Exploring Cognitive Processes to Develop Cybersecurity Defender Proficiency

by

David Schuster
Department of Psychology
San José State University
1 Washington Square
San Jose, CA 95192-0120
david.schuster@sjsu.edu
+1 408 924 5659

Exploring Cognitive Processes to Improve Cybersecurity Defender Proficiency

Practitioners and researchers are the target audiences.

Problem statement. Despite the mission-critical role of people in protect and defend roles, relatively little is known about how cognition supports defender proficiency, a pressing problem given a workforce shortage and skills gap (Crumpler & Lewis, 2019). Understanding of the cognitive processes relevant to cybersecurity roles could support strategies to develop the skills of defenders and increase workforce participation. These processes, called macrocognition, are the result of cognitive resources at work in operational environments (Klein et al., 2003).

Research questions. The cognition of defenders is not well understood, and defender proficiency is only starting to be defined; what cognitive processes support proficiency in defender roles? To address the skills gap, a critical question is how to develop the proficiency of defenders efficiently; how can understanding of defender cognition be used to strengthen the cybersecurity workforce?

Contribution. A methodology using cognitive task analysis (CTA) is presented to describe the macrocognition of defenders. CTA is a collection of "tools and techniques for describing the knowledge and strategies required for task performance" (Schraagen et al., 2000, p. xiii). This work is complementary to prior defender CTAs in that CTA is used to describe the cognition of individuals with the aim of generalizing those processes across related work roles.

Rationale. This approach connects work roles based on related cognitive skills. Understanding of macrocognition could help defenders connect how they think with their work outcomes and may unlock novel, evidence-based strategies for workforce development, especially in training and recruitment.

Investigative approach. Macrocognition's role in describing cognition at a useful layer of abstraction and complement to the NICE Framework (NIST SP 800-181; Newhouse et al., 2017) is introduced. CTA is discussed as a method of understanding proficiency in support of workforce development, and CTAs relevant to this perspective are reviewed. A use case with two industry defenders is presented, and lessons learned are offered to accelerate replication.

Lessons learned. The use case shows how concept mapping can lead to macrocognitive themes. The themes suggest practice implications and new research questions. Successes and failures in the use case are discussed so that researchers can more efficiently link CTAs to defender macrocognition.

Implications for practice. The NICE Framework defines knowledge, skills, abilities, and tasks mapped to work roles with emerging discussion of qualifications; the methodology complements the NICE Framework by establishing the cognitive mechanisms that support individuals performing the skills and abilities. Macrocognition for cybersecurity may predict performance across similar roles even when policies, departments, organizations, sectors, and technologies change. There is potential value in diagnosing macrocognition to improve performance

outcomes. This research can result in a more prepared cyber workforce, trained and recruited on the basis of cognitive skills relevant to their role.

Implications for research. This work serves as a call and framework for additional CTA research to understand cognitive processes, and replication is necessary. Through the methodology, quantitative researchers can benefit from better understanding of relevant contextual factors, which can lead to more meaningful experimentation and establish reliable measurement.

References

Crumpler, W., & Lewis, J. A. (2019). *The Cybersecurity Workforce Gap*. Center for Strategic and International Studies. https://www.csis.org/analysis/cybersecurity-workforce-gap

Klein, G., Ross, K. G., Moon, B. M., Klein, D. E., Hoffman, R. R., & Hollnagel, E. (2003). Macrocognition. *IEEE Intelligent Systems*, *18*(3), 81–85. https://doi.org/10.1109/MIS.2003.1200735

Newhouse, W., Keith, S., Scribner, B., Witte, G. (2017). *National initiative for cybersecurity education (NICE) cybersecurity workforce framework* (Report No. SP 800-181). Gaithersburg, MD: National Institute of Standards and Technology. https://csrc.nist.gov/publications/detail/sp/800-181/final

Schraagen, J. M., Chipman, S. F., & Shalin, V. L. (2000). *Cognitive task analysis*. Psychology Press.

Exploring Cognitive Processes to Develop Cybersecurity Defender Proficiency

I. Introduction

Problem Statement

Defending against cyber threats is a large and growing problem for companies, with a 67% increase in breaches in the last five years (Bissell et al., 2019) and an expected annual cost of $6 trillion by 2021 (Morgan, 2019). It also provides challenging, mission-critical work for defenders, cybersecurity professionals who “identify, analyze, and mitigate threats to internal information technology systems and/or networks” (Newhouse et al., 2017, p. 11). Defenders can be defined as a category of cybersecurity professionals within the NICE Framework’s Protect and Defend category (Newhouse et al., 2017); they include the work roles of analyst, defense infrastructure support specialist, incident responder, and vulnerability assessment analyst.

This work is challenging because of the asymmetry between attackers and defenders (Yurcik et al., 2003). Successful defense requires defenders to succeed in every instance, and a single failure can result in consequences to the business. This asymmetry is compounded by fact that organizational infrastructure can be large in scale and continuously changing as it supports the activity of the organization, resulting in a massive amount of data. Organizations are vulnerable to external attacks and insider threats. On their own or as part of a team, defenders must ensure that alerts are analyzed and understood, understand the severity of threats, and coordinate appropriate responses (Shah et al., 2018). They apply a great deal of knowledge to time-sensitive, high-stakes situations while juggling stakeholder requirements, available resources, and uncertainty (Rooney & Foley, 2018).

A cybersecurity workforce shortage underscores the critical need for people in effective cyber defense. In the United States alone, a shortage of over half a million workers existed in

2020, evidenced by the number of openings (NIST, 2020). Over 100,000 of these openings were for Protect and Defend roles (NIST, 2020). A contributing factor to the workforce shortage is a lack of available workers possessing the necessary skills for the job (Crumpler & Lewis, 2019). Solutions are needed to close the skills gap, including by developing the skills of defenders and increasing participation in the workforce.

NIST SP 800-181, known as the NICE Framework, provides a common language of Work Roles and knowledge, skills, and abilities (KSAs) that support cybersecurity roles (Newhouse et al., 2017). The NICE framework facilitates workforce development in several ways, including: (1) describing the elements of work common across organizations, sectors, and situations, defined as Work Roles and Tasks, and (2) describing what defenders know and do, defined through KSAs. In doing this, the NICE Framework also facilitates understanding of cognition by describing what cybersecurity professionals know (as knowledge) and do (as skills and abilities) for a given situation (as Work Roles and Tasks).

Despite the critical role of people in Protect and Defend roles, relatively little is known about the cognitive processes, the ways of thinking, that enable cyber defenders to perform their work. This problem is a knowledge gap between what highly proficient defenders do and how they are able to do it. A lack of understanding of cognition limits the ability to leverage human capabilities in cybersecurity work.

Research Questions and Contribution

This article addresses two research questions: (1) what cognitive processes support proficiency in defender roles? And, (2) how can understanding of defender cognition be used to complement the NICE Framework and develop the workforce? A research methodology, focused on the cognitive processes of individual defenders, is proposed to address these two questions.

Cognitive processes in cyber defense can be understood at a macrocognitive level, the result of cognitive resources at work in operational environments (Klein et al., 2003). Macrocognitive models have been used in other complex, high stakes environments, including transportation and healthcare to improve practice by describing how cognition supports performance (Klein & Wright, 2016). Cognitive task analysis (CTA) is a collection of qualitative research methods to develop and apply macrocognitive models to improve practice (Klein & Wright, 2016). The foundation of the research methodology presented in this article is the use of CTA to explain defender cognition. The methodology complements existing CTAs in that it uses CTA to describe the how cognitive processes of individuals support cyber defense with the aim of generalizing those processes to related tasks across situations and organizations.

Rationale

As the NICE Framework (Newhouse et al., 2017) provides a common language to further professionalize cybersecurity careers, so too could a better understanding of cognition further professionalize cybersecurity careers by connecting cognitive processes across situations, Tasks, and Work Roles. Identification of cognitive processes that apply to cybersecurity Work Roles addresses the gap between cognition and performance. It can describe the ways of thinking that predict success in identification, analysis, and mitigation of threats. This knowledge can be used for workforce development, specifically, in recruitment and proficiency development, through training. Recruitment and selection, and employee development and retention, are two of four components of the human capital management lifecycle, and they complement workforce planning and succession planning in the draft NISTIR 8193, the NICE Framework Work Role Capability Indicators (Stein et al., 2017).

