TECHNICAL REPORT 3184
June 2019

Situation Awareness in Defensive Cyberspace Operations: An Annotated Bibliographic Assessment Through 2015

Robert Gutzwiller, PhD

NIWC Pacific
San Diego, CA 92152-5001

This page intentionally blank.

TECHNICAL REPORT 3184
June 2019

Situation Awareness in Defensive Cyberspace Operations: An Annotated Bibliographic Assessment Through 2015

Robert Gutzwiller, PhD

Administrative Notes:

This Report was approved through the Release of Scientific and Technical Information (RSTI) process September 2018 and formally published in the Defense Technical Information Center (DTIC) in June 2019.

This Report's content represents work performed under Space and Naval Warfare Systems Center Pacific (SSC Pacific). SSC Pacific formally changed its name to Naval Information Warfare Center Pacific (NIWC Pacific) in January 2019.

NIWC Pacific
San Diego, CA 92152-5001

NIWC Pacific
San Diego, California 92152-5001

M. K. Yokoyama, CAPT, USN
Commanding Officer

W. R. Bonwit
Executive Director

ADMINISTRATIVE INFORMATION

The work described in this report was performed by the User-Centered Design and Engineering Branch of the Command and Control Technology and Experimentation Division of the Naval Information Warfare Center Pacific (NIWC Pacific), San Diego, CA. The NIWC Pacific Naval Innovative Science and Engineering (NISE) Program provided funding for this Basic Applied Research project.

Released by
Matthew A. Yanagi, Head
User-Centered Design and Engineering
Branch

Under authority of
Frank P. Calantropio, Head
Command and Control Technology
and Experimentation Division

ACKNOWLEDGMENTS

This work was funded in part by the Office of Naval Research Naval Innovative Science and Engineering (NISE) funding through Section 219, to Dr. Gutzwiller under Characterizing Human Limitations and Impediments to Cyber Situation Awareness, and by the Office of the Assistant Secretary of Defense funded project to SPAWAR Pacific.

EXECUTIVE SUMMARY

Situation awareness (SA) is a buzzword concept, but unlike most buzzwords, it has a robust and scientific research focus. Grounded in cognitive psychology and human factors research findings, SA is essentially characterization of what a person knows about their current and future environment, in the context of their current goals. Being aware of critical information, comprehending it, and even projecting the situation into the near future is a highly useful skill in many different domains, such as aviation, driving, and healthcare.

Cyberspace operators share many of the same concerns as more traditional roles of the pilot and the driver. In particular, cyber network defense (CND) operators must remain aware of different types of activity, comprehend what various sources of information mean and how they are pieced together, and project what effects on their network will be in order to stop or prevent threats.

For future Warfighter performance in cyberspace, substantial benefits are derived from improving defender SA, as we have in other domains. Improving SA is permanently connected with understanding how to measure awareness in the cyber environment. In support of this concept, a review of the available literature in cyber SA was conducted to determine how to begin the process of measurement and improvement, and to derive key points for further research.

This report finds that the utility of SA analysis and measurement has yet to be realized in cyberspace. With a few exceptions, almost no experimental work was found on measuring or characterizing the process or product in developing cyber situational awareness. In those few found, there are still occasional methodological or analytical deficits and concerns, which preclude any strong conclusions. The development of specific measurement techniques must be explored elsewhere.

Almost no research was found measuring SA in the CND environment, though many reports claimed that a new or unique interface *could improve it.* The willingness and rapacity of these claims in the literature mirrors claims by industrial software solutions, which many cyber professionals abhor and abandon in favor of Excel spreadsheets and command line interfaces. Together this suggests that demand is present for SA improvement, but capability has not risen to match it.

This page is intentionally blank.

CONTENTS

Figures

1. WHAT IS CYBER SITUATION AWARENESS, AND WHY DO WE CARE ABOUT IT?

Situation awareness (SA) is a concept most are familiar with by name or by experience. For example, drivers are acutely aware of their knowledge of the traffic and road environment around them. In order to drive safely, for example, one engages in an attentive process to select critical areas for survey (e.g., checking blind spots, visually scanning the road around the vehicle), integrating this into a "gist" of the environment, and even predicting the near future (e.g., can I make that gap between cars to fit in when changing lanes?). These general processes comprise situation awareness, and are important in many other domains, including aviation, air traffic control, medicine, manufacturing, supervisory control, and unmanned systems operation.

Anecdotally, SA may be conceptualized as "paying attention", or knowing everything important to know. Both are unsatisfactory explanations (Dekker, 2012). If an operator lacks awareness, it suggests the operator missed something critical. However, therein lies several fallacies: that operators are omnipotent (or that being so is necessary for safe performance); and that information is always valuable at all times. Just because some aspect of the environment was missed does not mean the operator was derelict of duty, but instead indicates changes in goals or information process. Hazard avoidance in driving, for example, would appear to a novice to require that the driver is fully aware of *all* surrounding cars. One can effectively argue, however, that only awareness of cars or objects *that could be hit* is important (and this is further supported in research; Gugerty, 1997). Clearly, **awareness is not just paying attention**: it is paying attention to the limited set of things at the right times, while building a comprehensive, predictive perspective on them as a whole. The building and projecting aspects are also why, depending on level of experience, SA as a process is mentally demanding, using resources such as working memory (Gutzwiller and Clegg, 2013; Sohn and Doane, 2004).

The general assumption is that **better SA leads to better decisions** because information and understanding are improved ahead of a decision opportunity (also upheld in research; Endsley, 1995b; Klein, 1997). One can imagine the driver again: if the driver fails to look in their blind spot (e.g., perception), or does not appropriately time their maneuver (e.g., projection), their merge (the decision point) into a busy lane on a highway could be disastrous.

1.1 SITUATION AWARENESS THEORY AND MEASUREMENT

The **usefulness of understanding individual's situation awareness** has run high, for example in determination of accident root cause (Jones and Endsley, 1996), understanding of complex dynamic environments such as air traffic control (Durso et al., 1998; Durso, Bleckley, and Dattel; 2006; Gronlund et al., 1997; Mogford, 1997); driving (Endsley and Kiris, 1995; Gugerty, 1997; Ma and Kaber, 2005), aviation (Sohn and Doane, 2004; Sulistyawati, Wickens, and Chui, 2009; 2011), medicine (Gaba and Howard, 1995), and in the design of displays (Endsley, 1995b).

A variety of theories conceptualize and measure situation awareness in the human operator differently. In the summaries below, I primarily present them as theoretical guidelines for new readers perhaps unfamiliar with situation awareness as a cognitive theory. They are by no means complete, or the only theory – however they are usually those found "in the wild" and thus easily discussed across a crowded table. Rising perspectives, which are perhaps more complex (such as distributed situation awareness or DSA; (Stanton, Salmon, and Walker, 2015; Stanton, 2016; Stanton et al., 2006)) may be more suitable in the long run. But DSA appears to be a skillset that HF/E at large has not yet adopted, for better or worse. I personally believe there to be a wealth of insight in

the methods as applied to cybersecurity but recognize the difficulty in conducting the research itself as well as analysis of the results. DSA may require more access to large amounts of people and information, a constraint that makes cyber assessments extra difficult. (It may yet be worth it; however, this is starting to go outside the goals of the review here).

1.2 TWO MAJOR THEORIES TO EDUCATE THE READER, AND FRAME THE FOLLOWING REVIEW.

1.2.1 Theory 1 - Endsley's Situation Awareness

Mica Endsley's situation awareness theory (1995a) is certainly the most widely known and supported perspective on situation awareness. She provides a formal definition as "the perception of the elements in the environment within a volume of time and space, the comprehension of their meaning, and the projection of their status in the near future" (36; Endsley, 1995a).

The definition comprises 3 established levels; level one is concerned with basic perceptual and attention process in situation assessment — what do people need to attend to in the environment? Level one then encompasses our look out the window or over our shoulder into the blind spot while driving. The second level, building on the first, refers to comprehension processes — how does information inform the operator about the global status or picture of the environment, and build understanding? Therefore, it is akin to building a "gist" understanding of the road environment. The final third level builds on levels 1 and 2, and is a process attempting to predict the near future based on expectancies and relationships in the mental model. For driving, this may predict where cars and road obstacles will be and how speed and trajectory move them in the near future.

Within the Endsley SA theory, memory is a large component. In particular, long term memory feeds the development of a mental model, a representation of the environment that includes all the parts and how they interact. With experience , these models become more accurate, though assessing them can be difficult. The mental model aids in directing attention to process relevant information, i.e., guides the creation of SA. This "process" of SA, or "how situation awareness is built" is different and usually somewhat unreportable, in comparison to the product of SA (explicit information or concepts at a specific time or 'slice'). Theoretically this means that a different process could produce the same product. Therefore, one may need to measure both, in order to better understand where improvements can be made. However, detailed task analyses can only go so far, and my emphasis here is on reviewing experimental work that attempts to measure and improve cyber SA specifically. CTAs should be used to identify important information that is candidate for SA assessment.

