his/her context is not captured well by the deep learning architecture. Using more than 30 min results in less accurate ISA assessments, with a small decrease in the AUC and F1 score. This could be explained by the fact that the user's behavior is recurrent over long learning time periods, thus the deep learning architecture cannot capture new information about the user. An interesting phenomenon observed for the longest learning time of 90 min, is the large decrease found in the AUC and F1 score. We conclude that for long learning periods, some of the data is very repetitive, which can be interpreted by the model as noise and make the behavior patterns less recognizable, thus harming the architecture's performance.

5. Discussion

In this section, we discuss the novelty of the context-based, data-driven approach and suggest different applications for continuous context-based ISA assessment.

5.1. Providing ISA assessments with a context-based, data-driven approach

In this study, we propose a novel context-based, data-driven approach for evaluating users' ISA. We showed that today's most common approach [3], static scores, are insufficient for evaluating users' ISA, since in our evaluation they performed the poorest, achieving the lowest AUC, and F1 results. These results emphasize the gap between the current ISA assessment methods and our proposed approach. We also demonstrated that our proposed approach reflects the users' actual behavior, achieving the highest results. This can be explained by the fact that the context-based, data-driven approach is able to capture users' dynamic behavior unlike traditional ISA scores, which are fixed and static. These results also indicate that the user's ISA can be learned based on his/her contextual information, e.g., the user's location, etc.

We also confirmed that browsing habits and POIs learned from the user's routine are crucial for evaluating users' ISA and that a half an hour is the optimal amount of learning time for evaluating a user's ISA in the next five minutes. These findings are important, since they can serve as the basis of future mobile phone sampling strategies and solutions.

Overall, all deep learning architectures achieved high performance for assessing users' ISA. This can indicate that deep learning based methods can capture hidden patterns between contextual factors and users' ISA. Moreover, these results highlight the use of our approach to use contextual information for assessing users' ISA. The high results achieved by the proposed deep learning architecture can be explained by the fact that the proposed architecture is designed specifically for the task of predicting users' ISA, i.e., we model each contextual factor with its own representation, based on the information type. For example, we model the browsing habits, POIs, and ISA indicators with a GRU and attention layer, since they are vary over time and are considered sequential information. Furthermore, we can see that the proposed architecture can also be used successfully for predicting the ISA of new users and that this architecture is robust for different users.

5.2. Applications

5.2.1. Personalized ISA training programs

Security awareness training programs are educational workshops that equip employees with tools to identify, mitigate, and report social engineering attacks [76]. These training programs utilize creative methods to deliver security awareness content, such as attack simulations, interactive web modules, videos, games, posters, and newsletters [77,78]. The social engineering attack landscape is rapidly evolving, with novel attack vectors that utilize different psychological manipulation techniques. With thousands of training modules available, enterprises are faced with two primary challenges: a limited training budget [76] and employees' lack of personal motivation for training [2]. The first challenge stems from the fact that user training is a very time-consuming task that is often performed during working hours. Consequently, the enterprise must allocate a dedicated budget for this task. The second challenge arises because training workshops are not tailored to the individual needs of each employee. Employees are often required to participate in training workshops that are not necessarily correlated to their vulnerabilities, resulting in user fatigue and reduced motivation [76]. Using the proposed context-based, data-driven approach, enterprises can continuously identify the contexts in which a person is more likely to be vulnerable. Furthermore, the proposed approach allows enterprises to identify which attack vectors an individual user is vulnerable to. With this knowledge, enterprises can tailor

training workshops to the individual needs of each employee. For instance, instead of requiring all employees to participate in the same training workshops, enterprises can wisely select the relevant training workshop for each employee, thus reducing training costs and user fatigue, resulting in more effective training.

5.2.2. Adaptive IT security policy

The information technology (IT) security policy defines the rules and procedures for accessing and using an enterprise's IT assets and resources. It is one of the most important controls available for managing and ensuring information security effectiveness. An effective IT security policy must consider users' needs for available, accurate, and reliable information, as well as an enterprise's need to secure its IT assets [79]. To enforce an IT security policy in real-time, enterprises can use security countermeasures, such as firewalls, network intrusion detection/prevention systems, and access control lists. Traditionally, IT security policies are static and do not change based on the context of use. For instance, when enforcing a network security policy using a firewall, the administrator creates rules to prevent/permit the inbound/outbound communication from/to a specific port and protocol. Using the proposed context-based, data-driven approach, enterprises can identify the contexts in which a person is more likely to be vulnerable to a social engineering attack. By knowing that, administrators can implement an adaptive security policy. For instance, when the system identifies a context in which an employee is more likely to be vulnerable, the administrator can adjust an employee's permissions, taking back sensitive permissions granted to the employee, thus improving system security.

