Computers in Human Behavior 48 (2015) 51–61

Contents lists available at ScienceDirect

Computers in Human Behavior

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

Effects of cyber security knowledge on attack detection

Noam Ben-Asher *, Cleotilde Gonzalez

Dynamic Decision Making Laboratory, Department of Social and Decision Sciences, Carnegie Mellon University, PA, USA

ARTICLE INFO

Article history:

Keywords:
Cyber security
Knowledge
Dynamic decision-making
Intrusion-detection system
Expertise

ABSTRACT

Ensuring cyber security is a complex task that relies on domain knowledge and requires cognitive abilities to determine possible threats from large amounts of network data. This study investigates how knowledge in network operations and information security influence the detection of intrusions in a simple network. We developed a simplified Intrusion Detection System (IDS), which allows us to examine how individuals with or without knowledge in cyber security detect malicious events and declare an attack based on a sequence of network events. Our results indicate that more knowledge in cyber security facilitated the correct detection of malicious events and decreased the false classification of benign events as malicious. However, knowledge had less contribution when judging whether a sequence of events representing a cyber-attack. While knowledge of cyber security helps in the detection of malicious events, situated knowledge regarding a specific network at hand is needed to make accurate detection decisions. Responses from participants that have knowledge in cyber security indicated that they were able to distinguish between different types of cyber-attacks, whereas novice participants were not sensitive to the attack types. We explain how these findings relate to cognitive processes and we discuss their implications for improving cyber security.

1. Introduction

Cyber-attacks—the disruption of computers' normal functioning and the loss of sensitive information through malicious network events—are becoming more widespread. Guarding against them is a significant part of the Information Technology (IT) governance done by cyber analysts, as many government agencies and private companies have moved to distributed systems (McHugh, 2001). The most important responsibility of a cyber-security analyst is to protect a network from harm. Many technological advances in information and network security have facilitated the advanced monitoring and threat detection for the analysts, but the tasks they perform cannot be completely automated. The analytical capabilities of the human decision maker are still needed and are indispensable (Cranor, 2008; Jajodia, Liu, Swarup, & Wang, 2010). However, although analysts are capable of performing cyber security tasks, our understanding of the cognitive processes that are required for effective network protection is relatively limited (Chen, Liu, Yen, & Mullen, 2012; Gonzalez, Ben-Asher, Oltramari, & Lebiere, 2014). Furthermore, it is unclear in what ways the analysts utilize their experience in cyber security to detect cyber-attacks.

One tool that security analysts heavily rely on is Intrusion Detection System (IDS). This tool can detect network intrusions and network misuse by matching patterns of known attacks against ongoing network activity. Once the IDS finds a match to a known type of attack or detects abnormal network activity, it produces alerts detailing the suspicious events (Goodall, Lutters, & Komlodi, 2009). In IDS, as in other alert systems, decreasing the number of missed events increases the number of false alerts (Green & Swets, 1966). Considering the amount of traffic in a mid-size corporate network and the ever-growing number and complexity of cyber-attacks, the number of alerts generated by an IDS can be overwhelming to a human analyst. Such systems can trigger thousands of alerts per day, up to 99% of which are false alerts (Goodall, Lutters, & Komlodi, 2004). Eventually, the high volume of intrusion alerts that needs to be processed and the high probability of false alerts make the process of accurately detecting a cyber-attack challenging for human cognitive capabilities.

There is a growing body of work within the cyber security field that is focused on understanding the work processes of security analysts (D'Amico et al., 2005; Goodall et al., 2009; Thompson, Rantanen, & Yurcik, 2006; Werlinger, Muldner, Hawkey, & Beznosov, 2010). Previous studies infer that the general cyber analysis work process model includes preparation, monitoring,

* Corresponding author at: Dynamic Decision Making Laboratory, Department of Social and Decision Sciences, Carnegie Mellon University, 4609 Winthrop Street, Pittsburgh, PA 15213, USA. Tel.: +1 412 268 9547.

E-mail addresses: noamba@cmu.edu (N. Ben-Asher), coty@cmu.edu (C. Gonzalez).

http://dx.doi.org/10.1016/j.chb.2015.01.039
0747-5632/© 2015 Elsevier Ltd. All rights reserved.

detection, analysis, and response to network events. Both monitoring and detection belong to a general process called *triage analysis*. When conducting triage analysis, the analyst screens a large number of IDS alerts and network events, identifies false alerts, and escalates suspicious events for further analysis, which can result in the appropriate response (D'Amico et al., 2005). Triage analysis is a knowledge-intensive activity in which an analyst's expertise is leveraged to promptly dismiss false alerts and to attend to alerts that provide true indications of a cyber-attack.

In this study, we investigate the basic cognitive processes involved in the detection of cyber-attacks with a specific interest in understanding the interplay between domain knowledge and cognitive skills. As one cannot play chess without knowing the rules of the game, some specific knowledge is required to detect cyber-attacks. Cyber security analysts and practitioners are required to have a broad knowledge of network operation and information security. They usually undergo extensive training and certification programs. However, it is not clear whether acquiring deep and detailed knowledge in cyber security is the main determinant of performance when detecting cyber-attacks or whatever the ability to efficiently apply general thinking strategies is at least equally crucial to this task. Furthermore, it is still unclear how aspects like information search and evidence accumulation, which serve as a basis for the detection of cyber-attacks, depend on the analyst's domain knowledge and on a general set of cognitive skills she apply (Perkins & Salomon, 1989). As the security analyst operates in a highly dynamic environment, domain knowledge can be incomplete or become outdated relatively fast. This type of environment highlights the dependability on thinking strategies for problem solving, inventive thinking, decision making, and learning. Thus, it is possible that mastering independent cognitive skills in such a context is a main component of cyber security expertise.

As an initial step in resolving these questions, we examine how the knowledge gap between experts and novices in cyber security influences their ability to detect cyber-attacks. A questionnaire allowed us to corroborate participants' knowledge in information and network security. Using a simplified IDS tool, we then conduct laboratory and online experiments with experienced individuals in cyber security and with participants with no significant knowledge of cyber security. We examined the intrusion detection process in different contexts (i.e., network scenarios) by presenting several types of cyber-attacks. For each network scenario, the intrusion detection process had two parts: the first included classification of network events as malicious or benign; and in the second part, a decision was made about whether or not the whole sequence of network events represents an ongoing cyber-attack. This allowed us to further examine the role of experience in different stages of the detection task. Overall, we predicted that a larger knowledge base would lead to better performance, and that experts would do better than novices that can only rely on their general cognitive skills. Therefore, we hypothesized that experts will be more accurate than novices, when judging a whole sequence of network events, and detecting a cyber-attack. We also expect that experts will decide more accurately whether a network event is malicious or not. Finally, we hypothesized that when judging a sequence of network events experts will be more confident in their decisions compared to novice. These differences, between experts and novices, are expected to be consistent across different network scenarios.

2. Knowledge and cognitive challenges of cyber security

The rate and the extent to which the cyberspace can change is extremely variable and unpredictable compared to other environments that are bound by physical constraints. The topology of the network, the services it provides, and the users who depend on these services are constantly changing. In parallel, new vulnerabilities that can be exploited continuously emerge, clever attack strategies are constantly developed and new counteracting protective measures are deployed. These challenges result in a continuous effort by the cyber security analyst to stay up-to-date on the knowledge needed to successfully defend a network.

An analyst continually monitors the network, identifies threats, and repairs each and any vulnerability; while the attacker only needs to find a single vulnerability that can be exploited (Yurcik, Barlow, & Rosendale, 2003). This simplified view highlights the asymmetric relationships between a security analyst, a complex environment, and an attacker. An analyst is constantly required to make multiple and interdependent decisions in a dynamic environment. Dynamic decision making is highly complex because it requires an understanding of multiple, interrelated attributes and the ability to anticipate the way that the environment will develop over time. A decision maker is also required to act at the right time to maximize the decision value (Brehmer, 1992; Edwards, 1962; Gonzalez, 2005; Gonzalez, Vanyukov, & Martin, 2005). Given the frequent and forcible changes in the cyber environment, an analyst has to make real-time decisions depending on past experiences and current knowledge.