As a recruitment outcome, potential cybersecurity professionals might be recruited on the basis of cognitive skills, which could be used to increase the diversity of people participating in the profession. That is, interests in other domains that utilize similar cognitive skills might suggest talent and interest in cybersecurity careers. This strategy is already being used informally. For example, the career website DICE lists “grasping the big picture” as one of six skills needed to succeed in cybersecurity (“Six skills,” n.d.). Research on defender cognition could provide further specificity and empirical evidence for these claims. It could also accelerate training, such that training of known cognitive skills readies people for a cluster of cybersecurity Work Roles, in complement to other training strategies. Training involves a process of learners acquiring KSAs; this can be done more efficiently and effectively with understanding of how such knowledge and skills are used in the field. Together, the cybersecurity workforce could be further developed through increased career participation and novel, evidence-based training strategies.

Investigative Approach

The purpose of this article is to provide a framework for researchers and practitioners to help them leverage understanding of cognition for workforce development. In the background section that follows, macrocognition’s role in describing cognition at a useful layer of abstraction and complement to the NICE Framework is introduced. CTA is discussed as a method of understanding proficiency to suggest interventions for workforce development, and CTAs relevant to this perspective are reviewed. Following this, a use case with two industry defenders is presented as an example of how CTA could be used to describe macrocognition. Immediate, albeit limited, practice implications of the use case are presented. As a use case, more value comes from lessons learned; increased participation in research on the human aspects of cybersecurity is necessary for the successful application of the methodology, and the lessons learned can accelerate implementation of the methodology. The article concludes with a discussion of the research and practice implications of the methodology on workforce development.

II. Background

The NICE Framework and Cognition

The NICE Framework was developed through an iterative process of participation from industry, government, and academic stakeholders with opportunities for public input (Newhouse et al., 2017). One outcome has been the definition of KSAs and Tasks mapped to Work Roles. KSAs help explain defender cognition by enumerating what defenders know and what they are able to do. The acquisition of knowledge and skill in a domain supporting high levels of performance are components of high proficiency, also called expertise (Feltovich, 2018).

Understanding of defender cognition can be used to augment the NICE Framework by explaining how proficiency develops and is used in operational settings. Knowledge in the NICE Framework is defined as, “a body of information applied directly to the performance of a function” (Newhouse et al., 2017, p. 5). Knowledge elements, for example as “Knowledge of encryption algorithms” (p. 59) refer almost exclusively to declarative knowledge or to the use of technologies (e.g., “Knowledge of virtualization products (VMware, Virtual PC),” p. 76). Skills are distinguished by being observable, defined as “observable competence to perform a learned psychomotor act.” For example, “Skill in using knowledge management technologies” could, with further specification, be demonstrated. Abilities are similarly defined as observable behaviors, albeit behaviors that result in an observable product. At present, a draft revision to the NICE Framework refactors Skill and Ability statements into Skill statements, defined by observability (Petersen et al., 2020). Knowledge and Skills describe capabilities of the defender; Tasks describe the work (Petersen et al., 2020). Given this, Skills and Abilities will be considered together and distinctively from Knowledge. Given its prevalence in the literature, the abbreviation KSA will be used to refer to the collective of Knowledge and Skills/abilities.

What is not provided by KSAs is understanding of *how* defenders are able to perform or demonstrate the KSAs. Understanding the cognition used by defenders in performance of their work would help the field understand how their proficiency develops. This requires an understanding of how interactions of NICE Framework Task requirements (Newhouse et al., 2017), defenders' thought and behavior, and the organizational and world context affect outcomes. One challenge to understanding cognition is that it is not directly observable; it occurs within an individual. This may be why Woods and Roth (1988) observed that the development of a technology generally outpaces our understanding of how to best use the technology.

Proficiency in Cyber Defense

Research on expertise across domains suggests that proficiency may be identified by differences in the process of thinking and reasoning (Feltovich et al., 2018). Proficiency goes beyond holding a greater amount of knowledge in a domain; it is also evident by thinking in different ways. Decades of expertise research have shown that, generally, highly proficient individuals can process larger and more integrated cognitive units, have deeper and more functional representations of tasks, are better able to apply their knowledge to problems, are more able to engage in self-monitoring, and can better recognize patterns in problem solving (Feltovich et al., 2018). To address the skills gap, a critical question is how to develop the proficiency of defenders efficiently.

The thinking of proficient defenders is not well understood, and defender proficiency is only starting to be defined. Agyepong et al. (2020) conducted a systematic literature review to identify challenges to the use of metrics in security operations centers. They found a combination of metrics that are objective and easily captured but limited, such as number of alerts analyzed, and the number of tickets closed per day (Agyepong et al., 2020). These contrast with metrics that are more comprehensive but subjective, and difficult to capture, such as the quality of analysis and quality of incident reports (Agyepong et al., 2020). They concluded that, "An understanding of how analysts addressing the difficult aspects of their work can be used will provide insights into their performance" (Agyepong et al., 2020, p. 14).

The NICE Framework Workforce Indicators (Stein et al., 2017), describe proficiency at three levels: entry, intermediate, and advanced. These levels are distinguished by the level of knowledge, the ability to perform successfully under limited guidance, the ability to serve as a resource for others, and the ability to perform successfully "in complex, unstructured situations" (Stein et al., 2017, p. 302). These levels map to proficiency categories Hoffman et al. (2014) adapted from craft guild terminology. The entry level of Stein et al. (2017) corresponds to the apprentice, a student who is working within the domain (Hoffman et al., 2014). The intermediate level of Stein et al. (2017) is referred to as journeyman by Hoffman et al. (2014), a person who can perform a day's work without supervision. The advanced level of Stein et al. (2017) corresponds to the expert level of Hoffman et al. (2014), distinguished with high regard from peers, the highest levels of task performance, and the ability to respond to rare or complex situations. This comparison usefully extends the NICE Framework Workforce Indicators (Stein et al., 2017) because Hoffman et al.'s (2014) characterization includes two additional levels of interest. The naïve individual is ignorant of the domain. This is the target population for interventions that bring new people into the domain. As a population, getting individuals to advance from naïve to higher levels means increasing participation in the domain. On the other end is the master level, distinguished by qualifications to teach. While there exist cybersecurity practitioners at all of these levels, the field currently has limited understanding of how to label and develop proficiency.

In several cases, general models of proficiency have been applied to cyber defense with limited success. One missing element is the understanding of cognition in context, specific to cyber defense. Ben-Asher and Gonzalez (2015) developed a survey to classify participants as cyber defense experts or novices based on domain knowledge and self-reported experience. Their novice category included individuals with no cybersecurity experience and more closely aligns with the naïve category. They observed differences in the dichotomized groups but relatively small differences in the ability of the two groups to detect attacks in a simulated task (Ben-Asher & Gonzalez, 2015). The authors concluded that pulling professionals from their operational environment to participate in the study may have resulted in fewer cues available to them, suggesting that defenders leverage cues from the operational environment that are not easily captured in simulation-based experiments (Ben-Asher & Gonzalez, 2015).

Saner et al. (2016) aimed to identify naïve individuals who may excel in cybersecurity careers based on relevant cognitive skills. They used an approach of inferring cognition from NICE Framework KSAs and then placing them in a framework with two dimensions. The first dimension was the difference between initiating (e.g., attacking) and responding (e.g., defending; Saner et al., 2016). The second dimension was the difference between real-time roles requiring action under time pressure and exhaustive roles allowing for greater planning and deliberation (Saner et al., 2016). This characterization provides a link between tasks and cognitive demands required to perform them. However, they struggled with the linkage from KSAs to cognition, noting that too few details about the steps involved in the operations were available in KSAs to make inferences about cognition (Saner et al., 2016). Together, these studies have established a need to better understand cognition of proficient defenders and suggest an approach that incorporates the complexity inherent in the work. Towards developing understanding of this context, Goodall et al. (2009) conducted a field study and found evidence of two aspects of defender proficiency: technical knowledge of the domain and knowledge of the specific network environment involved, which they called situated expertise. While efforts are emerging to label, measure, and understand defender proficiency, better understanding of human cognitive performance in defender Work Roles is needed.

Macrocognition: Managing Complexity

Rasmussen et al. (1990) recognized the issue of levels of abstraction in describing complex work. Different levels of abstraction provide different perspectives on the work. Models at a low level of abstraction are closely linked to specific circumstances in the physical world (Rasmussen et al., 1990). Models at higher levels of abstraction are closely linked to a specific purpose (Rasmussen et al., 1990). Rasmussen et al. suggested that a work function can be seen as a goal for a lower level of abstraction and an explanation for how higher levels of abstraction are realized (1990). The levels of abstraction, as summarized by Crandall et al. (2006), were goals, measures of the goals, general functions and activities, specific functions and activities, and workspace configuration. The work of defenders can be understood at each of these levels, from overall goals at the top, to physical processes at the bottom. Defending the organization is a high-level goal, supported by measures of the goals and general functions. The use of a tool, as a specific activity, explains how a general function is realized. Thus, understanding a task involves mapping lower level operations to higher level goals.