5.3. Limitations

The main limitation of our approach is its dependence on data for discovering meaningful behavioral patterns and thus for providing accurate ISA assessments. To demonstrate it in our case, we explore the classification accuracy for each timeframe, based on the ISA indicators. The TPR is 80% for the private information indicator, which is the most frequent indicator in our data (88.1%), and only 55% for the connection to an unsecured Wi-Fi network, which is only present in 3.2% of the data. However, this limitation could be addressed by collecting data both from more users and for a longer period of time.

In this study we use the dataset derived from the work of Bitton et al. [3]. While that data enables us to examine the use of our approach to assess ISA based on contextual information, it has two main limitations: (1) The population examined in this research is quite diverse in terms of activities and usage patterns, i.e., users used different mobile applications and visited in different POIs. However, the examined population age range is limited only to 18–30 years old. Thus, we are not able to generalize our method with the other age groups. (2) Using the collected dataset we are only able to examine our method with individuals. Thus, we are not able to present results and draw conclusions from a business perspective. Therefore, a future study should employ our method on users with other age ranges and as part of an organization training program.

6. Conclusions

In this paper, we proposed a novel context-based, data-driven approach for assessing users' ISA. The results emphasize the existing gap between our context-based, data-driven approach and most common approaches for ISA assessment, which provide fixed scores based on users' reported behaviors and are unable to capture the users' dynamic behavior. In contrast, our approach can capture users' dynamic behavior with respect to the users' context, enabling it to provide the most accurate ISA

Table A.7
Set of security questions used to measure the likelihood of performing an action.

#	What is the likelihood that you would perform the following actions?
1	Download an application from an unofficial application store.
2	Install an application that requires permissions that are not necessary for its functionality.
3	Install an applications with a low rating.
4	Install an application that requires root privileges.
5	Approve an application update that requires permissions that are not necessary for the application's functionality.
6	Verify an application update before approving it.
7	Click on an advertisement when using an application.
8	Click on an advertisement for a lottery.
9	Check which applications are installed on the device.
10	Check which applications are running.
11	Close applications that are running in the background.
12	Delete applications that are not in use.
13	Enter a website despite a security warning indicating that the site is dangerous (figure attached).
14	Download a file from a site that does not use an encryption protocol.
15	Insert private information on a site which does not use an encryption protocol.
16	Use your personal password on a site which does not use an encryption protocol.
17	Enter private information (e.g., phone number, email address) into a pop-up that appears when using an application.
18	Open an email classified as spam.
19	Enter a link sent from an unknown party (e.g., via Facebook, WhatsApp, SMS).
20	Download a file sent to you by email from an unknown sender.
21	Use a simple password that contains known personal details (e.g., name, date of birth, phone number).
22	Use a password that is constructed of many different digits and characters.
23	Use the same password for different services.
24	Save your password as plain-text on your device (e.g., in your contacts, documents, notes).
25	Update your password when you suspect that someone knows it.
26	Update your password at least once a year.
27	Use two-factor authentication for email/Facebook (a service in which another form of authentication is used besides the password).
28	Use password management services.
29	Jailbreak/root your device.
30	Use embedded security systems (e.g., firewall, encryption).
31	Install an anti-virus application.
32	Ignore a security alert.
33	Use screen lock (e.g., PIN code, pattern, fingerprint).
34	Use public (unencrypted) Wi-Fi networks.
35	Download files when using a public network.
36	Use a VPN service when using the Internet.
37	Use personal services (e.g., banking, online shopping, Facebook) when connected to a public Wi-Fi network.
38	Auto connect to blacktooth devices.
39	Connect your smartphone to foreign device (e.g., a friend's computer, wireless headphones).

assessments and outperform the traditional expert-based methods. In addition, the proposed deep learning framework, which implements our data-driven approach, outperformed all other ML classifiers and deep learning architectures. The data-driven approach was also shown to be beneficial, providing accurate ISA assessments in the common application scenario in which new users are introduced. An analysis of the ISA assessments derived from the data-driven approach reveals some interesting findings about users that have not been observed before, including the observations that half an hour is the optimal learning time for evaluating the user's ISA in the next five minutes, and the ISA assessments of users in the 25–30 age group are more accurate compared to those in the 18–24 age group.