Following Chi's (2006) view on the characteristics of expertise and the *relative* view of expertise (Chase & Simon, 1973), a cyber security analyst may be regarded as an expert with high levels of proficiency in information and network security when compared to a novice who is less knowledgeable. The term novice is used here in a generic manner, referring to a wide spectrum of individuals with relatively no knowledge of cyber security. The term "novices" also suggests that with proper training and with enough experience, individuals can become experts. More specifically, the relative view of expertise postulates that an expert is not expert due to some innate talent or cognitive ability that the novice cannot possess. Rather, a novice can become an expert with proper training. However, it is possible that some aspects of expertise depend on the ability to tune general cognitive skills, like sustained attention and information synthesis, to a specific context, providing contextualized ways to access and deploy domain specific knowledge (Perkins & Salomon, 1989).

Asgharpour, Liu, and Camp (2007) showed how individuals with various levels of knowledge in information security and years of experience, may have different mental models of cyber security. Higher proficiency in information security also suggests better performance in cyber detection than lower levels of knowledge. Experienced individuals are expected to make better decisions than inexperienced ones. An expert is expected to detect features and meaningful patterns that a novice cannot (Shanteau, 1987). Knowledge and previous experience should make an expert more sensitive to cues that are overlooked by a novice. Careful attention to these cues can foster the identification of patterns that construct a problem and should promote the choice of the appropriate courses of action. Such expertise appears to be domain specific, and it is built up through experience and intensive practice (Randel, Pugh, & Reed, 1996). However, expertise may be domain limited and context dependent. Expertise can also make individuals more rigid and result in problematic adaptation in more dynamic environments (Chi, 2006). Furthermore, depending only on domain knowledge and neglecting general cognitive skills and heuristics can harm the ability of experts to mitigate atypical problems.

Goodall et al. (2009) studied cyber security analysts and the practical aspects of intrusion detection. Their work particularly highlights the expertise required to successfully accomplish the intrusion detection task. It comprises of domain knowledge in information and network security, and also local knowledge grounded in the analyst's unique environment. In general, domain knowledge is the fundamental knowledge obtained through long and deliberate learning (Ericsson & Lehmann, 1996). It includes theoretical knowledge that the expert acquires through formal education, training, or certification (Chi, 2006). Domain knowledge also includes practical knowledge learned through hands-on practice and experience with tools, methods of operation, and workflows. Domain knowledge acquired through formal learning processes lays the essential foundation of requisite knowledge for the work of the cyber security analyst. However, domain knowledge may not be enough to detect cyber-attacks in operational environments. In addition to domain knowledge, the analyst may need situated knowledge (Goodall et al., 2004, 2009). Situated knowledge is implicit, hard to articulate, and organization-dependent (Schmidt & Hunter, 1993). This type knowledge tends to be dynamic and the expert acquires it through continued interactions with a specific operating environment. In the context of information and network security, effectively learning the nuances of a particular network is often achieved by tuning and adjusting the IDS so it will detect threats and meet the organization's security needs without standing in the way of legitimate network users. Thus, for effective threat detection in a network, the analyst should know how to operate an IDS in general and have experience in using the IDS in that specific network. Given that cyber-attacks are represented in abnormal network activity, an analyst should be able to define normal and abnormal network activity and utilize these definitions to detect attacks. As what can be considered normal network activity in one environment may be indicative of malicious activity in another, intrusion detection depends on the ability to integrate domain and situated knowledge in a dynamic environment (Yurcik et al., 2003).

To extend the qualitative ethnographic research methods that were used to understand the mental model and general workflows of cyber security analysts (e.g., D'Amico & Whitley, 2008; Paul & Whitley, 2013), there is a need for quantitative tools and measures of performance (Chi, 2006) that can be used to evaluate and analyze the intrusion detection process. We designed an intrusion detection task which resembles the task confronted by many cyber security analysts. In this task, participants use a simplified version of an IDS to detect cyber-attacks in a relatively small network. This setting allowed us to evaluate the general human performance of experienced analysts outside their regular environment of operation and of novice participants. Both experts and novices can perform the detection task in this study by applying general reasoning process. However, experienced participants are expected synthesize between their large knowledge in cyber security and general cognitive skills, to benefit from their extensive experience. As such, we expect experts to perform better when compared to novices, who mainly depend on their cognitive skills. An intrusion detection task involves detecting malicious network events in a sequence and then deciding whether the whole sequence represents a cyber-attack. We introduced various sequences of network events, representing different types of cyber-attacks and examined the interplay between the detection of malicious network events, the type of attack, and expertise. We hypnotize that experience will allow experts to make better judgments regarding the entire sequence of network events. Their past experience with different cyber-attacks will support the integration of the observed network events and evidence regarding malicious network activity when deciding whether there is a cyber-attack or not. It is also predict that experts will be more confident in their decisions, as they base their decisions on a large knowledge base of past experiences. However, it is possible that based on their past experiences, experts will also come up with several competing explanations to observed network behaviors. This might impair the ability to judge network events and decrease the confidence. Similarly, we hypnotize that experience will increase the accuracy of detection of malicious network events, expressed as a higher hit rate and a lower false alert rate when detecting malicious events, regardless of the specific network scenario. When it comes to classification of network events, our hypothesis is that experts will be able to identify and correctly interpret the relevant attributes of network events. Thus, we hypnotize that the experts will be consistently perform better than novices, regardless of the type of attributes that construct a network event.

3. The simplified intrusion detection task

In this study, participants served as security analysts of a fictitious online retail company. In this role, their duty was to protect the company's computer network from a malicious attacker located outside the company. Based on the network described by Lye and Wing (2005) and as illustrated in Fig. 1, we used a simple stereotypical computer network. This kind of network topology is common for local corporate networks that are connected to the Internet. Such corporate networks typically consist of a web server, a file server, and a cluster of workstations. This network setting is commonly used in cyber security research and training, as well as in the operative networks of real-world mid-size corporations (Dutt, Ahn, Ben-Asher, & Gonzalez, 2012; Lye & Wing, 2005; Xie, Li, Ou, Liu, & Levy, 2010). Without comprehensive training, it is unlikely that novices could interact with commercially available IDS like Snort (http://www.snort.org/) or even with a high fidelity IDS mockup. As such, the experimental system required some compromises in the representation of the cyber environment.

Detailed instructions stated that the local corporate network is connected to the Internet through a router that routes Internet traffic to and from the local network. The network has two zones or sub-networks: one containing a public web server, and the other containing a private file server (with payroll, accounting, sales, marketing data, etc.) and a private cluster of workstation computers that company employees use for their daily work. The public web server runs two services (httpd and ftpd) and enables shoppers on the Internet to buy products using the company's website. The fileserver stores the company's data and runs two services (ftpd and nfsd) that allow access to the data over the network. The employees of the company use their workstations to access the Internet, as well as the data stored on the fileserver. The

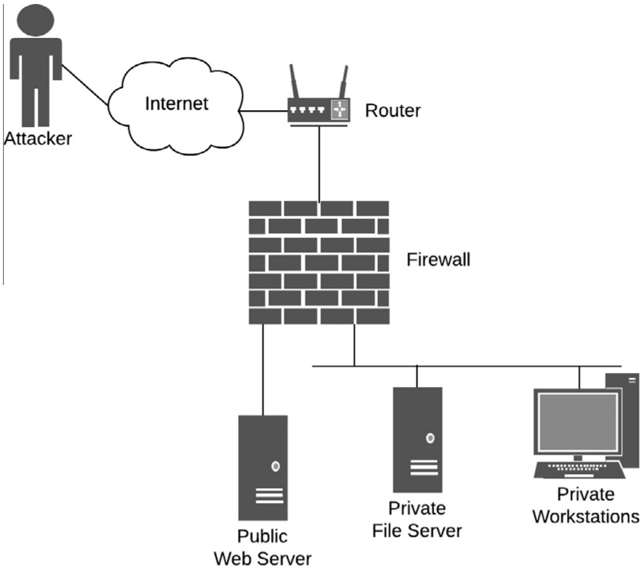

Fig. 1. The local computer network at an online retail company.

firewall prevents unwanted Internet connections from entering the local network, and it also checks the traffic between the different components of the local network.