1.2.2 Theory 2: Smith and Hancock's Adaptive Consciousness

Another conceptualization of SA provides a different perspective than that of Endsley. Smith and Hancock (1995) critiqued Endsley's early theories, and proposed an embedded-interactive model. They define situation awareness formally as "adaptive, externally directed consciousness" (138; Smith and Hancock, 1995). Consciousness in this case is "that part of an agent's knowledge-generating behavior that is within the scope of intentional manipulation" (138). SA is abstracted as the invariant at the center of a modified *perceptual-action cycle* as proposed by Neisser (1976), which is theorized to organize the information available in an environment, combined with the *knowledge needed to make relevant assessments* and the *knowledge of actions* that will allow the operator to achieve goals.

In the Smith and Hancock model, SA is not the "picture" of the situation, but rather, **SA builds the picture** (Smith and Hancock, 1995), reflecting a process approach. The interactive role of the perceptual-action cycle *requires* the operator to be pursuing a goal (specifically defined, again, by the arbiter of performance – an *explicit* goal) and **interacting** with the environment as it changes. Consequently, simply "attending" is not SA per se; one must be looking *for* something – in other words, attention under goal pursuit.

It should be noted that the Smith and Hancock conceptualization does not have an associated measure, given the emphasis on process. On the positive side, it does appear to resolve the problems recognized by Dekker regarding Endsley's theory, by discarding reliance on "snapshots" of memory and refusing to measure the "mind-matter" mismatch which is also not SA. They also formalize the need for an established arbiter of performance, which makes specific information important to a goal and allows SA to be evaluated.

While it is possible that Smith and Hancock's conceptualization is providing more construct validity, the need for formal measurement from it is difficult to achieve, especially in comparison to other methods available, including those of subjective, objective and distributed SA varieties. These are discussed in Section 1.4.

1.3 MEASUREMENT

Since I have defined 2 perspectives of what situation awareness is, I move on to different measures of SA. These measures comprise 3 major categories; **subjective** — what does the operator say, or think they know regarding information, comprehension and projection?, **objective** — what does the operator actually know of the moment, as assessed against ground truth and/or by external observations), and **implicit**, a form of mixed measurement in which the researchers determine what recordable behavior was exhibited, and what elements of it awareness can be inferred. These measures typically assess the *product* of SA, although some incorporate process measures as well. The key to creating valid versions of all of these measures is in determining what information is necessary for awareness, as discussed in the Smith and Hancock theory, and making sure the measures assess it. The operator may feel very "aware" but may actually be missing key information; or, the researcher can inappropriately design queries to ask about unimportant information.

1.3.1 Subjective Measures.

Subjective assessment of awareness has some validity, and many methods of reporting are available. Subjective responses in general provide valuable insight, as they are sometimes the first step into the domain by researchers.

The main downside of subjective measures of awareness is that of subjective reporting in general, perhaps exacerbated by the dynamics of situations where awareness is important; participant memory

is a faulty and reconstructive process. If a subjective assessment of SA is solicited after the awareness experience — a typical practice — it is difficult to determine how reflective a report is for an event that happened potentially minutes or hours previously. One cannot know for certain then exactly what a subjective SA measure captures and when during performance these subjective impressions were created (Jones, 2000). It is also possible that participants subjectively rating SA conflate awareness with performance; in other words, participants who did "well" on a task may end up rating their awareness as high regardless of whether it actually was (Jones, 2000). After all, **SA is not performance**, even though it should be related. **SA is also not workload** (Endsley, 1993), though it may be related. Separating out subjective impressions can be difficult and has generally led to the use of more consistent, objective measurement.

1.3.1.1 Objective Measures.

Contrasted to subjective measures, objective measures of SA do not rely on subjects to rate their own awareness. Instead objective methods typically focus on asking operators to recall or respond to experimenter-determined questions. These questions are designed (sometimes with expert help) to assess the operators' awareness at each stage of the Endsley model (e.g., perception, comprehension, and projection), though the literature has arguably measured perception more than comprehension or projection.

I discuss 2 techniques (SAGAT, and SPAM) commonly used to assess objective awareness, along with their relative strengths and weaknesses. Both SAGAT and SPAM methods provide an opportunity for the operator to be evaluated on their state of SA at each of the 3 main levels of awareness. Again, the point of discussing measurement is so that when approaching measuring SA in the cyber environment, it can begin from a grounded, theoretically optimal approach. There is probably no need to reinvent the SA wheel.

SAGAT – SITUATION AWARENESS GLOBAL ASSESSMENT TECHNIQUE.

Generally speaking, Situation Awareness Global Assessment Technique (SAGAT) measures situation awareness through assessing memory for information and understanding of the situation at the time of assessment; typically SAGAT is a series of pre-formed questions chosen to be asked at various intervals (Endsley, 1995b). They typically interrupt an ongoing task and remove all information from the screen or environment to avoid biased responding. SAGAT therefore rests on the assumption operators retain, and have available to recall, the requested information available in memory (Endsley provides evidence that this information can be retained up to 6 minutes, but this was in an aviation context), and second that the questions are not disruptive to the task (again Endsley provided evidence that they were not); and that the accuracy of responses is diagnostic of operator awareness (Endsley, 2000).

Because the time available to respond to SAGAT questions is typically controlled (e.g., 2 minutes provided to answer a series of questions), participants who answer faster or slower receive no explicit benefit in the analysis of responses. Yet one could suppose speed of response may be indicative of the process of awareness, especially as it relates to information retrieval (faster times indicating better awareness). Such an assumption is examined by using the Situation Present Assessment Measure (or SPAM) described next.

SPAM – SITUATION PRESENT ASSESSMENT MEASURE

Durso and colleagues (Durso, Rawson, and Girotto, 2007; Durso et al., 1998) posited that SA is probably best labeled situation *comprehension*, in the vein of reading comprehension. According to

Durso, framing as comprehension moves the field forward away from memory-based thinking, and toward a more dynamic assessment. His theory is based in pattern recognition, as found in chess (Ericsson and Kintsch, 1995), suggesting that "gist" – memory for meaning – is likely more important on average than memory for any one particular element (Durso Rawson, and Girotto, 2007). Memory for meaning and memory for specific elements can still co-occur, but which is most important may vary.

Using the analogy of reading comprehension then, SPAM divides SA measurement into assessments of **accuracy** (e.g., memory for meaning) and **latency** (measures during ongoing processing of information that serve as a "predictive" component of comprehension). Durso and colleagues measure SA using real-time probes – brief questions that pop-up during an operators' performance. Operators can negotiate when to answer them but must answer them within a certain time period (usually brief). In contrast to the SAGAT method, which removes the environment during SA query, in SPAM the environment is available. Operators, having read the query, could recall information; or they could use memory for *where information is located* to rapidly answer the question at lower levels of SA. In the theory viewpoint of Durso, either one counts as awareness and memory should be a faster process than searching for the information directly. For higher levels of SA (e.g., comprehension), the knowledge required to answer a query is less commonly or easily available in the environment.

Durso claims these differences make SPAM SA assessment more related to actual, current awareness (as it has no chance to decay or be forgotten in transition to queries as in SAGAT and post-test measures). This assumption also considers context. Operators may not have to remember information for them to become aware of it[1] (as not all information is relevant to remember), a distinction that moves theory toward a more embodied idea of awareness. Essentially, the SPAM measurement incorporates how operators benefit from knowing where to find information in the world, versus information solely "in the head" (a tactic that is reflected in other work; O'Regan, 1992). Nevertheless, these measures are still theoretically insufficient to fully address SA as it relates to consciousness (Dekker, 2012), and are still providing us only slightly improved "snapshots". A clear solution to this problem has not been proposed. Others have defended existing measures (Parasuraman, Sheridan, and Wickens, 2008) and it can be suggested that short of new measurement methods intended to capture larger systems, such as distributed situation awareness (DSA), the methods are closer than simply surveys and more powerful than subjective reports.

SPAM measures rely in part on avoiding operator interruption, as in SAGAT. Criticism has been levied at SPAM methodology, in that it requires additional operator workload or otherwise alters operator behavior during ongoing performance. A recent experiment (Pierce, 2012) critiqued SPAM, finding a small amount of interference in an air traffic control task, as the cuing paradigm for SPAM measurement in which the operator receives a notification to give a response, was not enough to keep the measure from interfering with primary task performance.

[1]Other experiments show memory-dependent answers may be difficult to dissociate between current awareness, and prior knowledge. In other words, some awareness measures may capture knowledge, not awareness. In ATC performance, for example, prior knowledge such as knowing all Southwest planes are 737s, may facilitate seemingly greater awareness even when the operator may not have attended to the plane type information (Gronlund et al., 1997).