By analogy, cognition can be understood at various levels of abstraction. Individual cognition can be understood at micro (low) and macro (high) levels of abstraction. At the micro level, of interest to cognitive psychologists, cognition is partitioned into resources, such as attention, perception, and memory, so that their function can be understood. A common assumption in cognitive psychology research is that cognitive resources can be understood independent from their cultural and societal context (Braisby & Gellatly, 2005). The benefit of this approach is that it can identify laws, which apply universally; however, because it is separated from a specific purpose, it leaves unanswered questions about how cognition supports goals and tasks.

The macro level, called macrocognition, is the result of cognitive resources at work in operational environments, stepping towards a specific purpose, incorporating context, and emphasizing a descriptive approach over normative models (Klein et al., 2003). Macrocognition is a set of cognitive functions that support human performance. Microcognition includes the mental operations that explain macrocognition. Klein et al. identified six macrocognitive functions (see Table 1) in support of six macrocognitive processes (see Table 2; 2003). In their model, goals drive the use of macrocognitive functions. Macrocognitive processes are employed to support the macrocognitive functions. Importantly, macrocognition does not attempt to describe physiological circumstances or operations. There is no physical mental model, for example. Instead, mental models reflect an application of cognitive operations towards a specific purpose (Klein et al., 2003). This is why macrocognitive constructs are defined by the outcomes they support; mental models are mechanisms to "generate descriptions of system purpose and form, explanations of system functioning and observed system states, and predictions of future states" (Rouse & Morris, 1986, p. 351).

Situation awareness (SA) is another macrocognitive construct frequently of interest in research on defender cognition (for a review, see Gutzwiller, 2019). In simple terms, it is the outcome of the process of sensemaking (see Table 1), which is also called situation assessment. Situation awareness is most frequently defined according to Endsley's (1988) model, involving perception of relevant elements in the situation, comprehension of the elements in the situation, and projection of the future status of elements. The relevance of elements in the specific situation is what distinguishes macrocognitive SA from microcognitive perception. The content of SA is defined by the goal and situational context. Because of this, the SA of an airline pilot cannot be meaningfully compared to the SA of a defender. Separated from a situational context, SA has no meaning. This also requires measurement of SA to be goal specific. This challenge may be one of the causes for dilution of this term in the literature. As Gutzwiller et al. (2016) noted, authors have increasingly been using situation awareness as a term to define the result of data fusion rather than a macrocognitive process. They suggest cyber-cognitive situation awareness (CCSA) to describe defender SA according to Endsley's (1988) model. Although this term is not yet widespread in the literature, using CCSA instead of SA identifies it as a construct of human macrocognition.

Macrocognition connects cognition to tasks and goals. But because macrocognitive processes and functions are specific to context, they need to be specified for cyber defenders. The last needed piece is a method to understand macrocognition applied to tasks. As a collection of applied, qualitative research methods, CTA provides this piece.

Table 1

Macrocognitive functions identified by Klein et al. (2003) with definitions quoted and adapted from Crandall et al. (2006)

Function	Definition
Problem detection	Spotting “potential problems at an early stage” (p. 139)
Coordination	“The way team members orchestrate the sequencing of their actions to perform a task” (p. 139)
Adaptation	“Modifying, adjusting, or replacing a plan that has already been implemented” (p. 138)
Planning	“Modifying action to transform a current state into a desired future state” (p. 138)
Sensemaking / situation assessment	Diagnosing “how the current state of affairs came about” (p. 138)
Naturalistic decision making	Relying “on experience to identify a plausible course of action” (p. 137)

Table 2

Macrocognitive processes identified by Klein et al. (2003) with definitions quoted and adapted from Crandall et al. (2006)

Function	Definition
Managing attention	Using “perceptual filters to determine the information a person will seek and notice” (p. 142)
Identifying leverage points	“Identify opportunities and turn them into courses of action” (p. 141)
Managing uncertainty and risk	Developing “skills for coping with uncertainty” (p. 141)
Mental simulation and storybuilding	“Enacting a series of events and pondering them as they lead to possible futures” (p. 141)
Developing mental models	“How sense is made of situations” (p. 140) involving “mental imagery and event comprehension” (p. 140)
Maintaining common ground	“The continuous maintenance and repair of calibrated understanding amongst members of a team” (p.140)

Cognitive Task Analysis (CTA)

CTA is a collection of "tools and techniques for describing the knowledge and strategies required for task performance" (Schraagen et al., 2000, p. xiii). Klein and Militello (2001) explained CTA in terms of its description of cognition, focus on tasks in natural settings, and attempt to explain the cognitive processes observed. CTA involves three main components: knowledge elicitation, data analysis, and knowledge representation (Crandall et al., 2006). Knowledge elicitation involves data collection with practitioners in the domain. Data analysis is the process by which the researcher synthesizes data and discovers meaning. Knowledge representation summarizes the meaning uncovered in data analysis.

Many CTAs have been conducted with defenders, though relatively few of the CTAs have provided data on individual defender macrocognition in organizations outside of government and defense. In the methodology, past CTAs inform future CTAs and suggest research questions for quantitative research. Integrating CTAs can be challenging. Because they are not testing theory, each provides its own insights and a complementary glimpse into the cognitive work of defenders. There are also a wide variety of techniques available (for a review, see Wei & Salvendy, 2004).

Categorizing CTAs based on their knowledge representation can be useful. Some CTAs have resulted in knowledge representations that are workflow-oriented (e.g., Erbacher et al., 2010). The CTAs that are most aligned to the methodology resulted in a list of goals and subgoals of defenders and/or reflect decision making through the questions asked by defenders (e.g., Buchanan et al., 2016). It should be noted that, although outside the scope of this review, work has been done to understand processes at the team-level (e.g., Cooke et al., 2013; Nyre Yu, 2019; Tetrick et al., 2016).

Erbacher et al.'s CTA (2010) resulted in a workflow representation. They conducted a seven-phase CTA for the purpose of developing visualization techniques. Participants included network analysts, network managers, and security researchers at Pacific Northwest National Laboratory. This work resulted in a task-flow diagram in four stages: assessment, detailed analysis/cleanup, response, and audit (Erbacher et al., 2010). The model is circular, reflecting an iterative process of re-assessment. It features a big picture construct at the center, supporting all steps and being affected by cleanup. The big picture includes a variety of constructs, including world view, known players, cyber-attacks, host information, and coordination (Erbacher et al., 2010).

Trent et al. (2019) used CTA to describe the workflow of US military cyber protection teams. As with Erbacher et al. (2010), this CTA emphasized the high-level process of the work. It was also idealized in that some steps are skipped in practice (Trent et al., 2019). One theme of the model is that the work does not necessarily proceed in sequence (Trent et al., 2019). Rather, the work, "needs to be described in terms of parallel tasks and feedback loops, not as a series of steps or stages" (Trent et al., 2019, p. 129). The role of periodic communication with intelligence sources was also highlighted.

Narrowing to CTAs that help explain the cognition of individual defenders provides a more succinct list. Many CTAs aimed to describe the content of CCSA for defender tasks. Utilizing a number of methods, including an extended period of observation, Paul and Whitley (2013) investigated how analysts establish and maintain awareness of large computer networks. They suggested two components: event detection and event orientation (Paul & Whitley, 2013). In a review of such studies through 2015, Gutzwiller (2019, p. 41) noted that CCSA needed to be defined for particular roles, measurement was still needed, there was limited understanding of the linkage between defender CCSA and performance, and there was a need for research to understand other macrocognitive functions and processes.

Some CTAs focused on defining the categories of defender CCSA. In an early example, Biros and Eppich (2001) suggested categories of recognition of nonlocal Internet protocol (IP) addresses, identification of source IP addresses, development of a mental image of normalcy, creation and maintenance of analyst situational awareness, and facilitation of knowledge sharing. D'Amico et al. (2005) described detection, situation assessment, and threat assessment being developed in a largely linear process of building understanding. D'Amico and Whitley (2007) developed a three-stage process model mapped onto Endsley (1988); CCSA was represented as a hierarchy of raw data being filtered to leave what is interesting, then what is suspicious, then events, then incidents, and finally intrusion sets, which are groups of related incidents (D'Amico & Whitley, 2007).