CRediT authorship contribution statement

Adir Solomon: Conceptualization, Methodology, Software, Investigation, Validation, Formal analysis, Visualization, Writing – original draft. **Michael Michaelshvili:** Software, Investigation, Validation, Formal analysis. **Ron Bitton:** Conceptualization, Methodology, Formal analysis, Writing – original draft. **Bracha Shapira:** Conceptualization, Methodology, Writing – review & editing, Supervision. **Lior Rokach:** Conceptualization, Methodology, Writing – review & editing, Supervision. **Rami Puzis:** Conceptualization, Methodology, Writing – review & editing, Supervision. **Asaf Shabtai:** Conceptualization, Methodology, Writing – review & editing, Supervision.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix. Questionnaire

Table A.7 presents the security questionnaire used to measure the likelihood that the user will perform a certain action. All of the questions were measured according to a five-point Likert scale with the following answers: "Never", "Unlikely", "Medium Likelihood", "Very Likely", and "Always".

References

[1] D. Ki-Aries, S. Faily, Persona-centred information security awareness, Comput. Secur. 70 (2017) 663–674.

[2] M.I. Mann, Hacking the Human: Social Engineering Techniques and Security Countermeasures, Gower Publishing, Ltd., 2012.

[3] R. Bitton, K. Boymgold, R. Puzis, A. Shabtai, Evaluating the information security awareness of smartphone users, in: Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, 2020, pp. 1–13.

[4] F. Mouton, L. Leenen, H.S. Venter, Social engineering attack examples, templates and scenarios, Comput. Secur. 59 (2016) 186–209.

[5] K. Parsons, A. McCormac, M. Butavicius, M. Pattinson, C. Jerram, Determining employee awareness using the human aspects of information security questionnaire (HAIS-Q), Comput. Secur. 42 (2014) 165–176.

[6] A. McCormac, D. Calic, K. Parsons, T. Zwaans, M. Butavicius, M. Pattinson, Test-retest reliability and internal consistency of the human aspects of information security questionnaire (HAIS-Q), 2016.

[7] R. Wash, E. Rader, C. Fennell, Can people self-report security accurately?: Agreement between self-report and behavioral measures, in: Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems, ACM, 2017, pp. 2228–2232.

[8] E.M. Redmiles, Z. Zhu, S. Kross, D. Kuchhal, T. Dumitras, M.L. Mazurek, Asking for a friend: Evaluating response biases in security user studies, in: Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security, ACM, 2018, pp. 1238–1255.

[9] P. Kumaraguru, J. Cranshaw, A. Acquisti, L. Cranor, J. Hong, M.A. Blair, T. Pham, School of phish: a real-world evaluation of anti-phishing training, in: Proceedings of the 5th Symposium on Usable Privacy and Security, 2009, pp. 1–12.

[10] K. Jansson, R. von Solms, Phishing for phishing awareness, Behav. Inf. Technol. 32 (6) (2013) 584–593.

[11] R.W. White, P. Bailey, L. Chen, Predicting user interests from contextual information, in: Proceedings of the 32nd International ACM SIGIR Conference on Research and Development in Information Retrieval, 2009, pp. 363–370.

[12] F. Long, Improved personalized recommendation algorithm based on context-aware in mobile computing environment, Wirel. Commun. Mob. Comput. 2020 (2020).

[13] W. Liu, X. Li, D. Huang, A survey on context awareness, in: 2011 International Conference on Computer Science and Service System (CSSS), IEEE, 2011, pp. 144–147.

[14] K. Cho, B. Van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, Y. Bengio, Learning phrase representations using RNN encoder-decoder for statistical machine translation, 2014, arXiv preprint arXiv:1406.1078.

[15] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A.N. Gomez, Ł. Kaiser, I. Polosukhin, Attention is all you need, in: Advances in Neural Information Processing Systems, 2017, pp. 5998–6008.

[16] V. Gkioulos, G. Wangen, S.K. Katsikas, G. Kavallieratos, P. Kotzanikolaou, Security awareness of the digital natives, Information 8 (2) (2017) 42.

[17] V. Gkioulos, G. Wangen, S.K. Katsikas, User modelling validation over the security awareness of digital natives, Future Internet 9 (3) (2017) 32.