1.3.1.2 Implicit Measures

Implicit measures of SA are one additional form of objective assessment. The data gleaned from them is impartial to subjective interpretation; however, as their name suggest they do not rely on direction questions or solicitations for information. Instead, implicit measures use indirect performance assessments, which are inconspicuous to the participant. These effectively measure awareness, as shown in several studies (Brickman et al., 1999; Drury, Scholtz, and Yanco, 2003; Gugerty, 1997; Gutzwiller and Clegg, 2013; Pritchett and Hansman, 2000; Yanco and Drury, 2004).

Examples of an implicit measure is measuring whether or not a user responds to the emergence of a set of information (thus indicating awareness of it), or how they react to a specific onset of an element (e.g., a runway incursion). The key strength is that these measures do not require the operator focus on anything other than the task at hand, nor do they rely on memory. Additionally, the results are sometimes well correlated with the direct questioning of objective measures (Gugerty, 1997). It may be that implicit measures more adequately represent the notion of situation awareness. Implicit measures are less likely "snapshots" and incorporate more of a "comprehension" perspective by examining behavior in context.

The difficulty with implicit measures is in determining what aspects of performance truly provide this diagnosis (though these issues are similar for the other measures, only differing in their explicit nature). These aspects may be determinable with the use of a SME. An additional weakness is that implicit measures tend to be "all or none", i.e., participants either "were" or "were not" aware. This deviates from the comprehension perspective of Durso and may restrict the level of understanding gleaned in analysis (Gutzwiller and Clegg, 2013).

I observe here that it may be the co-occurrence — proper perception and memory for an element *in context* that is most important. Without an emphasis on the gist, then, one may lose the important surroundings that contextualize a piece of information. The accidental grounding of the Royal Majesty contains this type of SA failure in that information was viewed that indicated a problem, but was dismissed because of contextual factors and assumptions (Lützhöft and Dekker, 2002). It suggests that measurement of SA should include several forms of objective assessment, perhaps as many as the domain can support for comparison and the differential value each provides. As I mentioned at the start, it also suggests the importance of more distributed measurement strategies, such as those suggested by Distributed SA (Stanton, 2016).

Regarding cyber defenders, the suspicion is that cyber defenders make similar kinds of contextualized judgments as intelligence analysts, and may be biased - worse, incentivized - toward dismissing potential attacks. Though defenders may construct awareness, any incoming information could be contextualized by the base rate of attacks they receive, and reinforced by their desire to confirm findings (both initially and in hindsight), leading them toward ignoring signs of malicious activity because it does not fit with their mental models (Lemay and Leblanc, 2018). Measuring awareness in this sense, then, may be a way to assess bias in defenders if it can be confirmed that information was attended to but dismissed, as in the Royal Majesty.

1.3.2 Summary

Based on the above outline of broad theory and measurement, it is clear that a variety of methods and viewpoints on SA exist. Choosing the appropriate ones for a given domain remains important. Different measures should be chosen based on their likelihood to disrupt the operator, their likelihood to capture behavior which has severe consequences if it is not noticed, and other methodological constraints such as time and administration. The best solutions may combine several different measures to gather a more fleshed out perspective of SA, though for similar (administrative) reasons this is rarely done.

2. UNDERSTANDING CYBERSPACE SITUATION AWARENESS

2.1 CYBERSPACE AS A DOMAIN OF OPERATION.

Defining the operation environment is a difficult process. The notion of cyberspace itself is vague. For my purposes, I take cyberspace to comprise communications between computing devices (networks), the devices themselves, and the interconnections of machines and networks to the physical world as both sensors and actuators.

Conti et al. provides a concise representation of cyberspace as the parallel dimension to the electromagnetic spectrum of reality. In their view, it is a dimension that permeates physical domains, but remains separate from traditional reality. I adopt Conti's perspective here. Cyber encompasses all of the traditional boundaries of operations (space, air, land and sea) in the physical world, and yet, it is distinct (Conti, Nelson, and Raymond, 2013). For example:

Cyberspace *activity* happens much faster - literally at the speed of light; or much slower, over years at times. Information in driving, or anesthesiology, changes observably over both short and long scales of time and situation awareness happens in shorter periods of performance, which are familiar to humans. Some cyber attacks are on the speed of light level, while strategies of some long-term persistent threat attackers in cyberspace, instead exert their behavior periodically over a very long period of time (weeks, months, and even years). Humans are not suited to act in real time on the fast side of the scale (< 10ms) and may easily miss information at the other (> days and weeks). Thus a significant amount of computer mediation, via software tools and sensors, is required to manage the collection and interpretation of current and prior network activity through recorded data streams. The reliance on these abstractions lends itself to weakness (e.g., Conti, Ahamad, and Stasko, 2005).

In the cyber world the intentions are much subtler and occur within a massive amount of "noise". Whereas a physical attacker may be hidden but strategically present themselves, probing the line for weakness (as in the conflict in the Ia Drang Valley, early in the Vietnam War), as a cyber defender, it is much harder to attempt to understand the intent of attacker actions (Li, Ou, and Rajagopalan, 2009). In this respect, cyber defense is more akin to recent armed conflicts in Iraq and Afghanistan, than to large-scale operational engagements.

Threats faced are very asymmetric. Yurcik highlights that the attackers have the advantage at nearly every turn; there is constant network exposure with internet connectivity and the rise of the internet of things. Attackers only need to find one vulnerability while defenders must find all of them (Yurcik, Barlow, and Rosendale, 2003). Defenders are further reliant on peer security, but attackers can work alone; defense must be multifaceted while attackers can be specific, and most importantly attackers will normally have the element of surprise. These same phenomena are generally true for the real world, though security in cyberspace is perhaps a worse-case example. (And, humans have much less experience and so fewer heuristic mechanisms in place for dealing with these strategies.)

Operators' perception of information is highly contrived in the cyber environment. Whereas real-world perception served the pilot in the cockpit (i.e., the capability to look out the window), the same representation is unavailable to the cyber operator. They must instead rely on an interface to communicate all relevant information. While we as humans are innately familiar with the real world, we are largely unfamiliar with how to navigate and monitor the cyber realm, and therefore, operators must accumulate a wealth of experience in order to correctly understand the environment and make decisions (M. Champion, Jariwala, Ward, and Cooke, 2014; Dutt, Ahn, and Gonzalez, 2013).

As there are differences, so are there some facets that are similar. Below is a brief list of areas regarding cyber defense and situation awareness that are comparable to the real physical world.

All levels of SA are likely to be important to both physical and cyberspace operations. For example, detection of a possible exploit or vulnerability falls in each of the levels of situation awareness. The operator must **notice** the exploit is present, **comprehend** what it means for his or her network, and **make the prediction** that a malicious actor may take advantage. The likely decision is to patch the vulnerability. While vulnerability scanning and patching computers is a largely administrative task, it also shares the need for SA in management and deployment. Other analysts may work at near real-time to examine network activity alongside automated classification systems and other tools; these roles too are predicated on building an accurate, updated understanding of the cyber environment in order to appropriately triage potential cyber incidents.

Kott, Buchler, and Schaefer (2014), highlighted that **a significant portion of time in both domains will be spent gathering information** (level 1 SA) — and that this information is fairly uncertain. The uncertainty then increases the role of experience and operator intuition in making the determination of what information is critical. True for both cyberspace and real space.

Another potential parallel is that, as in the physical world, operators may over-rely on static knowledge versus dynamic (Kott, Buchler, and Schaefer; 2014), in large part because dynamic understandings are very hard to gather. The difficulty with dynamic information ties back to similar cognitive load constraints in the real world in the results of Blalock et al. (2014) and general cognitive limitations. This distinction may play out in cyber network defense because dynamic understandings are also very difficult to mentally piece together, due to the variation in available data, and the time span over which attacks occurs (plus the attackers' purposeful obfuscation of evidence of their presence). In fact this issue has already been identified as the difficulty in "connecting the dots" in cyber incidents (Pfleeger and Caputo, 2012).

The **use of teams** is also common between physical and cyber domains, though the methods may initially differ between the two (Kott, Buchler, and Schaefer; 2014). Teaming has had interesting effects on situation awareness in the physical realm; namely, commander SA could degrade based on other team members reinforcing poor information or confirmation biases, or simply being a distraction. It stands to reason, however, that negative influences do not characterize all aspects of team performance. Here is where methodologies for awareness assessment that seek a more distributed, systems approach (as they have elsewhere in submarine command and control) may be very useful or even required in certain cases.