D’Amico et al. (2005) represented questions asked by defenders, the first example of an approach used by others. They also suggested site-specific knowledge as a challenge to proficiency development; defenders must know what is normal for their environment (D’Amico et al., 2005), a theme that emerged in other sources, such as Goodall et al.’s (2009) situated expertise. Mahoney et al. (2010) created six preconstructed scenarios, which they discussed with a single subject-matter expert. The result of the CTA was a list of nine preliminary categories of CCSA. Buchanan et al. (2016) conducted a goal-directed CTA to elicit the subgoals and decisions, also phrased as questions, under two high level goals: detecting threatening incidents and characterizing those incidents. Describing the content of CCSA as lists of questions can be understood as an answer to the question, *what should a defender attend to?* This level of abstraction can reveal cognitive process (i.e., how CCSA works for defenders) and offers implications for practice.

The approach of Zhong et al. (2015) represented defender behavior at a lower level of abstraction. They used automation to capture defender behavior in combination with participant self-reporting, which suggested 11 operations: browse, filter, search, inquire, select, selected, link, new hypothesis, modify hypothesis, switch context to a different hypothesis, and confirm or deny a hypothesis (Zhong et al., 2015). These operations could suggest building block operations of decision making but are less connected to goals; additional context is needed to connect the behaviors to goals.

Gutzwiller et al. (2016) conducted a CTA with six participants in three phases. They combined a semi-structured interview, a knowledge audit, and a concept mapping activity. Gutzwiller et al. provided a three-component model of CCSA: understanding and awareness of the network, the team, and the world (2016). The network includes elements of network architecture and behavior (Gutzwiller et al., 2016), which map onto elements of Mahoney et al.'s (2010) CTA. The world component includes awareness of novel threats and abnormal behavior (Gutzwiller et al., 2016). Finally, the team component represents awareness of team members to facilitate coordination and support (Gutzwiller et al., 2016). This model differs from others in that it gives more prominence to contextual factors.

This section has described how CTA has been used to better understand the work of defenders. CTAs in this area have contributed to understanding how cognition supports this challenging work. Despite this, more research is needed to describe how macrocognition of individual defenders in industry works to affect security outcomes. A methodology to achieve this aim is described next, followed by a case study to illustrate how the methodology can be applied to the work of defenders.

III. Research Methodology

--- Insert Figure 1 here ---

Figure 1 depicts the research methodology using CTAs to understand the cognition of individual defenders. While no single link in the methodology is novel, the methodology augments current research methods by suggesting mutual support between CTA and quantitative research. In the next section, a use case illustrates how the methodology can be applied to understand defender cognition.

IV. Use Case

Method

Participants

Two participants were cybersecurity professionals at technology companies in the San Francisco Bay area. Organizations with defenders were invited to participate in the research through convenience sampling. Organizations were discovered through professional networks and by attending cybersecurity-related conferences. CISOs and managers were contacted by e-mail to discuss the study. Upon organizational agreement to participate and IRB approval, organizations were invited to announce the study to employees, who participated voluntarily as part of their workday.

Defenders were defined broadly as professionals who monitor networks and/or respond to network threats on a daily basis. While defenders were the focus of the investigation, the inclusion criteria were broadened to any employee in an operational cybersecurity role to encourage participation of cybersecurity professionals regardless of their job title.

Of the 19 individuals who participated, two met our definition of a defender, resulting in a sample size of *N* = 2 from two different companies. One company was a large networking technology company. Participant 1's company employed between 50,000 and 100,000 individuals at the time of the interview. Participant 2's company provided cloud services and had between 100 and 500 employees. One interview was conducted in early 2018; the other was conducted in early 2019.

Materials

Participation involved two activities, a survey and a concept mapping interview. Both of these activities were exploratory, with the goal of identifying themes appropriate for future investigation. The purpose of the survey was to describe the expertise of the participants and qualify them for inclusion in the study. The purpose of the concept mapping interview was to elicit the knowledge of participants and represent it visually.

Exploratory Survey. A survey was used to qualify cybersecurity professionals as defenders and document their experience. Participants were sent a link to complete the survey using Qualtrics. The survey items were adapted from a survey used with an earlier sample (Schuster & Wu, 2018). As part of the survey, participants were asked their gender, age, job title, years they have been in their current role, and total number of years of experience working in cybersecurity. The survey then asked whether they respond to network threats on a regular basis, the highest level of education obtained, the name of degrees obtained, and certifications held.

Concept Mapping Interview. An individual concept mapping interview was conducted with each participant lasting approximately 60 minutes. The concept mapping protocol was adapted from Crandall et al. (2006). The interview was conducted over teleconference using audio and screen sharing but without video. IHMC Cmap Tools Knowledge Modeling Kit (version 6) was used to create the concept maps, and this software was made visible to the participant during concept mapping using screen sharing.

Members of the research team included the author and student research assistants. Training for the research team involved study of Crandall et al. (2006), review of the research protocol, and participation in mock data collection sessions with other members of the team. Data collection required three researchers. The facilitator led the interview and was the only member of the research team in regular communication with the participant. Meanwhile, a second researcher operated the concept map software, and a third researcher took notes. Participants did not manipulate the concept map software directly. To facilitate the mapping, the concept mapper could interject to slow or repeat parts of the interview. The notetaker took notes without interacting with the participant.

After an introduction of the members of the research team, the informed consent notice was displayed and discussed with the participant. Next, the concept map activity was introduced with an example using driving as a domain. This tutorial introduced concepts, linking words, and propositions. Concepts are the major concepts in the domain. Concepts are connected to other concepts by linking words. Together, two concepts form a complete sentence with a linking word, called a proposition. Propositions are directional and are read in the direction of the arrow. Participants were provided an overview of the process of generating the concept map and given a suggested list of linking words suggested by Crandall et al. (2006, p. 60).

Following the introduction and tutorial, participants were presented with a focus question designed to anchor the concept map. The guiding priority was to represent the participants' individual perspective of their work, not the work of their company as a whole or the basics of the field. Therefore, the map was anchored with a focus question of "What do you need to be aware of when monitoring for and/or responding to threats?" This question proved insufficient when the sample started to include non-defenders. Thus, the approach was modified to instead co-create a focus question tailored to the participant's job title. The focus question was structured in the format, "What do you need to be aware of when…," with a job description following. The default question was "What do you need to be aware of when creating secure network systems?" Participants were then asked to confirm that the question applied to their daily work. In both interviews, participants chose to revise the question. The focus question was collaboratively revised to, "What do you need to be aware of when responding to cyber security incidents?" for Participant 1 and, "What do you need to be aware of when supporting network security systems?" for Participant 2.

The revised focus question was entered as the highest-level concept in the map. From here, the concept map was constructed in four general steps. The first step was initial concept generation, in which participants were asked to list the most relevant concepts that came to mind after reading the focus question. Participants were asked to, "identify the most relevant concepts that you think of when you read the focus question" and that "it is important to know that these concepts are not final, and you can choose to add, remove, or change any concept at any moment. We will write them down as you say them aloud." The second step was to organize the terms of the map. Participants were asked to suggest the most general or important concepts, which were then moved toward the top of the screen. The goal of this step was to organize concepts so that they were descending from more general to more specific at the bottom. The third step was to link concepts, starting with one relationship. Participants generated propositions by forming a complete sentence from one concept to another, with the linking words in the middle. The fourth step was the refine the map. In this step, the facilitator navigated the map, reading and confirming propositions aloud. Before concluding, participants confirmed that they were satisfied with the map representation.

Results

Survey Responses

Participant 1 held the title of InfoSec Tier 2 analyst, worked in incident response, and had been in the role for five years. Participant 1 reported eight years of experience in cybersecurity. This participant held security-related bachelors and master's degrees. Participant 2 held the job title of security analyst, supported a web application firewall, and had been in that role for a single year. This participant reported three total years of experience in network security. Participant 2 participant held a non-security bachelor's degree. In terms of the NICE Framework, Participant 1's work mapped to the cyber defense incident responder role, while Participant 2's work mapped to the cyber defense infrastructure support specialist. Participants held two or three certifications each with no overlap.

In all, the two participants differed in the duration of their work experience and job title. Participant 1 had more experience in the current role and the field while working for a much larger organization. The participant from Company 2 had less experience in the current role and the field while working at a smaller organization.

Concept Map Analysis

The purpose of the concept map analysis was to identify elements of macrocognition in the work of defenders. This was done by analyzing map structure and content and identifying macrocognitive themes.