[18] I. Androulidakis, G. Kandus, Bluetooth® usage among students as an indicator of security awareness and feeling, in: Proceedings ELMAR-2011, IEEE, 2011, pp. 157–160.

[19] K. Onarlioglu, U.O. Yilmaz, E. Kirda, D. Balzarotti, Insights into user behavior in dealing with internet attacks, in: NDSS, 2012.

[20] A. Mylonas, A. Kastania, D. Gritzalis, Delegate the smartphone user? Security awareness in smartphone platforms, Comput. Secur. 34 (2013) 47–66.

[21] S. Egelman, E. Peer, Scaling the security wall: Developing a security behavior intentions scale (sebis), in: Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems, 2015, pp. 2873–2882.

[22] P. Kumaraguru, Y. Rhee, A. Acquisti, L.F. Cranor, J. Hong, E. Nunge, Protecting people from phishing: the design and evaluation of an embedded training email system, in: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, 2007, pp. 905–914.

[23] A. McCormac, T. Zwaans, K. Parsons, D. Calic, M. Butavicius, M. Pattinson, Individual differences and information security awareness, Comput. Hum. Behav. 69 (2017) 151–156.

[24] A. Wiley, A. McCormac, D. Calic, More than the individual: Examining the relationship between culture and information security awareness, Comput. Secur. 88 (2020) 101640.

[25] D. Dang-Pham, S. Pittayachawan, V. Bruno, Applications of social network analysis in behavioural information security research: Concepts and empirical analysis, Comput. Secur. 68 (2017) 1–15.

[26] D. Dang-Pham, S. Pittayachawan, Comparing intention to avoid malware across contexts in a BYOD-enabled Australian university: A protection motivation theory approach, Comput. Secur. 48 (2015) 281–297.

[27] M. Karyda, E. Kiountouzis, S. Kokolakis, Information systems security policies: a contextual perspective, Comput. Secur. 24 (3) (2005) 246–260.

[28] P. Ifinedo, An exploratory study of the relationships between selected contextual factors and information security concerns in global financial services institutions, J. Inf. Priv. Secur. 7 (1) (2011) 25–49.

[29] D. Canali, L. Bilge, D. Balzarotti, On the effectiveness of risk prediction based on users browsing behavior, in: Proceedings of the 9th ACM Symposium on Information, Computer and Communications Security, ACM, 2014, pp. 171–182.

[30] M. Aburrous, M.A. Hossain, K. Dahal, F. Thabtah, Predicting phishing websites using classification mining techniques with experimental case studies, in: 2010 Seventh International Conference on Information Technology: New Generations, IEEE, 2010, pp. 176–181.

[31] S. Gupta, A. Singhal, Dynamic classification mining techniques for predicting phishing URL, in: Soft Computing: Theories and Applications, Springer, 2018, pp. 537–546.

[32] I. Tjostheim, J.A. Waterworth, Predicting personal susceptibility to phishing, in: International Conference on Information Technology & Systems, Springer, 2020, pp. 564–575.

[33] M. Sharif, J. Urakawa, N. Christin, A. Kubota, A. Yamada, Predicting impending exposure to malicious content from user behavior, in: Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security, 2018, pp. 1487–1501.

[34] F. Foroughi, P. Luksch, A multi-agent model for security awareness driven by home user's behaviours, in: Proceedings of the Future Technologies Conference, Springer, 2018, pp. 185–195.

[35] V.S. Saridewi, R.F. Sari, Implementation of machine learning for human aspect in information security awareness, J. Appl. Eng. Sci. 19 (4) (2021) 1126–1142.