In summary, then, cyberspace activity happens much faster or much slower; the attacker may be hidden, and threats faced by operators are very asymmetric, made difficult by the contrived cyber environment. As a result operators rely on experience as they do in the physical world: noticing, comprehending and predicting activity in both areas, with the largest emphasis likely to be on gathering information. Another commonality is over-reliance on stale or static information, and the ways this may affect decision making, especially in teams. The importance for improving SA in the cyber domain, then, reflects its importance in the physical. Decisions are still likely to be based upon an operators current understanding of the environment (and their predictions about what might happen next; i.e., situation awareness). **This basic assumption is also supported in many of the cognitive task analyses performed to date on various roles of cyber analysts** (Alberts et al., 2004; D'Amico et al., 2005; Erbacher, et al., 2010; Killcrece, Kossakowski, Ruefle, and Zajicek, 2003; Mahoney et al., 2010).

2.2 Defining Cyberspace Situation Awareness

Having defined the area of interest as network defense and described how SA is viable within it, I turn to the more specific definitions for cyber situational awareness that have been proposed in the existing literature. These definitions are few, as most authors choose to cite or adapt Endsley's definition, but they are worth mentioning. (Conti, Nelson, and Raymond, 2013) provide a definition of cyber SA in a military setting as "the requisite current and predictive knowledge of the environment upon which operations depend- including physical, virtual and human domains — as well as all factors, activities, and events of friendly and adversary forces across the spectrum of conflict." One can imagine the assessment here would be monstrous and require the use of distributed SA measurement techniques. Yurcik et al., defined SA for a system administrator as "the ability to effectively determine an overall computer network status based on relationships between security events in multiple dimensions." (Yurcik, Barlow, and Rosendale, 2003). (Levin, Tenney, and Henri, 2001) has a more vague definition of situation awareness as it relates to cyber awareness displays; in their view, SA is "the perception of network events and data, the comprehension of their meaning in terms of mission, resources, connectivity, threats, and vulnerabilities, and the projection of their status in the near future" (Levin, Tenney, and Henri, 2001) which is very similar to Endsley but with cyber terminology and objectives.

Barford et al. (2010) highlights 7 facets of needed awareness, that comprise their definition of situation awareness for cyberspace. These are an assessment of the current network situation; understanding the impact of an attack and tracking situations evolving; understanding the behavior of the attacker, and the why / how the situation is caused; understanding the trustworthiness of any available awareness, including how current the awareness is; and using knowledge of defense and offense positions to examine potential future situations. Barford et al. (2010). therefore provides a more nuanced definition which broadly highlights the challenges for situation awareness in network defense and appears to address all 3 levels of awareness. However, the definition does take a "post-attack" perspective in many ways, assuming an attack has occurred. Awareness has no such dependency; good awareness could simply mean that you are aware you are not being attacked; poor awareness could mean that you think you are constantly under attack, when actually it is routine non-malicious activity. In fact, this is the struggle of IT security — keeping things secure but usable for customers.

None of the above definitions appears to go beyond what has been previously developed for awareness in the human factors literature. Therefore, I felt confident in claiming that a review could be done using common terminology from human factors.

This page is intentionally blank.

3. A FORMAL REVIEW OF THE CYBERSPACE SA LITERATURE

In my initial examination of the research literature, a formal review of cyber SA was found (Franke and Brynielsson, 2014). Franke and Brynielsson (2014) conducted a literature review of the cyber situational awareness domain and were able to categorize their findings into several main clusters. However, for the present purpose, the biggest finding from this review is twofold. First, their review focused on the literature at large and human interactions only encompasses a small part of their efforts. This would suggest that the review may be incomplete, and indeed I have found additional papers not covered by the earlier review. Second, within their review they only found 3 papers which experimentally examined the operator's cyber situational awareness. Many of these were still only observational studies, rather than experiments, and were often conducted within ongoing exercises. The current review continues to point to a lack of overall assessment of cyber situational awareness.

The discovery of the prior review did not invalidate the need for another review for 3 reasons. First, as cyber (especially human factors involvement in cyber) is a developing area, a review conducted in 2013 necessarily misses the most up to date information. Second, the focus of the prior review was top-level; the authors did not focus in depth on measures of SA, levels of awareness assessed, or truly in any depth on the experiments that they did find (which were few). Third, issues surrounding the human limitations in the cyber context would not be found by searching only for cyber situational awareness.

3.1 METHODOLOGY

I examined the available, published literature (up to March 2015 as reported here) to find reports which examined human operator situation awareness in cyberspace. A literature search was run using Google Scholar as well as several other databases including ACM, IEEE, and SAGE. Relevant search terms included concatenations of terms such as "cyber", "situation awareness", and "human performance". The articles found were also citation scavenged, leading to additional article inclusion. Researcher profiles and CVs were further sourced for relevant works that may have not been searchable using Scholar. Articles were then reduced in scope by me reading abstracts. In addition to the goals stated above, I looked for:

Development or discussion of any cyber SA theory, and how it might inform our pursuits. While cyber SA theory spans both techno centric approaches and human-centered, I focused on the human component.

(1) Identification of existing human factors gaps. Within cyber SA, what has been discovered, and what remains to be done? What other problems do CND operators have?

(2) Experimental results. These are the most informative reports we can find and serve to addressing facets of our overall project and experimental design. They inform our:

   (a) Measurement – what types of awareness measures have been used? What facets of tasks do they measure? To what degree of success?

   (b) Scenario and operator cognitive role – what operator roles exist, and/or have been examined? Do their roles necessitate situation awareness measurement?

The results of the reviewed articles are discussed next in 2 main sections. The first examines the literature concerning SA in the general cyber network defense environment and discusses the types of tasks and operator roles uncovered that have been previously linked to SA. Visualizations for cyber SA, a large area, as well as the potential metrics for SA are also discussed, as well as the incorporation of automation, and the emergent use of human performance modeling. In the second section, as my primary focus, are papers who have reported experiments using human subjects and measuring situation awareness in the cyber environment are described in detail with critiques.

3.2 NON-EXPERIMENTAL LITERATURE (n = 28)

3.2.1 A Review of Research Areas (n = 2)

Barford et al. (2010) brought together the work in cyber to determine the consistency of the major research areas[2], finding a wide range of research exists. Only one small facet is focused on the human operator, and traditional understanding of SA in terms of performance and measurement. This represents an unfortunate situation, as cyber technology relies on the operator to detect and respond to threats to be maximally effective (Tyworth, Giacobe, Mancuso, and Dancy, 2012).

A prior review of cyber SA (Franke and Brynielsson, 2014) also binned their findings into several categories, listed below, though their bins were still rather vague[3]. More details of the Franke and Brynielsson review will be covered later. It is clear that the domains of cyber situation awareness are vast and complex, spanning from emergency management, to algorithmic development, and human-computer interactions. However this characterization lacks focus on cognitive demands placed on human cyber operators. To understand demands, some researchers have conducted cognitive task analyses, and interviews with the operators themselves.

3.2.2 Details on Cyber Network Defense (n = 2)

3.2.2.1 Conducting Human Subjects Observational Research in a Network Operations Centre (NOC)

Paul provides insight into the challenges of gaining access to the operators and operational environment for cyber. In particular, she recommends embedding observers into the environment, and conducting longitudinal studies that have a minimal impact on operations. In effect she advocates for the types of CTAs that have been done previously, and that are covered elsewhere by our project, Paul (2014).

3.2.2.2 Theoretical SA for CND Operators

In order for the human to play a role in cyber operations, the authors (Dressler, Bowen, Moody, and Koepke, 2014) contend we will have to capitalize on visual, pattern recognition ability (as we do in traditional warfare). However, though it is a fair comparison, these aspects do not always hold

[2] (1) Techniques to reduce uncertainty, and mitigate risk (probabilistic reasoning); (2) Situation representation and modeling, based on transcription of human to machine-interpretable understanding; (3) Automating SA processes; (4) SA across abstraction levels; (5) Reasoning with incomplete information; (6) Using machine learning to gain better SA; (7) Integration of the three SA levels; (8) Identifying SA measures of performance and effectiveness; (9) Information fusion techniques; (10) Investigating the ability for machine "self-awareness"; (11) Attacker behavior analysis.

[3] (1) general cyber SA, (2) cyber SA for industrial control systems (power grid), (3) cyber SA for emergency management, (4) tools, architecture, algorithms for cyber SA, (5) information fusion, (6) HCI, design, work flows for cyber SA, (7) nation-wide large scale cyber SA, (8) information exchange for cyber SA, and (9) military cyber SA.

within cyber because we are limited by the interface. Further, the authors identify that with the currently available cyber tools, *mission impact lacks awareness*. Methods for presenting impact are fledgling, and the authors do not discuss ways to measure this awareness, although presumably one could assess what information is needed to determine impact, and assess how operators build that awareness. It may be that it occurs only during forensic assessment.