A security analyst (i.e., the participant) monitors the network by observing a sequence of network events, as seen in Fig. 2. Each sequence is independent and includes network events that represent a single network scenario. The order of the events within a scenario corresponds to the order in which they occurred over the network, and the security analyst's goal is to decide whether or not a network event represents an attack. At the end of each network scenario, an analyst had to determine if the entire network scenario represented a cyber-attack on the company's network or not.

The simplified monitoring tool, seen in Fig. 2, presents network events one by one. Each event has an identification number which corresponds to its presentation order in the scenario. The event can also include an Alert, which is generated by the IDS system. This alert provides information regarding a suspicious network activity; for example, indicating that a service has stopped or started. An analyst is informed that the IDS can generate false alerts and that the IDS can also miss malicious events. The last part of the display presents the event's description. It provides an analyst with important information regarding the specific component of the network where the event occurred. The structure of the description is consistent, and the description includes the network component name (e.g., web server, file server, etc.) and the processes which are currently running on that component (e.g., httpd, ftpd, etc.). The description informs an analyst the state of the network traffic between a pair of network components (e.g., between the web server and the fileserver). Following Lye and Wing's (2005) convention, the traffic information indicates the load between two network components and has one of the following values: 0 Mbps, 3.3 Mbps, 6.7 Mbps, and 10 Mbps. Where 3.3 Mbps is the normal traffic condition, 0 Mbps indicates that there is no traffic between two components, and 10 Mbps indicates maximal capacity. Additional information included in the description indicates whether or not an operation was executed on this network component. The simplified information for each network event resembled the presentation of events' signature in IDS, with simplifications allowing participants without experience with IDS to comprehend it.

An analyst can classify each event as malicious or not by checking or un-checking the corresponding "Is threat" box for each event, and it was possible to go back and check/uncheck any previously presented event before the network scenario had ended. At the end of each network scenario, an analyst had to decide whether the current scenario represents a cyber-attack or not and report their level of confidence about this decision. Upon completing a network scenario, a message informing analysts that a new network scenario is about to start appears.

3.1. Network scenarios

An attacker outside the corporation may try to gain access to the corporate network in order to obtain confidential information or to compromise an essential service. For this, the attacker can follow an attack called "island-hopping" attack (Jajodia et al., 2010), where the web server is compromised first, and then it is used to originate attacks on the file server or on the company workstations.

Based on the network structure illustrated in Fig. 1, we defined five network scenarios; each represents a different network behavior. Four of the networks scenarios represent the progress and escalation of a cyber-attack in the corporate network, and one network scenario represents normal network operation (Lye & Wing, 2005). The ability to analyze different types of cyber-attacks is important. It increases the ecological validity of the experimental systems, so that it resembles the diversity of real-world cyber-attacks. Furthermore, cyber-attacks differed in the severity of their outcomes: we introduced disastrous cyber-attacks where the network stops functioning by the end of the scenario and it is clear that this is not a normal behavior; we also introduced network scenarios that represent severe cyber-attack but without such apparent outcomes. For the latter type of cyber-attack, the decision of whether or not a network scenario represents a cyber-attack might prove more challenging to the decision maker and more crucial to the organization.

To maintain fine-grain control over the network events and attacks introduced to the participants, we used the same ground truth rule when constructing the network events for all the network scenarios in this study. An event is malicious if there is an alert and if the description of the event indicates irregular network traffic or an operation was executed on one of the network components (or both). Thus, a network event is malicious if the information regarding the event follows this rule: Alert $\cap$ (Operation $\cup$ Network Load). This indicates that the IDS is relatively reliable and does not miss malicious network events. However, the IDS could generate false alerts. Meaning that the event included an alert, but its description did not include any suspicious network activity (i.e., not operation nor network load). The IDS system in this study generated three false alerts in each of the following network scenarios. This rule is consistent with the idea that hardening the sensitivity of an IDS system allows it to correctly generate alerts for all malicious events (i.e., no missed events), while the strict set of rules generates some false alerts at the same time (D'Amico & Whitley, 2008).

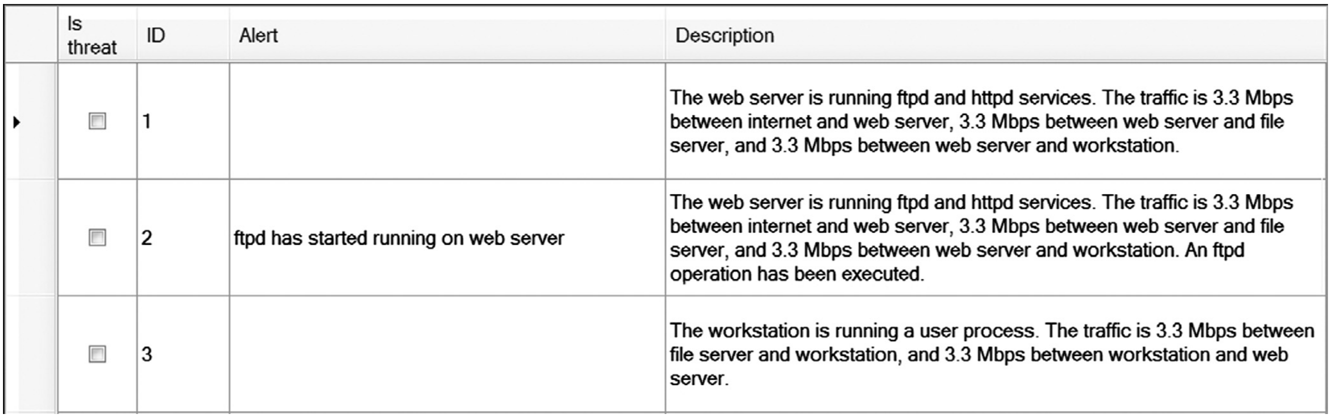

	Is threat	ID	Alert	Description
▸	□	1		The web server is running ftpd and httpd services. The traffic is 3.3 Mbps between internet and web server, 3.3 Mbps between web server and file server, and 3.3 Mbps between web server and workstation.
	□	2	ftpd has started running on web server	The web server is running ftpd and httpd services. The traffic is 3.3 Mbps between internet and web server, 3.3 Mbps between web server and file server, and 3.3 Mbps between web server and workstation. An ftpd operation has been executed.
	□	3		The workstation is running a user process. The traffic is 3.3 Mbps between file server and workstation, and 3.3 Mbps between workstation and web server.

Fig. 2. Network events with description and alerts from the IDS.

3.1.1. Deface website

A common target for cyber-attacks is the public web server that connects the corporate network with the Internet. The web server typically runs httpd and ftpd services, and the attacker gains control over the server by exploiting vulnerabilities in these services. Gaining control over the web server (i.e., root shell) allows the attacker to deface the corporate web site, which is changing the web site's content and then leaving. Usually, defacing a website has temporal and limited effect on its content. However, this kind of attack is challenging, as an analyst might mistakenly assume that there is a legitimate maintenance activity on the web server rather than a cyber-attack. Also, gaining control over the web server is a preliminary stage in more sophisticated and venturesome attacks. Thus, some of the malicious network events that appear in this scenario appear in other scenarios as well.

3.1.2. Sniffer detected

Once the attacker gains control over the web server in this scenario, he installs a sniffer and a backdoor program. The sniffer will sniff out passwords from the users as they access the file server or web server. Later, the attacker comes back through the backdoor program and collects the password list from the sniffer. If the attacker's action go unnoticed, the consequences of the attack can be disastrous. The attacker may have access to valuable information and can also sabotage the network services. In this scenario, the sniffer is detected before the attacker manages to sabotage the network services. The detection of the sniffer indicates that the network is under attack, as some of the network services were compromised and the attacker might have gained access to user accounts.