Map Structure and Content. Following the interview, concepts on the maps were adjusted so that all propositions were visible and directionality clear. Concepts that were listed by the participant but not used in the map were removed. Participant concept maps are shown in Figure 2 and Figure 3. As a first step, the maps were examined to see if they differed in the quantity of concepts and propositions. Concept frequency was assessed by listing concepts used in each proposition. If a concept was connected to more than one other concept, it was counted each time it appeared in a proposition. For example, “asset targets include cloud asset” and “asset targets include user asset” were counted separately.

The two maps were different in their complexity. Comparing across the two maps, Participant 1 generated a greater number of concepts (45) and propositions (46) than Participant 2 (23 and 36, respectively) but fewer links per proposition. The following concepts were common to both maps: Logs (as logs or web logs), assets (as high value assets, user asset, cloud asset, ownership of the asset, asset targets, host asset, or data asset), and monitoring (as monitoring or established monitoring plan deployed within your infrastructure).

--- Insert Figure 2 here ---

--- Insert Figure 3 here ---

Participant 1’s map was organized around the attribution of source and target and the meaning of the event. For Participant 1, assets were considered in the context of the attack, and four asset types were specified (cloud, user, host, and data) and associated with logs (as the target of the log). For Participant 2, the major elements were the infrastructure and the organization. Security controls, which include monitoring and segmentation, protect the infrastructure. For Participant 2, assets were part of infrastructure, although they are at the same level as the network, network applications, and hardware. By including it as a separate concept, the map suggests asset value is a relevant factor.

Evidence of Macrocognition. The concept maps were examined to identify macrocognitive themes. Using the definitions described in Table 1, two members of the research team independently identified potential macrocognitive elements in each concept map. The two lists were aggregated, and propositions using these concepts were examined. This process revealed themes for each participant. Table 3 lists collective themes of macrocognition created by aggregating the participant-level themes. Additionally, practice and research implications are suggested for each theme.

Table 3
Macrocognitive Themes and Implications

Theme	Implications
1. Defenders maintain mental models of assets: the elements of the infrastructure (such as targets, zones, and events) and their business context.	**Practice implication**: Defenders need to develop an understanding of relevant infrastructure. Defenders may benefit from awareness of other organizational perspectives so that they can best incorporate business context into their understanding and decisions. **Research questions**: How do defenders' representations of elements of the infrastructure relate to the reality of these resources? How do defenders' representations of elements of the infrastructure evolve over time? How do organizational goals and culture affect mental models of assets?
2. Selecting information from asset targets is an example of managing attention.	**Practice implications**: As defenders develop proficiency, they may be identified by their ability to manage attention; in support of CCSA, proficient defenders understand where to find needed information. **Research questions**: How do defenders learn and know where to find needed information?
3. Technology contributes to sensemaking by filtering and providing perceptual cues.	**Practice implications**: Technology affects sensemaking in the quality and relevance of cues provided. Defender proficiency may be identified by the ability to put relevant cues together to develop meaning. **Research question**: How do defenders adapt to the limitations of imperfect cues?
4. Sensemaking is a prerequisite for remediation but continues during remediation. Sensemaking in remediation pulls in organizational,	**Practice implications**: Defenders need to be able to draw connections between the organization, infrastructure, and situation. They must understand how the situation relates to the organizations' goals.

infrastructural, and situational factors.	**Research questions**: How does sensemaking before remediation relate to, and differ from, sensemaking during remediation? At boundaries between tasks, how does macrocognition change?
5. Event detection is an example of spotting anomalies.	**Practice implications**: Defenders could learn the critical cues experts use to spot anomalies. Defenders' use of critical cues may suggest strategies for detection of novel events. **Research questions**: What are the critical cues that hint at anomalies? How are previously unknown anomalies detected by defenders?
6. Planning occurs throughout the remediation process. Adaptation occurs during remediation when implementing procedures to account for time lost in remediation.	**Practice implications**: Defenders mental models help them adapt to the needs of the situation. Adaptation requires understanding of what the situation means for the organization. **Research questions**: Where is adaptation situated? How are handoffs managed? How do defenders decide to alter their plans?
7. Knowledge, capabilities, and skills of people in the organization are leverage points.	**Practice implications**: Organizational factors beyond the team affect the quality of their cyber defense because they are resources that can be drawn upon by defenders. All parts of the organization are potential resources in response; others in the organization may need awareness of what defenders do to maximize this resource. **Research questions**: How are knowledge, capabilities, and skills of people leveraged by defenders? How do differences in organizations affect how human resources are leveraged by defenders?

V. Discussion

The use case shows how CTAs can be used to learn about the macrocognition of defenders. Macrocognition informs the practice of cybersecurity by connecting the KSAs and Work Roles described in the NICE Framework to the strategies and processes of individuals. Implications and lessons learned through the use case are discussed next. Finally, implications of the methodology are presented to distinguish the potential of the larger methodology from the limitations of the specific use case.

Implications of the Use Case

The macrocognitive themes affirm findings from prior CTAs that CCSA, as developed through the process of sensemaking/situation assessment, is a salient component of defenders work to detect, analyze, and respond to events. The themes also suggest the importance of mental models; in combination with domain knowledge, defenders make use of a variety of continually evolving representations of interconnected elements. These include contextual factors such as who owns the asset and the business context (e.g., how people in the organization affect an asset). This representation extends the CCSA types of both Mahoney et al. (2010) and Gutzwiller et al. (2016). Mahoney et al. (2010), incorporated the business context in compromise extent awareness and situational factors as social/organizational/behavioral awareness. Gutzwiller et al. (2016) represented these as team and world components. Based on the present representation, these factors could broaden to infrastructure (extending beyond the network) and organization (extending beyond the team).

The biggest limitation in interpretation of the use case as standalone research is the combination of a small convenience sample with limited time with each participant. Most past CTAs enjoyed larger samples and/or more comprehensive observation. In this case, the immediate value of the use case is the suggestion of themes for further investigation. A number of practical lessons are discussed next.

Lessons Learned

Lessons were learned at each step of the use case, starting with the recruitment. Encouragingly, when organizations responded to discuss the study, there was no shortage of support expressed for research to improve the proficiency of cybersecurity professionals. Being affiliated with San José State University and working on a National Science Foundation project may have facilitated participation. Managers and CISOs were generally eager to offer their teams for participation, but barriers to participation included concerns about confidentiality and limited access to cybersecurity professionals. In discussing the aims of the study, members of the research team emphasized the goal of the research in describing cognition but not describing specific incidents or threats. As the use case illustrates, there is value for methods which do not elicit information about specific incidents; however, this also constrained the research, as many CTA methods rely on discussions of specific cases, especially challenging ones, to understand how experts address them. Finding cybersecurity professionals in the correct role was also a challenge, evidenced by two participants despite a multiyear recruitment effort. In some cases, defenders were outsourced or distributed worldwide, making access more difficult. This required a pivot to remote data collection. Future research can address these limitations by seeking more depth within one organization through close partnerships, rather than trying to sample across many organizations. This could allow more precision in job role selection. Researchers should be aware that building the collaborative partnerships for recruitment requires substantial time. Meanwhile, the research needs to be ready to run on a short timeframe, as short windows for participation can appear after long delays from due diligence and participant availability. This can challenge an academic research team due to the seasonality of the academic calendar.

While the survey roughly identified participants as cybersecurity professionals in a defender role, it provided a limited picture of the training and experience of cybersecurity professionals. Informal and self-directed learning also contribute to the experience of many cybersecurity professionals, including in childhood (Champion et al., 2014), and this information was not captured. Beyond degrees earned, information about experience outside of cybersecurity was not captured. A better survey would be more comprehensive in describing the cybersecurity and non-cybersecurity education and experience of the participants. A succinct survey is also desirable; using the NICE Framework's Categories and Specialty Areas (Newhouse et al., 2017) may provide a way to quickly categorize participants. Categories and Specialty areas are presently deprecated in the NICE Framework Revision draft (Petersen et al., 2020); an alternative approach could be to sample representative work roles and ask about frequency of performing in those roles.

In future research, a more comprehensive survey could allow researchers to unify the focus question across maps. In hindsight, the research team started each concept map interview with limited information about defenders' experience, so revising the focus question was necessary to understand the participant's role. Differences in experience of the participants may be evident in the content and organization of the concept maps. Participant 1 reported more years of experience and drew a more detailed map oriented around decision making. Participant 2's map suggests a greater role of automation, but inferences comparing the proficiency of these individuals is speculative, largely due to the limitations of the survey.