Although automation may help and is sometimes present, it is limited and unintelligent (e.g., IDS systems cannot do the entire intrusion task of determining whether an attack has been successful, or necessarily even occurred). Therefore, situation awareness is constricted by the information environment itself, though automation has the potential to help.

Dressler et al. (2014) argue effectively for the relevance of situation awareness in cyber defense from the organizational as well as the individual perspective. Their method of addressing the problem is related to information fusion. The authors stop short of describing in detail what such awareness may comprise of outside vague notions, such as providing awareness of the "threat environment" and "current capabilities" — just that such information needs to be fused to combat the traditionally stove piped data problems found in cyber. The authors point to a variety of additional solutions such as visualizations and improved ability for the operators to share information, but do not at any point evaluate a particular approach.

3.2.3 Operator Task and Role in Cyber Network Defense (n = 5)

3.2.3.1 Network Administration

The activities of a network admin are broad, as they are for analysts (e.g., Alberts et al., 2004; D'Amico et al., 2005; Killcrece et al., 2003; Yurcik, Barlow, and Rosendale, 2003). Admin engage in both the monitoring of the network, as well as staying advised of the latest security information, configurations and patching, attack response and user help. These activities may be split, and roles can be established for things like responding to attacks (the role of an incident response team; Killcrece et al., 2003). For cyber analysts, similarly, many different aspects of monitoring, reasoning and taking action against threats are described (D'Amico et al., 2005).

The role of admin may rely in part on situation awareness. However, much less was found on this role, and the available literature focused more saliently on the intrusion detection system (IDS) operator role.

3.2.3.2 Intrusion Detection

D'Amico and Whitley (2007) reviewed analyst processes. The role of the defensive information operations analysts is grappling with the high false alarm rate of intrusion detection, data overload and sifting through packets to ID suspicious or anomalous activity undetected and may also need to forecast new threats and determine the intents of possible attackers. In this realm, situation awareness is key in supporting the achievement of those activities. The authors interviewed/observed 41 defense information analysts (1 commercial and 6 DoD organizations). They dealt with the complicated, sometimes confidential nature of the environment by using hypothetical scenario construction, still enabling the knowledge elicitation of where to seek information, what information they need and the cognitive connections between pieces of information.

The authors found that as CND analysts work through raw data and categorize activity, an appropriate analogy is progression of confidence in the detection of a threat. This is particularly the case for the role of intrusion detection analysts, who were found to migrate generally from raw data,

pulling out potentially suspicious and interesting activity which was then upgraded to events and incidents with further analysis (Figure 2). Eventually these feed into a set of known methods of intrusion which can be fed back into the ways that IDS systems alert operators.

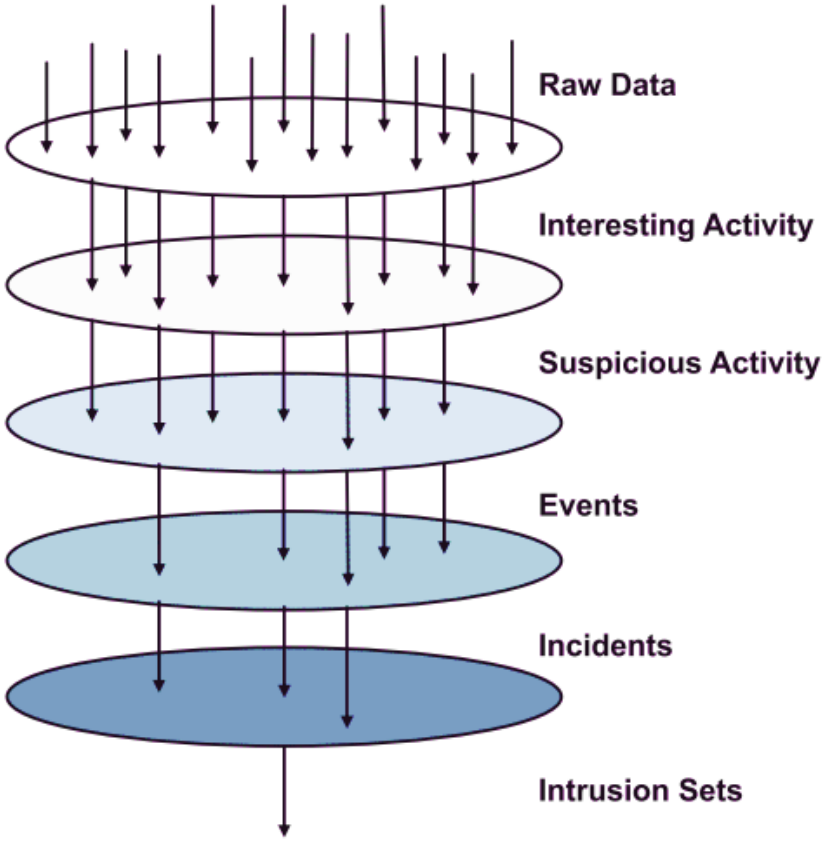

Figure 1. Reproduced from D'Amico and Whitley (2007), this figure illustrates the progression of information toward confidence in the goal state (e.g., correct intrusion detection).

Different analysts may work at different filter levels; *interesting activity* represents data that has already been run through a standard set of filters through an automated process to reduce the data load — interesting activity is data that has been flagged. Flags can be high in false alarms, so analysts at this stage must spend time to determine whether the flagging is real or not.

In order to make this determination the activity is viewed for information on the alert details (from an IDS) and its related data (maybe packet header data). Any information that makes it through this process is moved on to suspicious activity4 and may become known as an anomaly or a likely "event" D'Amico and Whitley (2007). Events are reportable, so influences on awareness that come from successful reporting of incidents originates from this phase or filter. Analysts may be in the role that specifically works to confirm/ gain the best possible understanding of the activity; others focus primarily on triaging potential activity, while others focus on the activity that has been escalated D'Amico and Whitley (2007).

Incidents require a report, which are distributed to wider chains of interested parties. Intrusion sets are related incidents; these are community-wide, known relationships and reflects an official decision.

[4] D'Amico and Whitley provide the following examples; "…a series of scans from the same source IP address; an unusual increase in traffic to or from a server inside the network; virus infections on several workstations in a short time period; and misuse of the monitored network by employees downloading inappropriate content"