3.1.3. Denial-of-service (DoS)

Another course of action that an attacker can take after gaining control of the web server is to launch a DoS attack from inside the network. In this scenario, the attacker installs a virus that gradually increases the traffic load within the network. Once the traffic load is beyond the network's full capacity, the network grinds to a halt and nothing productive takes place over it.

3.1.4. Stealing confidential data

This scenario represents the progression of a cyber-attack, starting from normal network operation and ending after the attacker attained confidential information and obstructed the normal operation of the network. Here, the attacker gains access to the web server by compromising one of the services it runs. Once the attacker has hacked into the web server, he installs a sniffer and a backdoor program. The sniffer collects passwords from the network and the attacker uses the passwords to steal confidential information stored on the file server or the workstations. To cause more damage to the network, the attacker shuts the network down using the obtained privileges at the end of this scenario.

3.1.5. No attack scenario

Unlike the above network scenario, this scenario included events that are part of a normal network operation.

4. Methods

4.1. Participants

We recruited 55 participants from the university participant pool and invited them to a computer laboratory at Carnegie Mellon University. As these participants were recruited from the general student population, they are considered novices in cyber security as none of them were part of the cyber security workforce. Participants were compensated with $10 as base payment and could earn additional monetary incentive based on task performance. Participants earned 1 cent for classifying each network event as attack and no-attack correctly, and lost 1 cent for each incorrect classification.

In parallel, we recruited 20 security professionals from technical communities, such as the Computer Emergency Response Team (CERT), professionally oriented social network (e.g., LinkedIn), and mailing lists. Expert participants performed the task online, and earned 1 point for each attack and no-attack correctly classified and lost 1 point for each attack and no-attack incorrectly classified. Each point earned was worth one ticket for a raffle with a prize of one $50 Amazon.com gift card that occurred after completing the data collection.

4.2. Expertise questionnaire

Although numerous studies examined the work of the cyber security analyst with a specific interest in intrusion detection and the use of IDS, there are no tools to evaluate expertise in this field. To measure or determine expertise, previous studies mainly used interviews (Botta et al., 2007; D'Amico et al., 2005; Goodall et al., 2004, 2009). Other studies combined interviews with participatory observation (Werlinger, Hawkey, Muldner, Jaferian, & Beznosov, 2008) or a card sorting task (Paul & Whitley, 2013). These studies required meeting an analyst in person and relied on job titles as a verification of expertise level.

In this study, we wanted to distinguish between experts with profound understanding in cyber security and novices who currently have little to no expertise in cyber security. For that, we developed a short questionnaire that aims at differentiating between these two populations. Goodall et al. (2009) identified several types of expertise that are required by a security analyst to detect threats, and the ways in which expertise may be developed. The first is domain knowledge and the other is situated knowledge. Given that our goal was a comparison of detection performance according to their type of knowledge in cyber security, our questionnaire assesses only domain knowledge through questions on network and information security. It evaluates the participants' domain knowledge in two independent dimensions. The first dimension is theoretical knowledge and the second is practical knowledge. Theoretical knowledge refers to participants' familiarity with the correct meaning of different technical terms, types of cyber-attacks, and the use of network security tools. For example, participants had to state what are the differences between passive and reactive IDS. In another question participants had to explain what is the consequence of blocking port 80 on a web server, without any other changes in its default configuration. The second dimension is practical knowledge, and it evaluates working experience in information and network security, specific training, secondary education, certifications, time spent on resolving information and network security, and daily usage of security tools. Here participants had to state how many years of working experience they have in network operation and security area, with the possible answers of: None (1), A few months (less than a year)(2), 1–5 years (3), 5–10 years (4) and 10+ years (5). Another question evaluated more specifically the interaction with an IDS with the possible responses: None (1), About once every month (2), About once every week (3), Once every day (4), Many times every day (5). As for the network scenarios, we designed and validated the questionnaire based on discussions with cyber security professionals, including a practitioner in the office of information and network security at the university and a faculty member in the department of Computer Science who studies intrusion detection.

4.3. Design and procedures

Based on the network structure illustrated in Fig. 1 and the attack scenarios described above, we defined specific sequences of network events that represent the progress and escalation of cyber-attacks. Each scenario was composed of 20 network events. The base-rate of the malicious events in all these scenarios was the same with the value of .35. Meaning that out of 20 network events, 7 events were malicious. We followed Lye and Wing's (2005) definitions for the potential state space of the network and the possible transitions in the network state and accordingly generated 10 network scenarios. The network scenario differed in the type of attack, how the type of service that the attacker compromised first to gain control over the web server (i.e., ftpd or httpd), and where the sensitive data was stored (i.e., file server or on one of the workstations). In addition to the attack scenarios, we also constructed one scenario that represented normal network operation. In general, these network scenarios are relatively short compared to monitoring ongoing traffic in a network. In addition, the proportion of false alerts in the network scenarios is relatively low, compared to the high proportion of false alert IDS usually generate. Due to these simplifications, the experimental environment carries some differences from the actual task performed by security professionals. However, these simplifications were carefully considered to accommodate novices in this task, and we expected that experts would benefit from the simplified environment as well.

The scenarios were presented using the IDS-tool (see in Fig. 2), and a new event appeared on the screen every 10 s. As network events appeared, participants classified each as an attack or no-attack. At the end of the scenario, participants had to determine if the entire network scenario represented a cyber-attack or not, and then state their confidence in this decision. Participants in the novices group saw the 10 scenarios in a random order, and then completed the expertise questionnaire. Participants in the experts group performed the experiment using an online version of the IDS-Tool and saw only 3 randomly selected scenarios of the original 10 and then completed the expertise questionnaire. The online version of the IDS-Tool allowed analysts to complete the task remotely and increased our participants pool. Also, keeping the task relatively short increased the experts' willingness to participate in the experiment and the task completion rate. Altogether, novices completed the experiment in about 60 min, whereas experts completed the experiment in about 25 min.

5. Results

Analyses of the expertise questionnaire indicate a clear distinction between the two groups of participants. Most of the participants in the experts group (80%) had 1 or more years of experience in network operation and information security, and 35% of them had more than 10 years of practical experience. In contrast, 93% of the participants in the novices group stated that they have no experience in network operation and information security. Similarly, 80% of the experts reported dealing with at least one cyber security incident each day and 60% of them spent one or more hours a day handling issues related to network operation and security, whereas 60% of the novices reported dealing with cyber security once a year or less and 96% do not spend any time handling issues related to network operation and security on a day-to-day basis. Regarding IDS usage, 60% of the experts used an IDS at least once a month, whereas 93% of the novices never used an IDS.

When evaluating theoretical knowledge in information security, 100% of the experts knew the right definition for a DoS attack, compared to only 36% of the novices. All of the experts knew the right definition for phishing attacks, and a relatively high proportion of novices (85%) also knew the right definition for phishing attacks. Phishing attacks can target any end-users, whereas DoS attacks target a network: this can explain the smaller difference found between experts and novices when comparing the responses to the phishing and DoS attack questions. Most of the experts (90%) knew what differentiates between a reactive IDS to a passive IDS compared to only 27% of the novices.

We integrated the responses regarding theoretical and practical knowledge separately, by calculating the individuals' theoretical and practical knowledge scores on a scale between 0 and 1. Results indicated that theoretical and practical knowledge are two independent dimensions of expertise for experts (Cronbach's $\alpha < .01$) and novices (Cronbach's $\alpha = .463$). Furthermore, it seems that the disassociation between theoretical and practical knowledge was higher for experts compared to novices. As illustrated in Fig. 3, the two groups of participants have distinct characteristics. Both the theoretical and practical knowledge of experts were significantly higher than the theoretical and practical knowledge of novices, $t(73) = 13.206$, $p < .001$; $t(73) = 14.179$, $p < .001$, respectively.