Finally, concept mapping worked well for this purpose but is not the only CTA technique. Concept mapping resulted in a knowledge representation that could be generated in a single meeting. Limited time and access prevented use of certain CTA methods that may provide more clarity into aspects of macrocognition or teamwork. These include direct task observation and critical decision methods. The critical decision method involves a subject-matter expert retroactively or concurrently walking the researcher through a specific incident (Crandall et al., 2006). Examination of standard operating procedures and lists of specific tools were also unavailable. Crandall et al. (2006) suggested that the representations from CTA be refined in a collaborative process with the original participants; this was only done immediately after constructing the map. A second interview with our participants to vet our conclusions was not part of the research protocol, resulting in a lack of data about how the resulting representation is viewed by the participants.

Research Implications of the Methodology

This work serves as a call and framework for continued CTA research to understand cognitive processes. By applying the methodology, quantitative researchers can benefit from better understanding of relevant contextual factors, which can lead to more meaningful experimentation and establish reliable measurement. For example, by enumerating and refining the types of CCSA, a richer picture of proficiency development in defenders can emerge. The methodology can explain the cognition behind proficiency. This can suggest strategies for its measurement, which can be developed and validated through quantitative research. Measured reliably, cognitive process measures would complement measures of KSAs in the NICE Framework to describe what defenders do and how they are able to do it. This could facilitate training for defenders in how to think, addressing a gap in the current state of practice.

The use case shows how research questions could be tested in quantitative research. CTAs are resource-intensive and can generate more questions than they answer. New questions can be supplied to quantitative researchers, who need testable hypotheses and the ability to isolate a limited number of variables of interest. This work suggests variables that may be able to be measured, such as CCSA for an incident; experimentation is needed to establish reliable measurement. Hoffman (2020) described challenges to conducting experiments to understand cyber operations including accounting for all important variables, especially ones embedded in the operational environment. As part of the methodology, researchers can use existing CTAs to list and prioritize variables for experimentation. Hoffman (2020) suggested a mixed-method approach of implementing CTA methods into their experimental protocols, essentially giving participants the opportunity to explain the cognition used in an experimental task. When this is not feasible, the methodology suggests how CTAs of defender cognition and quantitative studies may inform each other.

CTAs have great potential to augment each other. As an example, rather than offering a focus question tied to a person’s role, a concept map could be made of asset targets, zones, and events. Insights can be derived even if participants reject the categorization and describe their own categorization and mental model. Further themes may emerge by analyzing concept maps with the same prompt from employees with, near, and outside a defender team.

Practice Implications of the Methodology

Defenders are the innovators in the practice of their profession, and this is especially evident in their central role in CTA. The scientific study of situation awareness emerged from interviews of pilots (Endsley, 1988), who discussed it as a familiar concept. The methodology shows how research can contribute to professionalization of cybersecurity careers by informing practice through understanding the cognition of high performing individuals.

The methodology and use case offer both near-term and long-term implications for defenders. In the near term, defenders could improve their practice by focusing on what happens before decisions are made. Outcomes are more salient than cognitive processes, which are not directly observable. Given this, there is value in discussing the hidden-yet-valuable role of macrocognition in performance outcomes, even as understanding of how it works is still emerging. That is, talking about macrocognition could help defenders connect how they think with their work outcomes. Encouraging discussion and debate of macrocognition among practitioners may help them reflect on their own thinking and development, a process called *meta*cognition. Metacognition has been shown to predict performance and training effectiveness (Cuevas et al., 2004). Thinking about their thinking may help defenders better apply their skills and match them to the KSAs in the NICE Framework. This may help proficient defenders mentor novices. In this way, individual defenders themselves will continue to contribute insights about how their work is conducted.

In the long-term, much interdisciplinary research is needed to realize the aims of the methodology, which itself is only one perspective on the work of cybersecurity professionals. In this regard, the methodology serves as a call for replication and participation by researchers and practitioners within and outside of cybersecurity. Envisioning this possibility, well-developed understanding of macrocognition may unlock novel and validated strategies for workforce development, especially in training and recruitment. Measurement of macrocognition would allow practitioners to diagnose human performance and to more fully utilize people as part of cyber defense. For training, it could provide an explanation for why and how skills are missing and suggest interventions. It would provide a layer of understanding of cybersecurity work that is relevant to, but not specifically tied to, an organization, Work Role, or Task.

Understanding macrocognition can also augment recruitment strategies by better answering the question of what qualities help people succeed in cybersecurity careers. It may allow people in other professions and pathways, or people who have not yet picked a profession, to be recruited to participate in cybersecurity careers on the basis of skill or interest in other tasks that involve thinking like a defender. Together with training, this approach may support broader participation in the field.

Conclusion

The methodology presented in this article complements the NICE Framework with research to understand defender cognition. The NICE Framework provides the language of Work Roles, Tasks, and KSAs. The methodology supplies explanation of how KSAs are used and develop. Together, they can help close the cybersecurity skills gap by connecting the elements of work to the capabilities of people.

VI. Acknowledgement

This work is based upon work supported by the National Science Foundation under Grant No. 1553018. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author and do not necessarily reflect the views of the National Science Foundation.

The author gratefully acknowledges assistance with data collection and coding by student members of his lab.

VII. References

Agyepong, E., Cherdantseva, Y., Reinecke, P., & Burnap, P. (2020). Challenges and performance metrics for security operations center analysts: A systematic review. *Journal of Cyber Security Technology*, *4*(3), 125–152. https://doi.org/10.1080/23742917.2019.1698178

Bissell, K., Lasalle, R. M., & Dal Cin, P. (2019) *Ninth Annual Cost of Cybercrime Study.* Dublin, Ireland: Accenture https://www.accenture.com/us-en/insights/security/cost-cybercrime-study. Also available at https://www.accenture.com/_acnmedia/PDF-96/Accenture-2019-Cost-of-Cybercrime-Study-Final.pdf#zoom=50

Ben-Asher, N., & Gonzalez, C. (2015). Effects of cyber security knowledge on attack detection. *Computers in Human Behavior*, *48*, 51–61. https://doi.org/10.1016/j.chb.2015.01.039

Biros, D. P., & Eppich, T. (2001). Human element key to intrusion detection. *Signal.* https://www.afcea.org/content/human-element-key-intrusion-detection

Braisby, N. & Gellatly, A. (2005). *Cognitive psychology*. Oxford University Press.

Buchanan, L., D'Amico, A., & Kirkpatrick, D. (2016). Mixed method approach to identify analytic questions to be visualized for military cyber incident handlers. *2016 IEEE Symposium on Visualization for Cyber Security (VizSec)*, 1–8. https://doi.org/10.1109/VIZSEC.2016.7739578

Champion, M., Jariwala, S., Ward, P., & Cooke, N. J. (2014). Using cognitive task analysis to investigate the contribution of informal education to developing cyber security expertise. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, *58*(1), 310–314. https://doi.org/10.1177/1541931214581064

Cooke, N. J., Champion, M., Rajivan, P., & Jariwala, S. (2013). Cyber situation awareness and teamwork. *ICST Transactions on Security and Safety*, *1*(2). https://doi.org/10.4108/trans.sesa.01-06.2013.e5

Crandall, B., Klein, G. A., & Hoffman, R. R. (2006). *Working Minds: A Practitioner's Guide to Cognitive Task Analysis*. MIT Press.