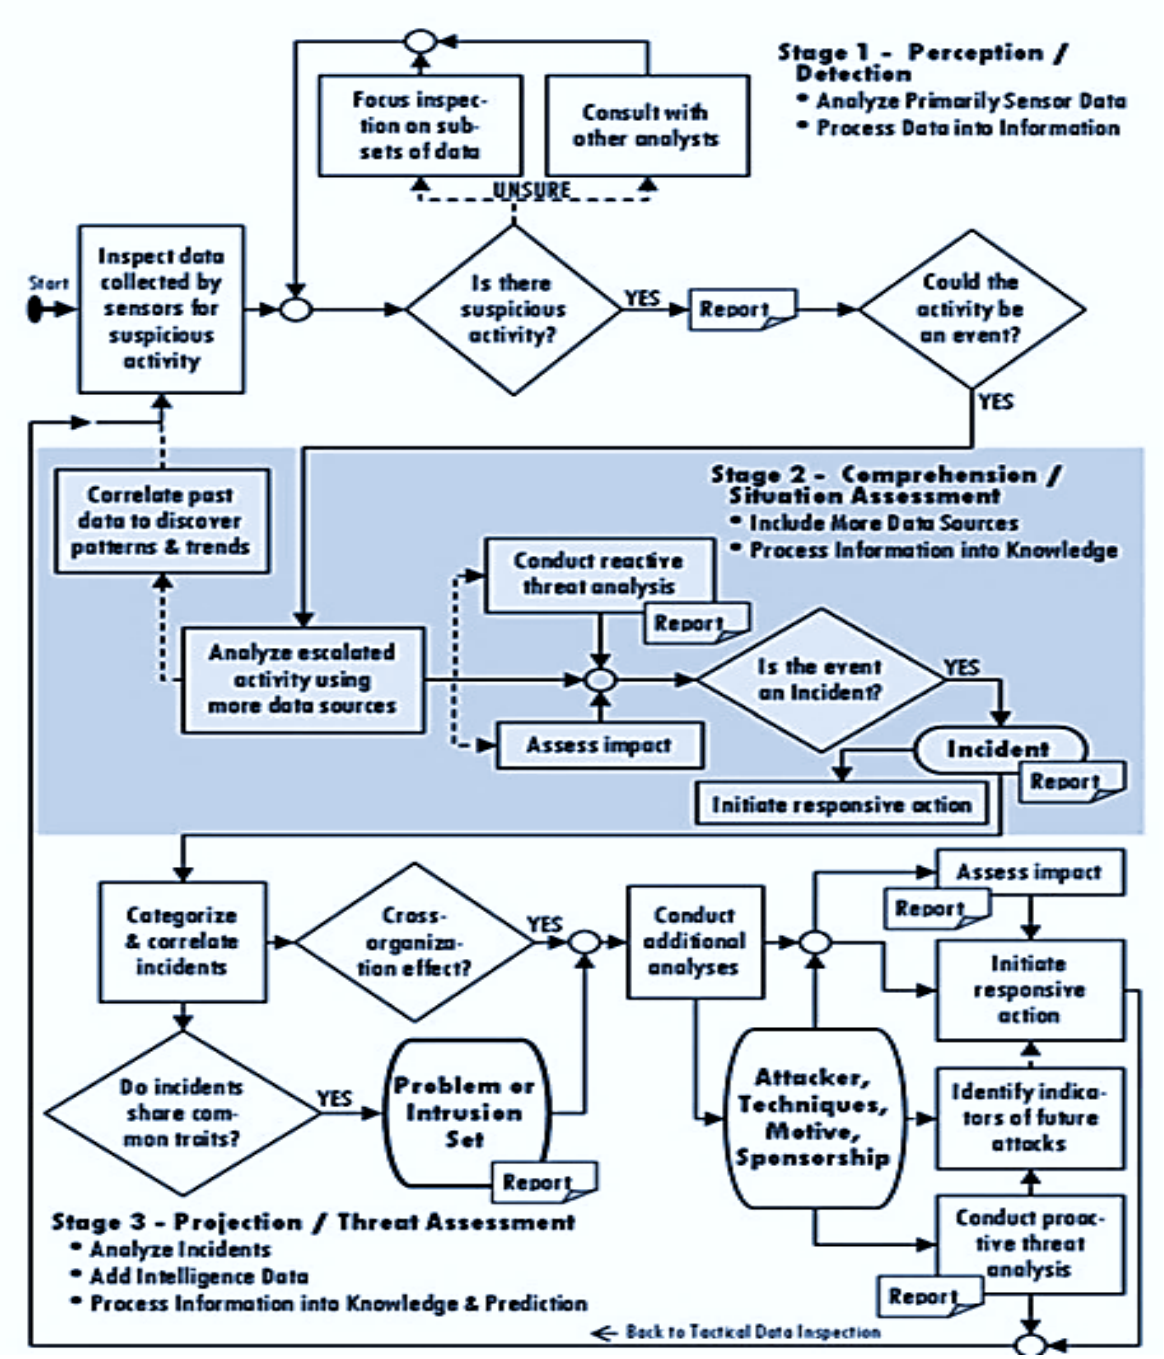

Figure 2. Reproduced from D'Amico and Whitley (2007), this figure shows a linkage between the stages of activity of a cyber-analyst and that of the overall situation awareness levels using Endsley's conceptualization.

D'Amico and Whitley (2007) specifically tied levels of awareness with the activity of the network intrusion operator (Figure 3). A similar representation has not been provided of any other CND task. The authors also identified cognitive tasks for SA in this domain, including "identification and recognition of IP addresses; developing and updating a mental model of normal network traffic; and acquiring and sharing a variety of knowledge including knowledge of emerging vulnerabilities and new exploits."

INTRUSION DETECTION SYSTEM OPERATOR

The authors (Paul and Whitley, 2013) focused on an observation of a specific role at a network operations center (NOC), that of the **intrusion detection system operator**. Operators completed a card-ordering task based on categories of potential operations activities that helped the researchers to understand the general tasks that they did. From this, the authors provided a taxonomy of IDS related work. Of interest were the 2 main categories of **event detection** (e.g., understanding network baseline; noticing changes; and monitoring the network activity) and **event orientation** (e.g., identification of threats; mission impact assessment; damage assessment).

While the examination provided some insight into the operators, once again the observations did not measure awareness. However, by providing more light into the processes operators focus on, it may be possible to derive awareness measures. For example, asking operators about events that may have occurred (e.g., frequency) may assess awareness of the current baseline; asking whether certain information has changed also considers the ability of an operator to be aware in event detection.

SMALL USER EXAMINATION OF IDS SYSTEMS

In this article the authors briefly describe the results from a small user examination of IDS system operators Goodall et al. (2004a). IDS operators focus primarily on monitoring the network; and in doing so they build awareness in part through learning their distinct networks, but also by configuring the IDS system's series of alarms/alerts. There is a sense of community beyond the software. In various online communities, information can be spread rapidly about emerging attacks, and this appears to be expanding. Analysts visit online information sources as part of their daily "startup" activities, .

DISTRIBUTED CYBER SITUATION AWARENESS IN IDS

The authors (Tyworth, Giacobe, Mancuso; 2012) describe why situation awareness in CND is important. Of interest, they describe the results of interviews and semi-structured CTA of cyber analysts, which resulted in evidence toward a distributed SA perspective. Tyworth's findings suggest **SA is only shared when there is task overlap** - and correlated with how much individual's **goals** overlap, and that SA is distributed across humans and the technology. The authors also describe how reports are the main method for transmitting information across the boundaries. **IDS analysts in particular**, used reports to (1) track status, (2) tell others, and (3) serve as an archive. The unique contribution of this manuscript is that it discussed the distributed nature of the work, and that it highlights the role of the IDS operator. No experimentation is presented that back these assertions up,

PHASES OF THE INTRUSION DETECTION TASK

Investigated the process of intrusion detection in a university network operations center by conducting interviews and a task analysis. Interestingly, intrusion detection incorporates a large amount of automation, but operators still play a major role. The intrusion task can be split into 3 main phases (according to Goodall, Lutters, and Komlodi (2004b); **monitoring** — looking for suspicious activity, though on average only 1% of alerts will be actual incidents; **analysis** — some event occurs, signaling a potential intrusion, and then there is some **response** — if an attack is identified, a decision is made about how to respond, Thompson, Rantanen, and Yurcik (2006).

These phases align with general information processing stages (information acquisition and analysis, decision making and action selection). They are also similar to the first 2 levels of situation awareness defined by Endsley (i.e., perception and comprehension). The task analysis also revealed how difficult the intrusion detection process truly is because of its time-critical nature. However,

(Thompson, Rantanen, and Yurcik; 2006) suggests that the decision making does not create high workload on the IDS operators. The focus on the first 2 phases aligns with the emphasis on early levels of SA in other domains (e.g., aviation; Jones and Endsley, 1996).

3.3 VISUALIZATIONS FOR CYBER SITUATION AWARENESS (n = 11)

Situation awareness relies critically on the availability of information that is useful to the operator. The operator uses their perception of this information to comprehend a situation and make an informed decision. Unlike other domains in the physical world, the cyber domain is completely dependent on the interface for understanding this information and sampling it. To the extent that the interface can increase or improve the presentation, understanding, or projection from information, then SA may also reap benefits from better visualizations. However a review in this area, though not focused on SA, has already shown a glut of interfaces that do not receive user evaluation or experimental examination (Staheli et al., 2014).

3.3.1 Requirements for Visualizations and a Survey

A survey of types of visualizations for types of network demands. Not many visualizations meet all three of the following requirements (relevance to network security; contribute to system and visual techniques; and have a satisfactory evaluation in place), Shiravi, Shiravi, and Ghorbani (2012).

The authors champion the need to go to CTAs and users to make sure that at each step, the user is consulted in what they need to perform. 3D displays are pointed out as *particularly* difficult to perceive and interact with, in addition to requiring a burden of needing fairly constant interaction, to an already overburdened analyst. 3D displays also hold the potential for obfuscation of information, thus reducing level 1 SA.

3.3.2 Challenges of Visualizations and User Observation

The authors use the majority of the paper to outline several cyber operations problems, and describe how operators work, and explaining the role of confidence, and data aggregation methods. They then highlight *different aspects of existing visualizations that do not work well for cyber analysts*. In particular, most tools do not aid reporting or collaboration, but the analysts need that feature, showcasing a usability gap that could be addressed by user-centered designers, Fink, North, Endert, and Rose (2009).

The second portion discusses their observation of 4 operators attempting to solve scenarios generated for the 2009 VAST contest, based on physical building access and network traffic data (a correlation problem). Users could typically pick out abnormalities easier using the visualizations; but each of the 4 analysts appeared to go through a different process to test their hypotheses. An interesting problem of **goal-resumption difficulties** was identified. The authors observed that operators "...following hunches may… [travel] down different paths of analysis, but often stated that getting back to where they decided to pursue one of these branches could be very difficult" (Fink, North, Endert, and Rose (2009).

Four design principles were then derived from the observations: (1) provide history and traceability of the humans' activities; (2) help support multiple concurrent cases; (3) design visualization tools for flexibility, broad spectrum support; and (4) enable direct manipulation of the visual space. In effect, these recommendations are generally to improve situation awareness. The authors do not address how to evaluate this measure, and their recommendation validity is left unanswered as none of their recommendations were evaluated. They must be treated as anecdotes.