[36] W. Shafqat, Y.-C. Byun, A context-aware location recommendation system for tourists using hierarchical LSTM model, Sustainability 12 (10) (2020) 4107.
[37] A. Livne, E.S. Tov, A. Solomon, A. Elyasaf, B. Shapira, L. Rokach, Evolving context-aware recommender systems with users in mind, Expert Syst. Appl. 189 (2022) 116042.
[38] C. Huang, J. Zhang, Y. Zheng, N.V. Chawla, DeepCrime: Attentive hierarchical recurrent networks for crime prediction, in: Proceedings of the 27th ACM International Conference on Information and Knowledge Management, 2018, pp. 1423–1432.
[39] T. Soikkeli, J. Karikoski, H. Hammainen, Diversity and end user context in smartphone usage sessions, in: 2011 Fifth International Conference on Next Generation Mobile Applications, Services and Technologies, IEEE, 2011, pp. 7–12.
[40] A. Tammewar, A. Cervone, E.-M. Messner, G. Riccardi, Modeling user context for valence prediction from narratives, 2019, arXiv preprint arXiv:1905.05701.
[41] R. Bitton, A. Finkelshtein, L. Sidi, R. Puzis, L. Rokach, A. Shabtai, Taxonomy of mobile users' security awareness, Comput. Secur. 73 (2018) 266–293.
[42] D. Damopoulos, G. Kambourakis, S. Gritzalis, iSAM: an iPhone stealth airborne malware, in: IFIP International Information Security Conference, Springer, 2011, pp. 17–28.
[43] N. Virvilis, N. Tsalis, A. Mylonas, D. Gritzalis, Mobile devices: A phisher's paradise, in: 2014 11th International Conference on Security and Cryptography (SECRYPT), IEEE, 2014, pp. 1–9.
[44] C. Brubaker, S. Jana, B. Ray, S. Khurshid, V. Shmatikov, Using frankencerts for automated adversarial testing of certificate validation in SSL/TLS implementations, in: 2014 IEEE Symposium on Security and Privacy, IEEE, 2014, pp. 114–129.
[45] Z. Li, G. Xiong, L. Guo, Unveiling SSL/TLS MITM hosts in the wild, in: 2020 IEEE 3rd International Conference on Information Systems and Computer Aided Education (ICISCAE), IEEE, 2020, pp. 141–145.
[46] J. Du, X. Li, H. Huang, A study of man-in-the-middle attack based on SSL certificate interaction, in: 2011 First International Conference on Instrumentation, Measurement, Computer, Communication and Control, IEEE, 2011, pp. 445–448.
[47] A. Ranjbar, M. Komu, P. Salmela, T. Aura, An SDN-based approach to enhance the end-to-end security: SSL/TLS case study, in: NOMS 2016-2016 IEEE/IFIP Network Operations and Management Symposium, IEEE, 2016, pp. 281–288.
[48] J.H. Park, K.J. Yi, Y.-S. Jeong, An enhanced smartphone security model based on information security management system (ISMS), Electron. Commer. Res. 14 (3) (2014) 321–348.
[49] J. Beekman, C. Thompson, Man-in-the-middle attack on T-Mobile Wi-Fi Calling, in: Electrical Engineering and Computer Sciences, University of California at Berkeley, 2013, https://www2.eecs.berkeley.edu/Pubs/TechRpts/2013/EECS-2013-18.pdf.
[50] Y. Zhou, X. Jiang, Dissecting android malware: Characterization and evolution, in: 2012 IEEE Symposium on Security and Privacy, IEEE, 2012, pp. 95–109.
[51] S. Peng, S. Yu, A. Yang, Smartphone malware and its propagation modeling: A survey, IEEE Commun. Surv. Tutor. 16 (2) (2013) 925–941.
[52] H.A. Kruger, W.D. Kearney, A prototype for assessing information security awareness, Comput. Secur. 25 (4) (2006) 289–296.
[53] I. Reyes, P. Wijesekera, J. Reardon, A. Elazari Bar On, A. Razaghpanah, N. Vallina-Rodriguez, S. Egelman, et al., "Won't somebody think of the children?" examining COPPA compliance at scale, in: The 18th Privacy Enhancing Technologies Symposium (PETS 2018), 2018.
[54] E. Pan, J. Ren, M. Lindorfer, C. Wilson, D.R. Choffnes, Panoptispy: Characterizing audio and video exfiltration from android applications., Proc. Priv. Enhanc. Technol. 2018 (4) (2018) 33–50.
[55] A. Majid, L. Chen, G. Chen, H.T. Mirza, I. Hussain, J. Woodward, A context-aware personalized travel recommendation system based on geotagged social media data mining, Int. J. Geogr. Inf. Sci. 27 (4) (2013) 662–684.
[56] G. Adomavicius, A. Tuzhilin, Context-aware recommender systems, in: Recommender Systems Handbook, Springer, 2011, pp. 217–253.