Crumpler, W., & Lewis, J. A. (2019). *The Cybersecurity Workforce Gap*. Center for Strategic and International Studies. https://www.csis.org/analysis/cybersecurity-workforce-gap

Cuevas, H. M., Fiore, S. M., Bowers, C. A., & Salas, E. (2004). Fostering constructive cognitive and metacognitive activity in computer-based complex task training environments. *Computers in Human Behavior*, *20*(2), 225–241. https://doi.org/10.1016/j.chb.2003.10.016

D'Amico, A., & Whitley, K. (2007, October 9). The Real Work of Computer Network Defense Analysts. *2007 IEEE Symposium on Visualization for Cyber Security (VizSec)*, 19–37. https://doi.org/10.1007/978-3-540-78243-8_2

D'Amico, A., Whitley, K., Tesone, D., O'Brien, B., & Roth, E. (2005). Achieving cyber defense situational awareness: A cognitive task analysis of information assurance analysts. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, *49*(3), 229–233. https://doi.org/10.1177/154193120504900304

Endsley, M. R. (1988). Situation Awareness Global Assessment Technique (SAGAT). *Proceedings of the National Aerospace and Electronics Conference (NAECON), 3*, 789-795. http://doi.org/10.1109/NAECON.1988.195097

Erbacher, R. F., Frincke, D. A., Wong, P. C., Moody, S., & Fink, G. (2010). A multi-phase network situational awareness cognitive task analysis. *Information Visualization*, *9*(3), 204–219. https://doi.org/10.1057/ivs.2010.5

Feltovich, P. J., Prietula, M. J., & Ericsson, K. A. (2018). Studies of expertise from psychological perspectives: Historical foundations and recurrent themes. In K. A. Ericsson, R. R. Hoffman, A. Kozbelt, & A. M. Williams (Eds.), *The Cambridge handbook of expertise and expert performance* (2nd ed., pp. 59–83). https://doi.org/10.1017/9781316480748.006

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

Gutzwiller, R. (2019). *Situation Awareness in Defensive Cyberspace Operations: An Annotated Bibliographic Assessment Through 2015* (No. TR-3184). NIWC Pacific San Diego United States. https://apps.dtic.mil/sti/citations/AD1074248

Gutzwiller, R. S., Hunt, S. M., & Lange, D. S. (2016). A task analysis toward characterizing cyber-cognitive situation awareness (CCSA) in cyber defense analysts. *2016 IEEE International Multi-Disciplinary Conference on Cognitive Methods in Situation Awareness and Decision Support (CogSIMA)*, 14–20. https://doi.org/10.1109/COGSIMA.2016.7497780

Hoffman, R. R. (2019). The Concept of a “Campaign of Experimentation” for Cyber Operations. *The Cyber Defense Review*, *4*(1), 75-84. Retrieved August 31, 2020, from https://www.jstor.org/stable/26623068

Hoffman, R. R., Ward, P., Feltovich, P. J., DiBello, L., Fiore, S. M., & Andrews, D. H. (2014). *Accelerated Expertise: Training for High Proficiency in a Complex World*. Psychology Press.

Klein, G., & Militello, L. (2001). Some guidelines for conducting a cognitive task analysis. In *Advances in Human Performance and Cognitive Engineering Research* (Vol. 1, pp. 163–199). Emerald. https://doi.org/10.1016/S1479-3601(01)01006-2

Klein, G., Ross, K. G., Moon, B. M., Klein, D. E., Hoffman, R. R., & Hollnagel, E. (2003). Macrocognition. *IEEE Intelligent Systems*, *18*(3), 81–85. https://doi.org/10.1109/MIS.2003.1200735

Klein, G., & Wright, C. (2016). Macrocognition: From Theory to Toolbox. *Frontiers in Psychology*, *7*. https://doi.org/10.3389/fpsyg.2016.00054

Mahoney, S., Roth, E., Steinke, K., Pfautz, J., Wu, C., & Farry, M. (2010). A Cognitive Task Analysis for Cyber Situational Awareness. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, *54*(4), 279–283. https://doi.org/10.1177/154193121005400403

Morgan, S (Ed.). (2019). *2019 Cybercrime Report*. Sausalito, CA: Cybersecurity Ventures. https://www.herjavecgroup.com/wp-content/uploads/2018/12/CV-HG-2019-Official-Annual-Cybercrime-Report.pdf

National Institute of Standards and Technology (NIST). (2020). Cyberseek [Web-based heat map of cybersecurity supply and demand]. http://cyberseek.org/heatmap.html

Newhouse, W., Keith, S., Scribner, B., Witte, G. (2017). *National initiative for cybersecurity education (NICE) cybersecurity workforce framework* (Report No. SP 800-181). Gaithersburg, MD: National Institute of Standards and Technology. https://csrc.nist.gov/publications/detail/sp/800-181/final

Nyre-Yu, M., Gutzwiller, R. S., & Caldwell, B. S. (2019). Observing Cyber Security Incident Response: Qualitative Themes from Field Research. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, *63*(1), 437–441. https://doi.org/10.1177/1071181319631016

Paul, C. L., & Whitley, K. (2013). A Taxonomy of Cyber Awareness Questions for the User-Centered Design of Cyber Situation Awareness. In L. Marinos & I. Askoxylakis (Eds.), *Human Aspects of Information Security, Privacy, and Trust* (Vol. 8030, pp. 145–154). Springer Berlin Heidelberg. https://doi.org/10.1007/978-3-642-39345-7_16

Petersen, R., Santos, D., Smith, M., & Witte, G. (2020). *Workforce framework for cybersecurity (NICE Framework;* Report No. SP 800-181 Revision 1). Gaithersburg, MD: National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-181r1-draft

Rasmussen, J., Pejtersen, A. M., & Schmidt, K. (1990). *Taxonomy for cognitive work analysis*. Risø National Laboratory.

Rooney, V. M., & Foley, S. N. (2018). What You Can Change and What You Can't: Human Experience in Computer Network Defenses. In N. Gruschka (Ed.), *Secure IT Systems* (Vol. 11252, pp. 219–235). Springer International Publishing. https://doi.org/10.1007/978-3-030-03638-6_14

Rouse, W. B., & Morris, N. M. (1986). On looking into the black box: Prospects and limits in the search for mental models. *Psychological Bulletin, 100*(3), 349–363. https://doi.org/10.1037/0033-2909.100.3.349

Saner, L. D., Campbell, S., Bradley, P., Michael, E., Pandza, N., & Bunting, M. (2016). Assessing Aptitude and Talent for Cyber Operations. In D. Nicholson (Ed.), *Advances in Human Factors in Cybersecurity* (Vol. 501, pp. 431–437). Springer International Publishing. https://doi.org/10.1007/978-3-319-41932-9_35

Schraagen, J. M., Chipman, S. F., & Shalin, V. L. (2000). *Cognitive task analysis*. Psychology Press.

Schuster, D., & Wu, S. (2018). Toward Cyber Workforce Development: An Exploratory Survey of Information Security Professionals. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, *62*(1), 1242–1246. https://doi.org/10.1177/1541931218621285

Shah, A., Ganesan, R., Jajodia, S., & Cam, H. (2018). A methodology to measure and monitor level of operational effectiveness of a CSOC. *International Journal of Information Security*, *17*(2), 121–134. https://doi.org/10.1007/s10207-017-0365-1

Six skills you need to succeed in cybersecurity. (n.d.). In *Dice Insights*. Retrieved August 31, 2020, from https://insights.dice.com/cybersecurity-skills/

Stein, D., Scribner, B., Kyle, N., Newhouse, W., Williams, C., & Yakin, B. (2017). *National Initiative for Cybersecurity Education (NICE) Framework Work Role Capability Indicators: Indicators for Performing Work Roles* (Report No. Draft NISTIR 8193). Gaithersburg, MD: National Institute of Standards and Technology. Retrieved from https://csrc.nist.gov/CSRC/media/Publications/nistir/8193/draft/documents/nistir8193-draft.pdf

Tetrick, L. E., Zaccaro, S. J., Dalal, R. S., Steinke, J. A., Repchick, K. M., Hargrove, A. K., ... Wang, V. (2016). *Improving Social Maturity of Cybersecurity Incident Response Teams*. Fairfax, VA: George Mason University. https://calctraining2015.weebly.com/the-handbook.html

Trent, S., Hoffman, R. R., Merritt, D., & Smith, S. (2019). Modelling the Cognitive Work of Cyber Protection Teams. *The Cyber Defense Review*, *4*(1), 125-136. https://www.jstor.org/stable/26623071

Wei, J., & Salvendy, G. (2004). The cognitive task analysis methods for job and task design: Review and reappraisal. *Behaviour & Information Technology*, *23*(4), 273–299. https://doi.org/10.1080/01449290410001673036

Woods, D. D., & Roth, E. M. (1988). Cognitive Engineering: Human Problem Solving with Tools. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, *30*(4), 415–430. https://doi.org/10.1177/001872088803000404

Yurcik, W., Barlow, J., & Rosendale, J. (2003). Maintaining perspective on who is the enemy in the security systems administration of computer networks. *Proceedings of the ACM CHI Workshop on System Administrators Are Users (*pp. 345-347). ACM Press.