These findings indicate that even when members of the experts group are not constantly engaged in IDS monitoring task, they had advanced knowledge and experience in network operation and cyber security, especially when compared to the novice group. Furthermore, most of the experts achieved optimal or near optimal score in the theoretical knowledge questions and the variability between experts was mainly found in the practice dimension. Novices had limited practical knowledge with limited variability and higher variability in their theoretical knowledge. These findings suggest that the questionnaire can be used to identify groups with different knowledge levels. For experts, however, it does not provide a fine-grain classification of different levels of expertise like journeyman, expert, and master (Chi, 2006).

5.1. Detection of attack scenarios

The overarching goal of the security expert is to protect the network from attacks. However, we find no difference in performance between the expert and novice groups. Participants in the experts group correctly detected 67% of the attack scenarios with only 12% of false detections ($d' = 1.60$); while participants in the novices group detected 68% of the attack scenarios with 14% of false detections ($d' = 1.52$).

After classifying a network scenario as an attack scenario or not, participants provided a short explanation of their decision and stated their confidence in the decision on a Likert scale ranging from 'Highly Unconfident' (1) to 'Highly Confident' (5). On average,

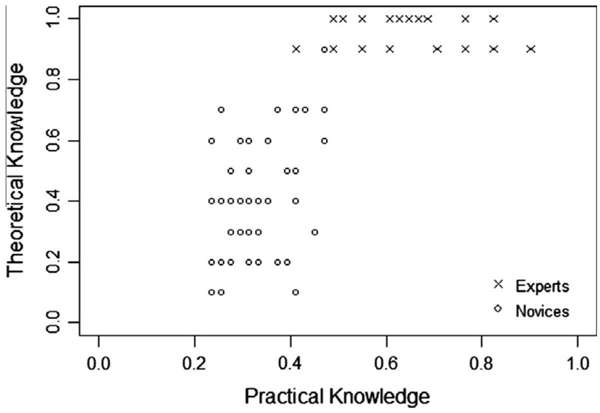

Fig. 3. Practical and theoretical knowledge in network security of experts and novices.

experts (3.35) and novices (3.12) had a similar confidence in their responses, $t(608) = 1.647$, p = ns. However, experts were more confident in their decision that a sequence of network events represented an attack (mean = 3.58, SD = .91) than they were in their decision that it does not represent an attack (mean = 3, SD = .78), $t(58) = 2.578$, $p = .012$. While participants in the novices group had similar levels of confidence for decisions in the attack scenarios (mean = 3.15, SD = 1.02) and in the no attack scenario (mean = 3.05, SD = 1.12), $t(548) = 1.080$, p = ns. The difference between experts and novices suggests that experts' confidence was sensitive to the decision (i.e., attack or no attack scenario), and they were more careful and less confident when deciding that a network scenario does not represent an attack.

To better understand how experts and novices performed in the task and to evaluate whether the type of cyber-attack influenced each group differently, we examined detection in each of the network scenarios separately (see Fig. 4). Analysis indicated that for the Deface Website attack, expert performance was about chance level (42%), and novices performed significantly below chance level (33%), $\chi^2(1, N = 12) = .333$, p = ns. and $\chi(1, N = 110) = 6.564$, $p = .005$, respectively. Participants in both groups were better at detecting DoS attacks (Experts = 70%, Novices = 76%) and performed significantly above chance level, $\chi(1, N = 10) = 4.8$, $p = .028$ and $\chi(1, N = 110) = 30.582$, $p < .001$, respectively. For experts (59%) and novices (62%) alike, classifying the Sniffer Detected network scenario correctly was not trivial and detection rates for it did not differ significantly from chance level, $\chi(1, N = 8) = 1.5$, p = ns. and $\chi(1, N = 110) = 3.636$, p = ns., respectively. Finally, experts (82%) and novices (77%) were more successful in classifying Stealing Confidential Data network scenarios correctly, and both groups performed significantly above chance level, $\chi(1, N = 22) = 8.909$, $p = .003$ and $\chi(1, N = 220) = 63.291$, $p < .001$, respectively.

We used a logistic regression model to examine the relationship between the decision whether a sequence of network events represented a cyber-attack or not, and the classification of network event as malicious. The model included the binary decision, Attack or No Attack as a dependent variable with the participant's group and the number of network events that the participant classified as malicious as the independent variables. The analysis indicated a significant main effect of the number of detected events, ($z = 9.152$, $p < .001$). Each event that the participant considered as a threat increased their likelihood of declaring a cyber-attack. The model also yielded a significant interaction between the participant's group and the number of events considered as malicious, as seen in Fig. 5. This finding indicates that after classifying a relatively small number of events as malicious, novice participants were more likely than experts to decide that there was a cyber-attack ($z = -3.816$, $p < .001$). However, once a relatively larger number of events was classified as malicious, experts were more likely to decide that there was a cyber-attack. Also, classifying an additional event as a threat had more influence on the experts' decisions than on the novices' decisions.

Examining the connection between the number of events correctly classified as malicious and the decision that a network scenario represented a cyber-attack revealed a similar gap in the performance of experts and novice. A logistic regression model with the number of correctly detected threats and the participant's group as independent variable and the decision regarding the whole scenario as the dependent variable indicated that as the number of detected malicious events increase, the likelihood of both experts and novices deciding that there was a cyber-attack increased significantly ($z = 8.194$, $p < .001$). However, as seen in Fig. 6 and as indicated by the significant interaction between the number of detected threats and the participants' group, novices were more likely to decide that a sequence of network events represented a cyber-attack after detecting a small number of actual threats. In contrast, experts were less likely than novices to decide that there is a cyber-attack when detecting a small number of threats ($z = -2.070$, $p = .038$). When most of the threats were detected, however, both experts and novices were likely to decide that there was a cyber-attack.

5.2. Detection of malicious network events

Participants classified each of the 20 network events in each scenario as malicious or not. In general, there were 1200 network events for experts and 11,000 network events for novices (participants × number of Scenarios × number of events in a scenario). Experts classified a significantly larger proportion of the malicious network events as threats and a significantly lower proportion of benign events as threats compared to novices, $\chi(1, N = 3829) = 15.651$, $p < .001$ and $\chi(1, N = 8371) = 15.068$, $p = .024$, respectively. As seen in Fig. 7, cyber security professionals correctly detected 55% of the malicious events with 15% of false detection ($d' = 1.18$), and novices detected 44% of the malicious network events with 18% of false detection ($d' = .78$).

Next, we calculated the proportion of correct (i.e., hits) and incorrect (i.e., false alerts) threat classification for each combination of participant and network scenario (see Fig. 8). Using a linear regression model, we analyzed the detection rates of malicious events with scenario and participants' group as independent variable. The results indicated that detection rate varied significantly across the different network scenarios, $F(3, 539) = 6.767$, $p < .001$. Detection rates of malicious network events in the Sniffer Detected

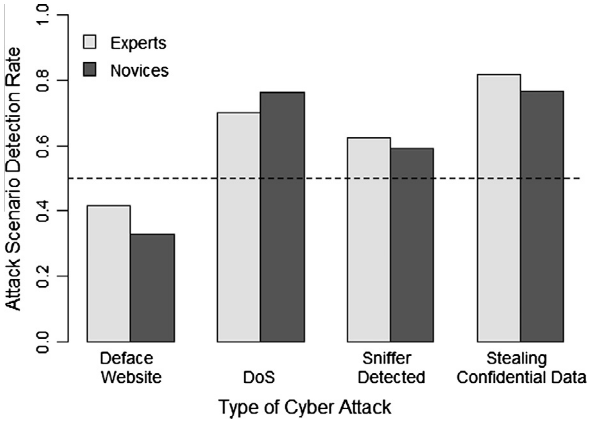

Fig. 4. Experts and novice attack detection rates for the four cyber-attack network scenarios.

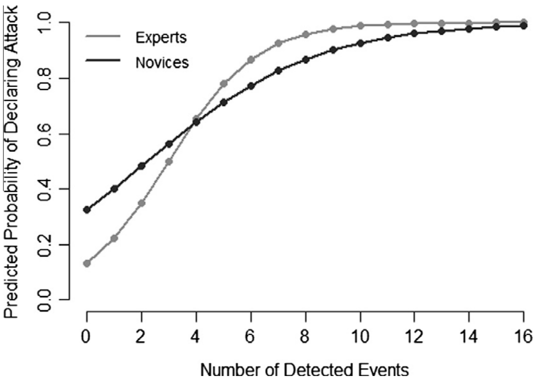

Fig. 5. The probability of experts and novices to declare a cyber-attack depending on the number of events classified as malicious.