3.3.3 Three-Dimensional Visualizations

The authors outline the 2 main types of mission impact analysis for cyberspace, "bottom-up" and "top-down". Each type of analysis is focused on determining the assets - related to a mission - that must be operational in order for the mission to be successful, either by starting from the asset and generalizing to a variety of missions, or vice versa, D'Amico and Salas (2003).

A 3D representation was developed to aid operations, and the relationship to elements and the mission critical functions was illustrated. A layer could be added to reflect the types of operational missions (and what mission critical tasks they depend on), and the system allowed for constraints in data based on time, location, mission phase, or protocol. Finally, a layer of IA "alerts" was imposed based on reporting, which could also be layered into a map display to examine physical / geographic points.

The design of this interface is intriguing, appears to support the need for operators to hypothesis test in network defense, and draws from similar efforts in decision support in military settings. However, as is the case for almost all of the work examined in this review, no experiment is conducted that would solidify the utility of such an interface. Nor is it compared to baseline performance or awareness, as neither has metrics associated with it.

3.3.4 Challenges of Visualization Design

The authors present the case for why examining visualizations for cyber analysts is important, but still has many challenges. They mention that the **visualizations are not always helpful**, and that there is a need to explore the domain for why they are not readily adopted. The authors do this by user questioning. The interactive, communication-based role of some operators is emphasized; the authors describe the need to establish common ground between the system and the human operator (though this is fraught with complication; Conti, Nelson, and Raymond, 2013), Best, Endert, and Kidwell (2014).

Guidelines for designing displays are derived. For most visualizations, time is recommended to be held as a constant in the system to allow for operators to make comparisons. **Comparing across time frames is noted as particularly difficult for operators**; there is also a clear need for feedback from the task. In place of feedback, operators tend to rely on a sort of network "cadence", a measure of operational tempo. A visualization could attempt to more easily present what is 'normal' versus what is abnormal, but these types of representation are difficult to articulate.

Aspects of visualization design are further emphasized. First, making things **observable** (aligning with level 1 SA), in terms of providing patterns and relationships in context and over time. This may increase the ability of the operator to use their experience in developing an understanding of the environment. Second, the visualization should create salient indications of change (also level 1). Echoing others, the authors showcase the need to aid **team work,** suggesting that a visualization can depict relationships between the operator and others operating on the network. Utility of the proposed design is described as being assessed at the time of publication.

3.3.5 The VIAssist Visualization

These papers present and summarize the functioning of their interface for CND analysts, VIAssist. In particular, **VIAssist is one of the few visualizations developed from a user-centered approach.** The authors based design on cognitive task analyses conducted with actual representative analysts, thereby incorporating better organization of information to allow for natural search and discovery processes; methods to "tag" data for other analysts to evaluate; data fusion techniques to link data

together; a “details on demand” approach; customization of the interface; and the use of geographic visualizations, D’Amico, Goodall, Tesone, and Kopylec (2007); and Goodall and Sowul (2009).

However, where the papers fall short in both cases is in evaluating the interface, though perhaps this has been conducted since. There is no baseline with which to compare the system; and no evidence to suggest that the system is indeed more effective, and the tool claims lie in their development pedigree (which is stronger than others found in this review).

3.3.6 Use of Virtual Worlds for CND

The authors suggest that we can use virtual worlds to examine network traffic, and that this should therefore improve cyber SA. However, no awareness theory is discussed or integrated, and no experiment is run. No sense of how to measure that SA is claimed either, Michel, Helmick, and Mayron (2011).

3.3.7 Cyber Defense Management

Authors point out that a relevant cyber SA problem is management of the cyber defense mission (and the people performing it). Developing shared SA then becomes an issue of importance for communication and delegation. The sample interface used in the paper shows mission-level information; when people do things and respond to events, the interface shows this through tasks across the network. This interface appears to address a need for cyber operators, and that is communication as highlighted in other papers. However, it makes no progress on actually evaluating the need a priori, or quantifying the potential help the tool could provide, Paul, Rohrer, Sponaugle, Huston, and Nebesh (2013).

3.3.8 Information Framework for Enhancing SA

The intrusion detection role is to detect potential threats, thus providing some helpful awareness. A lot of cognitive load is wrapped up in using current systems. Providing visualizations is speculated to reduce that load, but most visualizations do not have user studies associated with them; the authors provide a *user* study, though it is not a traditional experiment. Komlodi et al.’s user study emphasized the need to have users view and interact with the tool; the study uncovered that analysts do not simply focus only on the initial intrusion detection task - they have other duties as well, Komlodi, Goodall, and Lutters (2004). As operators transition phases from monitoring to analysis, ***a tool should therefore allow for more flexibility***. Essentially, the need for exploration reinforces that awareness is built through the search for information (selection in context, exploration, correlation and multiple sources to allow for multiple views of the same information). Though they do not link to Endsley explicitly, the viewpoint is that of the Endsley SA theory.

The authors are unique in my review for pointing out that current visualizations focus on the monitoring (level 1 SA) phase. The important analysis phase (and response) also need help. By incorporating a facilitating tool **for *both gathering information and linking it***, they are suggesting the tool will improve levels 1 and 2 SA. However, they do not measure any SA, nor do they experiment with the proposed system.

3.3.9 Moving from Text-based to Visualization of IDS Information

The authors review and discuss the tools used by intrusion detection analysts to help them monitor, analyze and response to potential threats on the network. In part these comprise tools designed to enhance situation awareness. **Many found in this examination are text based**, and encompass global tools, logs, outside information from web-based resources on latest intrusions and attacks, and samples of code. However Thompson et al. pointed toward many research efforts attempting to develop visualizations of these types of information, in an effort to reduce the effort that overloaded

operators contend with in building situation awareness of the network. As has been the consistent case, the work does not evaluate the claim, nor does it lay out a plan for examining SA in the intrusion detection task. However later experimental work (see the experiments section in the current paper) does test some of these ideas, so there is a nice continuation of the conceptual question, Thompson et al. (2006).

3.3.9.1 Usability Study for a 3-Dimensional, IDS Visualization Tool

The authors describe in detail a visualization tool, and briefly discuss the usability study results for an IDS tool. There was uniform agreement on **the importance of time** as an aspect of ID data (also echoed in their other work). They attempted to include flexibility, something that user surveys agreed with, in their tool by designing a 3-dimensional glyph-based display, Komlodi et al. (2005.

However, their usability study showed that participants had a great deal of difficulty with the basic use of the display (e.g., zooming, panning, and rotation), suggesting that UX and interaction designers may play a much-needed role in tool improvement. In particular, these were difficult tasks because of the **concurrent demands to monitor other information sources outside of the display**. I suspect this is another case of a goal-resumption problem, in which operators are side-tracked and have difficulty recalling what they were doing before interruption (e.g., Altmann and Trafton, 2002). The authors have done a good job at highlighting the cognitive constraints that IDS operators in particular face in a task management context.

In summary of the papers found here, information visualizations can provide advantages, and there are good reasons to believe this. However, the advantages must be documented, they must be shown to exist. Cyber situation awareness is rife with visualization efforts; but almost no evaluations at all (Staheli et al., 2014). The above findings indicate that visualizations are mostly ideations - some visualizations found in survey were essentially a proof of concept (e.g., Williams, Faithfull, and Roberts, 2012). The design goals of these visualizations are unfortunately, and likely unintentionally divorced from the user and their goals. Most visualizations are not based on a nuanced cognitive task analysis of real operators (Komlodi, Goodall, and Lutters; 2004), with some exception, and almost all are not evaluated for usability or for its influence on operator SA.

3.4 CYBER SA METRICS AND MEASUREMENT (n = 3)

As I described earlier in the sections on SA theory and measurement, a variety of established methods exist to assess SA. Often, authors have sought to establish their own measures, which lack validity testing even if they are potentially useful. It is instead those papers that suggest what aspects of behavior *comprise* SA, that are useful and are discussed in this section.

3.4.1 *Levels of Cyber Situation Awareness*

The authors begin with describing the brief history of examining situation awareness and the cyber defense task. **They define specific examples of different levels of awareness.** Level 1 could include perception of IDS alerts, their timing and origin. Level 2 may comprise the sorting process, moving from many alerts of activity to determining that certain servers or mission-critical resources require attention. Level 3 is essentially what helps an operator determine what to do next, in their understanding, D'Amico and Kocka (2005).