Zhong, C., Yen, J., Liu, P., Erbacher, R., Etoty, R., & Garneau, C. (2015). An integrated computer-aided cognitive task analysis method for tracing cyber-attack analysis processes. *Proceedings of the 2015 Symposium and Bootcamp on the Science of Security*. https://doi.org/10.1145/2746194.2746203

Figure 1

Research Methodology

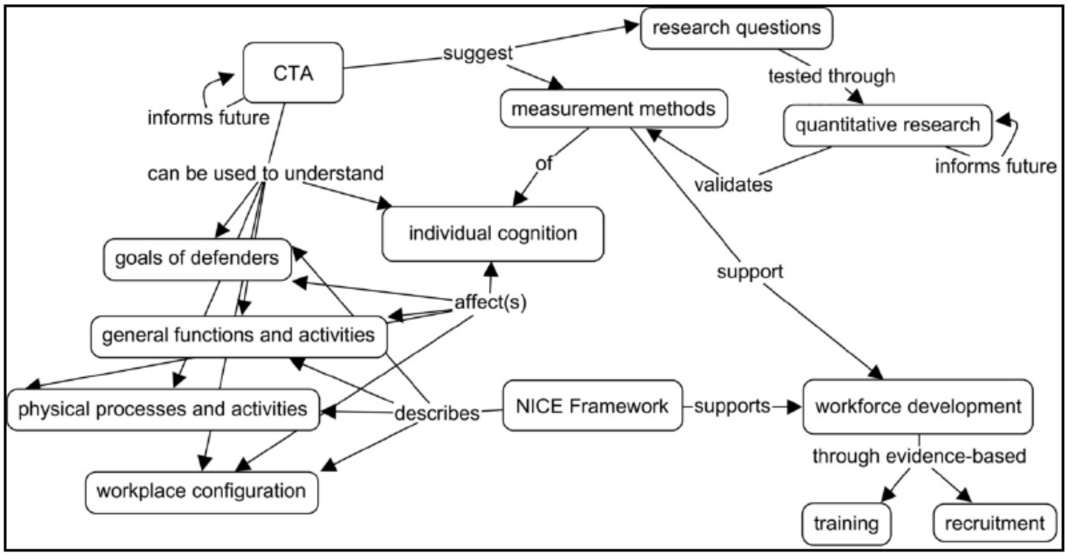

Figure 2

Concept Map for Participant 1

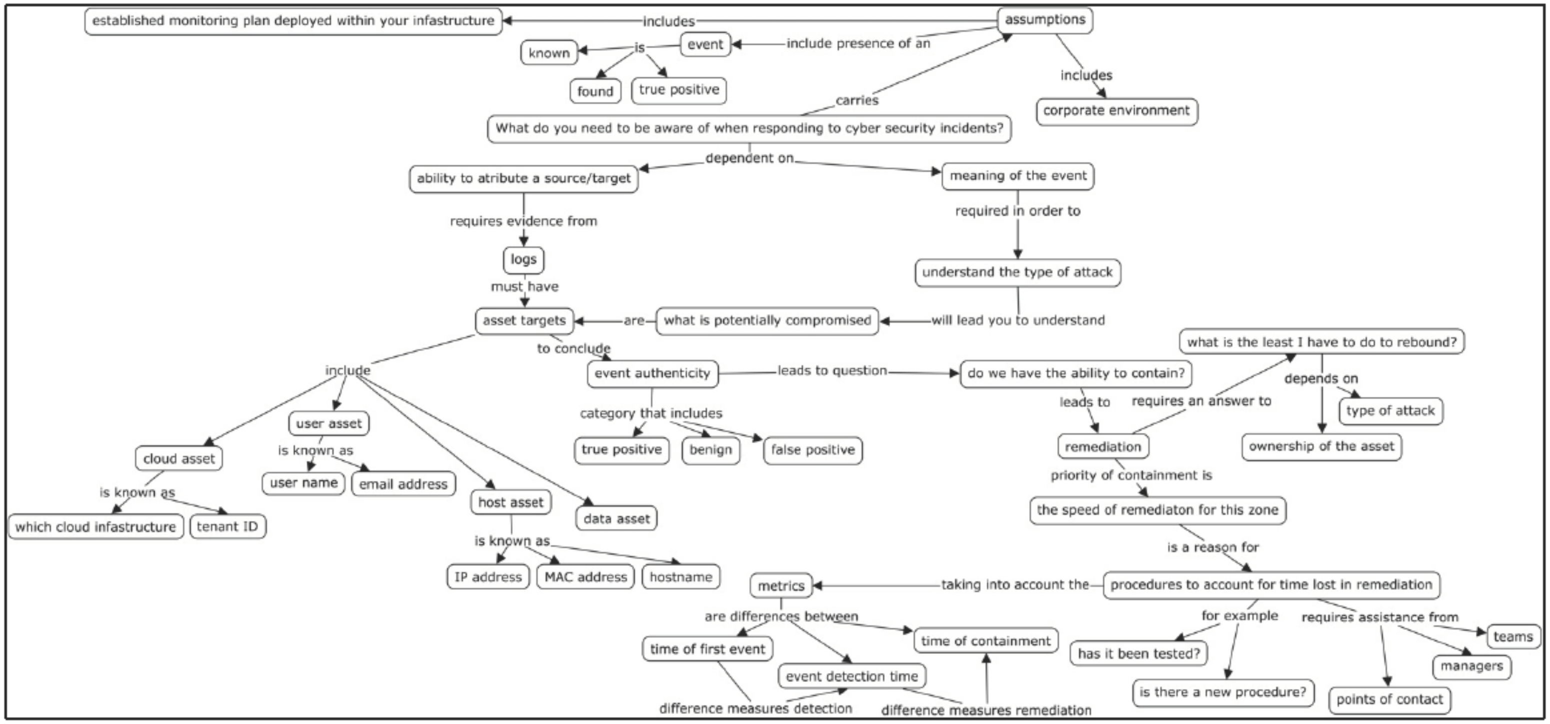

Figure 3

Concept Map for Participant 2

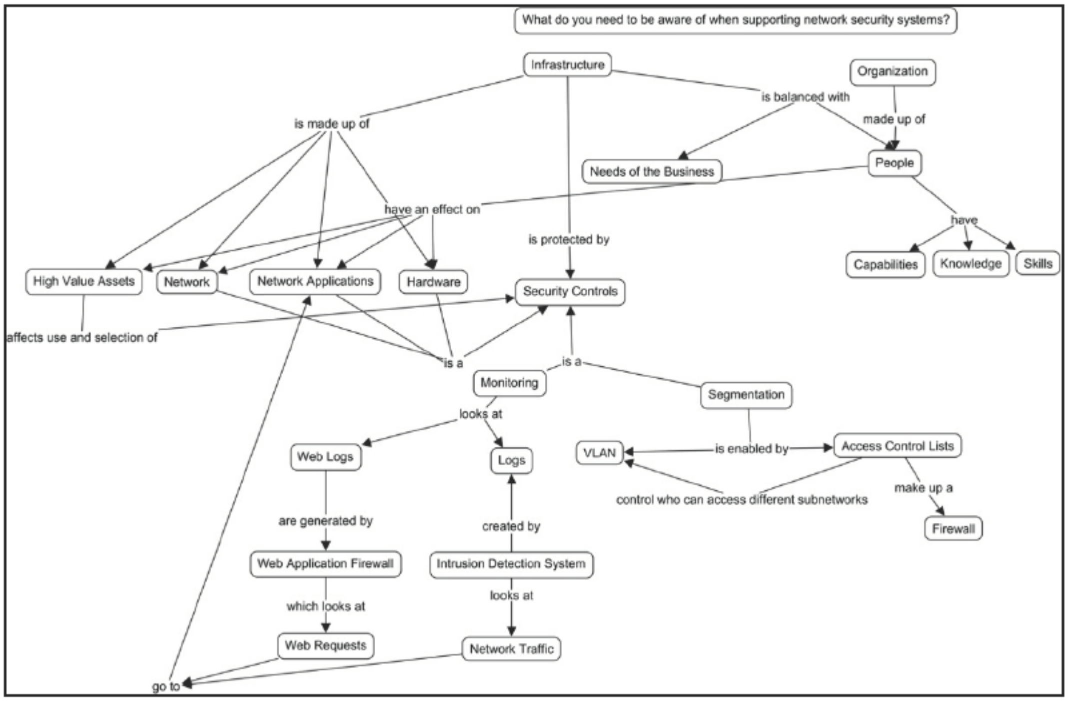
What do you need to be aware of when supporting network security systems?
Infrastructure
Organization
is balanced with
made up of
is made up of
Needs of the Business
People
have an effect on
have
is protected by
Capabilities
Knowledge
Skills
High Value Assets
Network
Network Applications
Hardware
Security Controls
affects use and selection of
is a
is a
Monitoring
Segmentation
looks at
Web Logs
Logs
VLAN
is enabled by
Access Control Lists
control who can access different subnetworks
make up a
Firewall
are generated by
created by
Web Application Firewall
Intrusion Detection System
which looks at
looks at
Web Requests
Network Traffic
go to