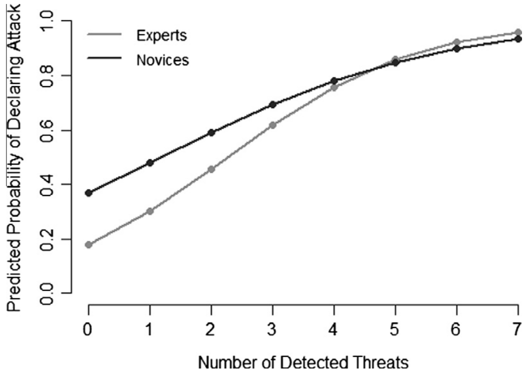

Fig. 6. The probability of experts and novices to declare a cyber-attack depending on the number of correctly detected threats.

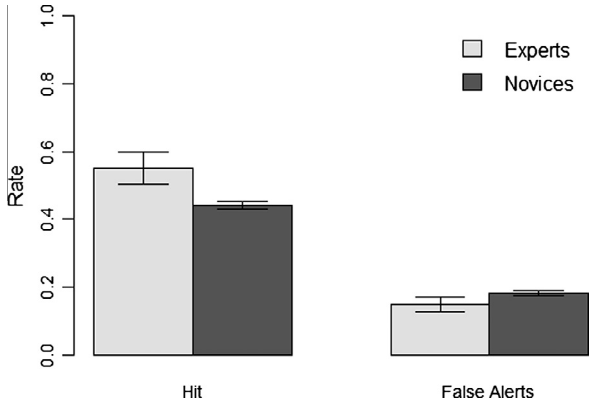

Fig. 7. Overall hit and false alerts rates for experts and novices when detecting malicious network events.

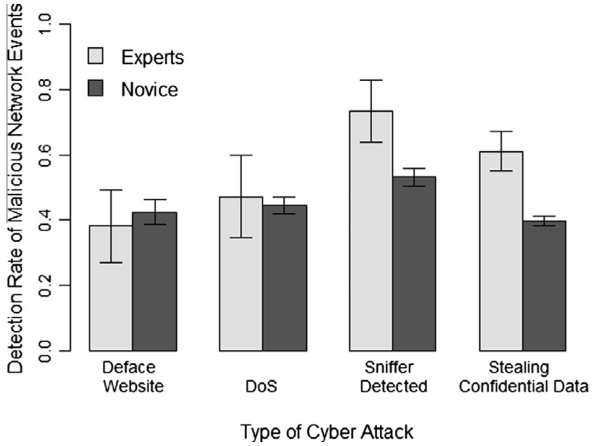

Fig. 8. Expert and novice detection rates of malicious network events for different types of cyber-attacks.

scenario (mean = 54%, SD = 28%) were significantly higher than the detection rates in the DoS (mean = 45%, SD = 28%) and the Deface Website (mean = 41%, SD = 30%) scenarios, $t(236) = 2.666$, $p = .008$ and $t(183) = 2.906$, $p = .004$. A significant two-way interaction was found between the participants' group and the scenario type, $F(3,539) = 2.715$, $p = .044$. This finding suggests that the cyber security experience influenced detection rates in the Stealing Confidential Data and Sniffer Detected scenarios, $t(240) = 4.105$, $p < .001$ and $t(116) = 1.948$, $p = .054$, respectively. Thus, experts performed better than novices in these scenarios, while there were no differences between the two groups in the other scenarios.

Analyses of the proportion of false alerts for each combination of participant and network scenario revealed that the type of network scenario also had a significant effect on the false classification of benign network events as malicious, $F(4,600) = 5.429$, $p < .001$. The false alert rate in the DoS scenario (mean = 23%, SD = 19%) was significantly higher than in the Stealing Confidential Data (mean = 18%, SD = 16%), the Deface Website (mean = 14%, SD = 15%), and the No-Attack (mean = 12%, SD = 14%) scenarios, $t(360) = 2.410$, $p = .016$, $t(185) = 3.109$, $p = .002$ and $t(181) = 4.018$, $p < .001$, respectively.

5.3. Attributes to network events

When constructing the networks scenarios and the events that they are composed of, we used a relatively simple ground truth rule to generate malicious events. A network event was malicious if the description of the event followed this rule:

$$\text{Alert} \cap (\text{Operation} \cup \text{Network Load})$$

Thus, the description of malicious events always had an alert on it, and it could also indicate irregular network traffic or an operation that was executed on one of the network components (or both). Events of the form *Alert, Network Load, and Operation* included all the indications that they are malicious and were similarly detected by participants in the expert (74%) and novice (64%) groups, $\chi(1,N = 545) = 1.904$, p = ns. Also, participants in the expert (57%) and novice (49%) groups correctly detected a relatively similar proportion of malicious events in the form of *Alert and Network Load*, $\chi(1,N = 1578) = 3.670$, $p = .055$. However, malicious events in the form of *Alert and Operation* were more likely to be detected by experts (47%) than by novices (33%), $\chi(1,N = 1706) = 13.158$, $p < .001$. This suggests that both experts and novices were sensitive to the network load. However, experienced cyber professionals were more aware of the relevancy of operations combined with an alert than were novices.

Analyses of the network events that participants falsely classified as malicious (i.e., false alerts) support the claim that cyber security experience influences how participants make connections between the attributes that constructed a network event. When the description of the event included only one possible indicator for a malicious event (i.e., *Alert* or *Network Load* or *Operation*), we find no differences between experts and novices participants. Events in the form of *Alert* were equally misclassified as malicious by experts (29%) and novices (26%), $\chi(1,N = 1830) = .761$, $p <$ ns. Events in the form of *Load* were equally misclassified as malicious by experts (19%) and novices (21%), $\chi(1,N = 2078) = .341$, $p <$ ns. Likewise, events in the form of *Operation* were equally misclassified as malicious by experts (7%) and novices (10%) alike, $\chi(1,N = 1218) = 1.519$, $p <$ ns. However, novice (30%) participants were more likely to generate false alerts by classifying benign *Network Load and Operation* events as malicious, compared to experts (15%), $\chi(1,N = 1299) = 14.500$, $p < .001$. Thus, experts were aware that the operation can be the cause of the load and that load combined with an operation can be a benign network activity that does not represent an attack. This suggests that experts evaluated the network load in the context of the network status and integrated several attributes in the event's description together.

6. Discussion

There is increasing awareness of the importance of human decision making for the safety and security of IT systems (Ben-Asher, Meyer, Möller, & Englert, 2009; Möller, Ben-Asher, Engelbrecht, Englert, & Meyer, 2011). Effective and efficient design of information security systems that support the work of the security analyst, like an IDS, may benefit from a better understanding of the capabilities and limitations of the human decision maker who uses them (Kaufman, Perlman, & Speciner, 2002). In particular, it is important to understand how past experiences and knowledge influence decision making in a highly dynamic environment like the cyber world.

In this study, we contribute toward a better understanding of the human decision making process in the detection of cyber-attacks. We developed an expertise questionnaire that helps to discriminate between domain experts and novices. Participants that were cyber security professionals had extensive experience and significant knowledge in information and network security. Even though not all of them were qualified security experts that operated an IDS on a daily basis, comparing them to participants in the novices group yielded distinctions in their approach to network security tasks. The most substantial difference in performance between expert and novices was found in the detection of malicious events within a sequence of network events, while only minor differences were found with respect to the decision of whether or not the entire sequence represented a cyber-attack. Knowledge of information and network security contributed more to the detection of malicious events than to the detection of a cyber-attack.