[57] OpenStreetMap contributors, 2017, Planet dump retrieved from https://planet.osm.org, https://www.openstreetmap.org.
[58] A. Singh, N. Goyal, A comparison of machine learning attributes for detecting malicious websites, in: 2019 11th International Conference on Communication Systems & Networks (COMSNETS), IEEE, 2019, pp. 352–358.
[59] C. Singh, et al., Phishing website detection based on machine learning: A survey, in: 2020 6th International Conference on Advanced Computing and Communication Systems (ICACCS), IEEE, 2020, pp. 398–404.
[60] K. Cho, B. Van Merriënboer, D. Bahdanau, Y. Bengio, On the properties of neural machine translation: Encoder-decoder approaches, 2014, arXiv preprint arXiv:1409.1259.
[61] G. Kobayashi, T. Kuribayashi, S. Yokoi, K. Inui, Attention module is not only a weight: Analyzing transformers with vector norms, 2020, arXiv preprint arXiv:2004.10102.
[62] T. Mikolov, I. Sutskever, K. Chen, G.S. Corrado, J. Dean, Distributed representations of words and phrases and their compositionality, in: Advances in Neural Information Processing Systems, 2013, pp. 3111–3119.
[63] A. Solomon, A. Bar, C. Yanai, B. Shapira, L. Rokach, Predict demographic information using word2vec on spatial trajectories, in: Proceedings of the 26th Conference on User Modeling, Adaptation and Personalization, 2018, pp. 331–339.
[64] D. Wang, P. Cui, W. Zhu, Structural deep network embedding, in: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 2016, pp. 1225–1234.
[65] F. Schroff, D. Kalenichenko, J. Philbin, Facenet: A unified embedding for face recognition and clustering, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2015, pp. 815–823.
[66] J. Bergstra, D. Yamins, D. Cox, Making a science of model search: Hyperparameter optimization in hundreds of dimensions for vision architectures, in: International Conference on Machine Learning, PMLR, 2013, pp. 115–123.
[67] T. Fawcett, An introduction to ROC analysis, Pattern Recognit. Lett. 27 (8) (2006) 861–874.
[68] D.J. Hand, R.J. Till, A simple generalisation of the area under the ROC curve for multiple class classification problems, Mach. Learn. 45 (2) (2001) 171–186.
[69] C. Goutte, E. Gaussier, A probabilistic interpretation of precision, recall and F-score, with implication for evaluation, in: European Conference on Information Retrieval, Springer, 2005, pp. 345–359.
[70] S. Yue, C. Wang, Power of the Mann–Whitney test for detecting a shift in median or mean of hydro-meteorological data, Stoch. Environ. Res. Risk Assess. 16 (4) (2002) 307–323.
[71] Y. Xin, L. Kong, Z. Liu, Y. Chen, Y. Li, H. Zhu, M. Gao, H. Hou, C. Wang, Machine learning and deep learning methods for cybersecurity, IEEE Access 6 (2018) 35365–35381.
[72] S. Mahdavifar, A.A. Ghorbani, Application of deep learning to cybersecurity: A survey, Neurocomputing 347 (2019) 149–176.
[73] C.H. Kim, E.K. Kabanga, S.-J. Kang, Classifying malware using convolutional gated neural network, in: 2018 20th International Conference on Advanced Communication Technology (ICACT), IEEE, 2018, pp. 40–44.
[74] F. Farahnakian, J. Heikkonen, A deep auto-encoder based approach for intrusion detection system, in: 2018 20th International Conference on Advanced Communication Technology (ICACT), IEEE, 2018, pp. 178–183.
[75] A.F.M. Agarap, A neural network architecture combining gated recurrent unit (GRU) and support vector machine (SVM) for intrusion detection in network traffic data, in: Proceedings of the 2018 10th International Conference on Machine Learning and Computing, 2018, pp. 26–30.
[76] H. Aldawood, G. Skinner, Reviewing cyber security social engineering training and awareness programs—Pitfalls and ongoing issues, Future Internet 11 (3) (2019) 73.
[77] A.U. Zulkurnain, A. Hamidy, A.B. Husain, H. Chizari, Social engineering attack mitigation, Int. J. Math. Comput. Sci. 1 (4) (2015) 188198.
[78] O.J. Olusegun, N.B. Ithnin, People are the answer to security: Establishing a sustainable information security awareness training (ISAT) program in organization, 2013, arXiv preprint arXiv:1309.0188.
[79] K. Höne, J. Eloff, What makes an effective information security policy? Netw. Secur. 2002 (6) (2002) 14–16.