An explanation for experts' superior detection of malicious events is that their knowledge allowed them to interpret the description of a network event, to understand the connections between the different attributes that compose an event, and to judge the event within the context of the specific network and network activity. Previous experience allowed experts to understand the connection between an IDS alert, the load in the network, and what operations are executed in the network. This resulted in an overall higher hit rate and lower false alarm rate compare to the novices. The ability of domain experts to detect features and meaningful patterns that a novice cannot is in-line with Shanteau (1987) and Randel et al. (1996).

However, the performance of experts and novices was relatively similar when detecting cyber-attacks. Specific types of cyber-attacks that were challenging for the novices were also challenging for the experts, and those that were relatively simple for the experts were also relatively simple for the novices. Shanteau (1992) suggest that experts will do well in tasks where the stimulus is static compared to dynamic stimulus; and better when decision relates to one judgment rather than a process or behavior. Although this may explain the relatively low performance of experts, it is also possible that some aspect of the environment hindered the experts and obstructed them from taking full advantage of their past experience and knowledge. Experts were taken away from their own familiar environment of operation and thus may have lacked situated knowledge. This finding quantitatively confirms the importance of situated knowledge as proposed by previous qualitative studies (Botta et al., 2007; Goodall et al., 2004, 2009). Losing a substantial amount of the situated knowledge and having to rely on domain knowledge impaired experts' ability to correctly judge a sequence of network events. Some verbal explanations may illustrate how a lack of situated knowledge influenced the intrusion detection processes. For example, expert (P02) did not classify the Deface web site scenario as a cyber-attack and explained:

> *"I think that the user is updating the website. The sequence of shutdown events and transactions, and the lack of transfer between the webserver and the file server seem to indicate that nothing too suspicious is going on. Unless the attacker is being slow about it."*

This expert considered two alternative explanations for the network behavior, one corresponding to an attack and the other to legitimate maintenance of the web server. Here, the expert could have benefited from having situated knowledge to disambiguate the status of the network. Some explanations provided by experts included explicit requests for additional situated information like IP logs that would help them make a more correct or confident decision.

It is also possible that when detecting cyber-attacks, novices benefited from highly indicative events that appeared toward the end of some network scenarios (e.g., DoS and Stealing Confidential Data scenarios). As seen in the results, novices were more likely than experts to judge a network scenario based on a limited amount of evidence (i.e., number of malicious events). Because the cyber-attack decision was made after observing all network events, it is possible that differences between experts and novices cannot be observed. However, we may expect that experts would detect the cyber-attacks earlier than novices and would be able to prevent damage by stopping the propagation of the attack through the network. This is illustrated by the explanations expert provided below. They seemed sensitive to the order in which the events appeared within a scenario. For example, an expert (P01) provided the following explanation for classifying a DoS network scenario as a cyber-attack:

> *"Based on the order of connections, it seems that an attacker compromised a user workstation through the webserver, and then used the workstation to access the fileserver and install some sort of backdoor on the webserver."*

Interestingly, this expert refers to a possible course of action that the attacker took and does not refer to specific attributes of network events like load. This is in contrast to the common explanations novices provided which were mostly concerned with the attributes of network events (e.g., network load and alerts). For example, a novice (P10) provided the following explanation to a cyber-attack decision:

> *"There were many alerts, traffic was high at a lot of points, and the traffic went down to 0 at the end."*

The similar confidence levels experts of novices seem to be in contrast to the idea that experts tend to be overly confident (Chi, 2006). There are two possible explanations for this. First, it is possible that experts' situated knowledge is a major contributor to their overconfidence. Alternatively, it is possible that like other domains experts (e.g., weather forecasting), cyber security experts have the tendency to be cautious and conservative (Chi, 2006). The cautious behavior of experts can be also associated with a more accurate understanding of the severe consequence of a wrong decision in cyber security. As demonstrated by an explanation expert (P03) gave for classifying Deface Web Site scenario as an attack scenario and assigning the decision with the lowest confidence score:

> *"It's safer than saying no."*

The higher levels of confidence that experts exhibited when deciding that there is an attack compared to their lower levels of confidence when concluding that there is no attack also support the notion that cyber security analysts adopt cautious behaviors when monitoring a network.

7. Conclusion

In summary, expertise and practical knowledge play an important role in triage analysis: the task of classifying a network event as a threat or not and the connections between these small decisions and the overall attack decisions based on a sequence of network events. It is likely that in cyber security, the accumulation of information regarding network events drives the decision process regarding the whole sequence of events (Dutt et al., 2012). In this study, we treat all network events in a similar manner, meaning we do not consider the possibility that some events provide more diagnostic information than others. It is likely that expertise support the identification of such critical events and attending to them can bias the decision maker. Furthermore, in the current experiment the network scenarios contained a relatively small number of events, compared to real-world network traffic. This can provide an advantage to the experts, especially if they detect correctly the most informative events. To address this limitation in our future work, we plan to focus on experts and to use more complex and long network scenarios. Also, further research is required to specifically evaluate the information accumulation process before making a conclusions, and how the decision regarding multiple events is synthesized. This line of research will also consider the possibility that the each network event has a different weights or contribution to the final judgment of a network scenario. In this context, evidence accumulation models proposed by Busemeyer and Townsend (1993) or Ratcliff and Smith (2004) should be considered. However, information integration models that use automatic-intuitive processes of memory and perception and dynamically weight the contribution of each event to the final decision might be more suitable for modeling experts' decisions (Raab & Johnson, 2007).

A security analyst needs situated and domain knowledge to benefit from all available data sources and visualizations. Furthermore, situated knowledge should be considered in an analysts' training process. In addition to theoretical knowledge and practical experience, analysts should also be trained to quickly learn and adapt to novel and dynamic environments. An analyst should constantly update and expand her situated knowledge regarding the operational environment. Such information regarding the importance and function of servers in the network is rarely systematically collected into a repository and even when collected, it is static and becomes outdated rather quickly as the network constantly changes with new equipment being added and the existing equipment being modified, upgraded, or retired. Such situated knowledge is a prerequisite for more comprehensive and mission-oriented situation awareness. Finally, considering the increasing number of personal networks that end-users deploy by themselves (e.g., home network), the growing number and variety of devices connected to these networks (e.g., computers, smartphones, tablets, media smart-TV, etc.) and their complexity, intrusion detection can become a concern of many end-users without extensive domain knowledge in information and network security.

Acknowledgments

This research was partly supported by a Multidisciplinary University Research Initiative Award (MURI; # W911NF-09-1-0525) from Army Research Office and by the Army Research Laboratory under Cooperative Agreement Number W911NF-13-2-0045 (ARL Cyber Security CRA). The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation here on. The authors would like to thank Hau-yu Wong, Dynamic Decision Making Laboratory, for help with editorial work in the paper. We also would like to thank Rob Floodeen for his help in gathering expert participants for this study.

References

Asgharpour, F., Liu, D., & Camp, L. J. (2007). Mental models of computer security risks. In *Proccedings of the 6th annual workshop on the economic of information security (WEIS 2007)*.

Ben-Asher, N., Meyer, J., Möller, S., & Englert, R. (2009). An experimental system for studying the tradeoff between usability and security. *Proceedings of the international conference on availability, reliability and security: ARES'09* (pp. 882–887). Los Alamitos, CA: IEEE. http://dx.doi.org/10.1145/1280680.1280693.

Botta, D., Werlinger, R., Gagné, A., Beznosov, K., Iverson, L., Fels, S., et al. (2007). Towards understanding IT security professionals and their tools. In *Proceedings of the third symposium on usable privacy and security* (pp. 100–111). New York, NY: ACM. http://dx.doi.org/10.1145/1280680.1280693.

Brehmer, B. (1992). Dynamic decision making: Human control of complex systems. *Acta Psychologica, 81*(3), 211–241. http://dx.doi.org/10.1016/0001-6918(92)90019-A.

Busemeyer, J. R., & Townsend, J. T. (1993). Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment. *Psychological Review, 100*(3), 432–459. http://dx.doi.org/10.1037/0033-295X.100.3.432.

Chase, W. G., & Simon, H. A. (1973). The mind's eye in chess. In W. G. Chase (Ed.), *Visual information processing* (pp. 215–281). New York, NY: Academic Press.

Chen, P. C., Liu, P., Yen, J., & Mullen, T. (2012). Experience-based cyber situation recognition using relaxable logic patterns. In *Proceedings of the 2012 IEEE international multi-disciplinary conference on cognitive methods in situation awareness and decision support (CogSIMA)* (pp. 243–250). IEEE. http://dx.doi.org/10.1109/CogSIMA.2012.6188392.

Chi, M. T. H. (2006). Two approached to the study of experts' characteristics. In K. A. Ericsson, N. Charness, P. J. Feltovitch, & R. Hoffman (Eds.), *Cambridge handbook of expertise and expert performance* (pp. 121–130). Cambridge, UK: Cambridge University Press.

Cranor, L. F. (2008). A framework for reasoning about the human in the loop. In *Proceedings of the 1st conference on usability, psychology, and security* (pp. 1–15). USENIX Association.

D'Amico, A., Whitley, K., Tesone, D., O'Brien, B., & Roth, E. (2005). Achieving cyber defense situation awareness: A cognitive task analysis of information assurance analysts. Proceedings of the human factors and ergonomics society annual meeting (Vol. 49, No. 3, pp. 229-233). doi: http://dx.doi.org/10.1177/154193120504900304.

D'Amico, A., & Whitley, K. (2008). The real work of computer network defense analysts. In G. Conti, J. R. Goodall, & K. L. Ma (Eds.), *Proceedings of the workshop on visualization for computer security* (pp. 19–37). Berlin, Germany: Springer-Verlag. http://dx.doi.org/10.1007/978-3-540-78243-8_2.

Dutt, V., Ahn, Y. S., Ben-Asher, N., & Gonzalez, C. (2012). Modeling the effects of base-rates on cyber threat detection performance. In N. Rußwinkel, U. Drewitz, & H. van Rijn (Eds.), *Proceedings of the 11th international conference on cognitive modeling (ICCM 2012)* (pp. 88–93). Berlin, Germany: Universitaetsverlag der TU Berlin.

Edwards, W. (1962). Dynamic decision theory and probabilistic information processing. *Human Factors, 4*(2), 59–73. http://dx.doi.org/10.1177/001872086200400201.

Ericsson, K. A., & Lehmann, A. C. (1996). Expert and exceptional performance: Evidence of maximal adaptation to task constraints. *Annual Review of Psychology, 47*(1), 273–305. http://dx.doi.org/10.1146/annurev.psych.47.1.273.

Gonzalez, C. (2005). Decision support for real-time, dynamic decision-making tasks. *Organizational Behavior and Human Decision Processes, 96*(2), 142–154. http://dx.doi.org/10.1016/j.obhdp.2004.11.002.

Gonzalez, C., Ben-Asher, N., Oltramari, A., & Lebiere, C. (2014). Cognition and technology. In A. Kott, C. Wang, & R. Erbacher (Eds.), *Cyber defense and situation awareness* (pp. 93–117).

Gonzalez, C., Vanyukov, P., & Martin, M. K. (2005). The use of microworlds to study dynamic decision making. *Computers in Human Behavior, 21*(2), 273–286. http://dx.doi.org/10.1016/j.chb.2004.02.014.

Goodall, J. R., Lutters, W. G., & Komlodi, A. (2009). Developing expertise for network intrusion detection. *Information Technology & People, 22*(2), 92–108. http://dx.doi.org/10.1108/09593840910962186.

Goodall, J. R, Lutters, W. G., & Komlodi, A. (2004). I know my network: Collaboration and expertise in intrusion detection. In J. Herbsleb & G. Olson (Eds.), *Proceedings of the 2004 ACM conference on computer supported cooperative work* (pp. 342–345). New York, NY: ACM. http://dx.doi.org/10.1145/1031607.1031663.

Green, D. M., & Swets, J. A. (1966). *Signal detection theory and psychophysics*. New York: Wiley.

Jajodia, S., Liu, P., Swarup, V., & Wang, C. (2010). *Cyber situational awareness: Issues and research*. New York, NY: Springer.

Kaufman, C., Perlman, R., & Speciner, M. (2002). *Network security: Private communication in a public world*. Upper Saddle River, NJ: Prentice Hall Press.

Lye, K., & Wing, J. M. (2005). Game strategies in network security. *International Journal of Information Security, 4*(1–2), 71–86. http://dx.doi.org/10.1007/s10207-004-0060-x.

McHugh, J. (2001). Intrusion and intrusion detection. *International Journal of Information Security, 1*(1), 14–35. http://dx.doi.org/10.1007/s102070100001.

Möller, S., Ben-Asher, N., Engelbrecht, K. P., Englert, R., & Meyer, J. (2011). Modeling the behavior of users who are confronted with security mechanisms. *Computers & Security, 30*(4), 242–256. http://dx.doi.org/10.1016/j.cose.2011.01.001.

Paul, C. L., & Whitley, K. (2013). A taxonomy of cyber awareness questions for the user-centered design of cyber situation awareness. In M. Louis & A. Ioannis (Eds.), *Human aspects of information security, privacy, and trust 2013* (pp. 145–154). Berlin, Germany: Springer-Verlag.

Perkins, D. N., & Salomon, G. (1989). Are cognitive skills context-bound? *Educational Researcher, 18*(1), 16–25. http://dx.doi.org/10.3102/0013189X018001016.

Raab, M., & Johnson, J. G. (2007). Expertise-based differences in search and option-generation strategies. *Journal of Experimental Psychology: Applied, 13*(3), 158–170. http://dx.doi.org/10.1037/1076-898X.13.3.158.

Randel, J. M., Pugh, H. L., & Reed, S. K. (1996). Differences in expert and novice situation awareness in naturalistic decision making. *International Journal of Human–Computer Studies, 45*(5), 579–597.

Ratcliff, R., & Smith, P. L. (2004). A comparison of sequential sampling models for two-choice reaction time. *Psychological Review, 111*(2), 333–367. http://dx.doi.org/10.1037/0033-295X.111.2.333.

Schmidt, F. L., & Hunter, J. E. (1993). Tacit knowledge, practical intelligence, general mental ability, and job knowledge. *Current Directions in Psychological Sciences, 2*(1), 8–9. http://dx.doi.org/10.1111/1467-8721.ep10770456.

Shanteau, J. (1992). Competence in experts: The role of task characteristics. *Organizational behavior and human decision processes, 53*(2), 252–266.

Shanteau, J. (1987). Psychological characteristics of expert decision makers. In J. L. Mumpower, O. Renn, L. D. Phillips, & V. R. R. Uppuluri (Eds.), *Expert judgment and expert systems* (pp. 289–304). Berlin, Germany: Springer-Verlag.

Thompson, R. S., Rantanen, E. M., & Yurcik, W. (2006). Network intrusion detection cognitive task analysis: Textual and visual tool usage and recommendations. In *Proceedings of the human factors and ergonomics society annual meeting* (Vol. 50, No. 5, pp. 669–673). doi: http://dx.doi.org/10.1177/154193120605000511.

Werlinger, R., Hawkey, K., Muldner, K., Jaferian, P., & Beznosov, K. (2008). The challenges of using an intrusion detection system: Is it worth the effort? In *Proceedings of the 4th symposium on usable privacy and security* (pp. 107–118). New York, NY: ACM.

Werlinger, R., Muldner, K., Hawkey, K., & Beznosov, K. (2010). Preparation, detection, and analysis: The diagnostic work of IT security incident response. *Information Management & Computer Security, 18*(1), 26–42. http://dx.doi.org/10.1108/09685221011035241.

Xie, P., Li, J. H., Ou, X., Liu, P., & Levy, R. (2010). Using Bayesian networks for cyber security analysis. In *Proceedings of the 2010 IEEE/IFIP international conference on dependable systems and networks (DSN)* (pp. 211–220). Los Alamitos, CA: IEEE.

Yurcik, W., Barlow, J., & Rosendale, J. (2003). Maintaining perspective on who is the enemy in the security systems administration of computer networks. In *ACM CHI workshop on system administrators are users*. ACM Press.