The authors describe a particular problem in network defense, in building good Level 1 SA. Operators need to look for any source IP "scanning many destination IPs within the same site's network traffic." This issue may be solved using traditional scatter plots and a filtering capability. They identify a Level 2 problem with determining what the attacker has contacted and/or exploited — a way of showing a network path, helping the operator visualize and share a sequence of attacker actions. The main Level 3 problem for CND is forecasting potential new exploits — the authors say this is currently addressed in part by tools like the *Analyst Notebook* but could be done using timeline and link analysis. They point out that displays/visualizations are likely to be situationally dependent, and efforts to create a global one are misguided. The contribution is one that is one moving forward the need to measure, and experiment with cyber defense SA.

3.4.2 Teaming for Situation Awareness

The authors present their conceptualization of SA, based on analyst interviews, as distributed across multiple individuals. They are one of the few papers (in addition to D'Amico and Kocka, 2005) who provide fairly explicit examples of how a cyber task maps to Endsley's definition of SA, Tyworth, Giacobe, and Mancuso (2012).

"For example, an intrusion detection analyst who detects a host on the network "beaconing" (transmitting data a regular interval) to an unknown host outside the network can be thought of as having Level 1 SA (perception). If the analyst correctly determines whether the activity is malicious or benign they have Level 2 SA (comprehension). Finally, if the analyst takes the correct action based on their comprehension — blocking the traffic if malicious / allowing if benign — they have achieved Level 3 SA (projection). The higher the level of SA, the more likely the human operator is to take the appropriate action; conversely a lack of SA is associated with operator error" (Tyworth, Giacobe, Mancuso; 2012).

(Tyworth, Giacobe, Mancuso; 2012) also point out that a significant amount of awareness could be provided to different roles of operators by improving the communication between them. Thus, they are arguing for a team and distributed SA perspective.

3.4.3 Difficulty in Cyber Interface Evaluations

Levin and colleagues wrote a coherent paper on the need, and difficulty, of evaluating interfaces for cyber situation awareness. The authors provide a look at the potential methodologies for evaluating SA and performance in cyberspace (Levin, Tenney, and Henri; 2001).

The authors outline several constraints brought on by assessing cognitive performance in the cyber realm. Normally these are controlled aspects of a laboratory experiment, such as operator familiarity with the task, comparable trials of controllable complexity, and elimination or control for carry-over effects. Further, the need to replicate real-world conditions like complacency were stressed, pointing out that the system should be tested and not operator abilities; but this must be weighed against establishing control for operator experience. They also describe a small set of data displays that may be candidates for comparative evaluation: (a) integrated data displays; (b) correlation displays; (c) course of action recommendations.

Several candidate measures of SA are also proposed: (1) choice of the correct response or category of the existing situation; (2) time take to reach correct understanding; (3) measuring the information type, and frequency of consultations (such as which screens or applications) were considered; (4) establishing a degree of SA through various other methods, such as self-assessed, with probes, or observed through SME ratings.

Finally, **following the methodologies, the authors outlined why they did not perform an experiment**. These were (1) credibility and fidelity concerns about scenario realism; (2) an absence of a true baseline of performance with a display; (3) underdeveloped display needs and concepts of operations; (4) Difficulties determining what is a correct response: It could be determined by the policies that may be in place, or learned through individual experience; (5) a complete lack of a paradigmatic approach or doctrine for CND; and (6) the heavy abstraction of the interface itself.

In future attempts to measure and evaluate cyber tools, it seems there are serious challenges to overcome. Scenario realism is always a concern in applied study, and a fine line is generally walked between fidelity and cognitive realism. It is indeed difficult to baseline when tools are highly specialized. Certainly, the needs of the defender are growing in clarity because of the work done on task analyses and work domain analysis. As with all SA research the objective truth and scoring can be difficult as well.

3.5 USE OF AUTOMATION TOWARD CYBER SA (n = 3)

The use of automation is unsurprising in the cyber environment. The use of automation may potentially improve operator awareness by enhancing the process of perception, comprehension or prediction.

3.5.1 Attempting to Automate Cyber SA

Authors discussed the notion that many aspects of cyber situation awareness could be automated. The authors claim that since information can answer many of the main questions that operators may have about the network environment, such as “how did the situation arise” and “what relationships between objects are important”, then these processes could be automated, Holsopple et al. (2010).

However, though they describe visualizations (as so many of the discovered papers do), they do not show any experimental evidence to support an assertion, i.e., that the processes can be automated, or that this would be an effective method to help CND. In fact, others have specifically suggested this is perhaps the wrong approach to cyber operations, opining that we rely on human experts versus automated systems or devices (Yurcik, Barlow, and Rosendale, 2003). As discussed in the visualization section, there is a significant worry that cyber SA has garnered too much of a technocentric focus, losing part of its etymological roots (e.g., Bass, 2000). That is, work is more concerned with developing more tools and technological approaches but is largely ignoring the operator and especially their assessment.

3.5.2 Using Automation to Organize SA Information

The authors argue for a broad SA model merging individual, shared, and “complementary” situation awareness. The automation component appears to be primarily related to a “data center” that aids in these processes, Cain and Schuster (2014).

3.5.3 An Autonomics-Like Approach

The paper is a proposal to develop a system. In this short paper, the authors describe the notion of cyber situation awareness and suggest a system that would detect the general health of a network, and then be able to understand how changes in that health affect missions, Okolica, McDonald, Peterson, Mills, and Haas (2009).

The idea begins to sound like an autonomics framework, however, no mention of that term is presented. The authors suggest use of pattern recognition ability to determine how to proceed given certain information; and in doing so, improve situation awareness. They suppose that a cyber SA system needs a language to describe the data and fuse it together; but they do not provide it.

One conclusion from the ideas here is that automation will not soon provide a clear method for improving SA. Importantly, these endeavors appear completely divorced from the user; how can automation improve a process if we do not have a baseline measure, or even a measure for comparison?

3.5.4 Modeling Efforts (n = 3)

Modeling the human poses many potential advantages to the researcher. In addition to forcing a clarification of cognitive processes and the tasks needing to be performed, once sufficiently reflective, a model can help predict how operators will handle a wide variety of situations without running human subject experiments. Models also guide theory development. The current modeling work in cyber is just developing, and for situation awareness specifically the work found was sparse. Not all modeling focused directly on the cognition of the operator, and not all modeling truly addresses the situation awareness component, but some is useful for the purposes of this review.

THREE MODELING METHODS:

MODELING THE CSIRT ANALYST

The authors are attempting to model the human operator as part of the CSIRT analyst process. The work is interesting, but very early in the research process. It was unclear if their model is valid, as it appeared to have a very low correlation with their other data set. Importantly the model appears to be missing components for the cognitive costs of task suspension and resumption, though they do attempt to account for operator experience in general, and with specific tools used (Reed, et al.,2014).

This work has promise, as cognitive modeling helps understand the processes that the operator uses to develop an understanding of the environment and make decisions. However, it is difficult to see the clear tie to SA here and the model is in early development.

A MODELING APPROACH TO DETECTING INFORMATION FOR CYBER SA

The authors point out the necessity to consider the user in the cyber environment, and the importance of SA. They describe the need for better approaches to presenting situation awareness information to the user, based on a modeling approach and briefly describe it. They note that visualizations should be useful, based on a discussion of different, basic, and sometimes vague attention and perception phenomena, mostly situated in anomaly detection. Of course, not all cyber situation awareness information is based in anomalies, Klein et al. (2010).

The authors also describe screen space allocation issues in cyber as being related to difficulties in creating visualizations, but yet rely on visual interfaces. Operator / information overload is further mentioned as one of the key issues for visualizations to help mitigate, though this could easily turn into *visual* overload more than *cognitive* overload.

MODELING AND VALIDATION OF THE SA OF CYBER DEFENDERS

The authors built a cognitive model of a cyber-defender based on instance-based learning theory. The model was designed to examine the role of operator task experience (e.g., what threats the operator had seen previously) and tolerance (how much evidence has to accumulate before an event is called an attack by the operator). In addition, they examined the role of attacker strategy (fast, or low and slow). According to their model results, attack strategy plays a large role in whether attacks are successfully detected; the “low and slow” type are much less likely to be detected. When the

model represented operators as risk-averse, and to have seen prior threats frequently (e.g., a threat-prone Bayesian representation), the outcomes were far superior to the converse when attacks were fast but resulted in worse recognition of the low and slow type. They conclude the best training recommendation would be to make sure operators' stance is threat-prone, (Dutt, Ahn, and Gonzalez, 2013).

However, as the focus was the model, not as much was said about actual situation awareness. While experience (e.g., Bayesian understanding here) is theorized to play a large role in the ability of an operator to build situation awareness in other domains, this link is not clearly made in the current work. No measures of specific situation awareness were solicited, and no human participants were used in this early work. The understanding is that this is foundational, and that it will be validated and expanded upon.

3.6 EXPERIMENTAL LITERATURE (n = 7)

3.6.1 Observing a Cyber Exercise

The pseudo experiment was run during a basic red-blue teaming exercise. The competition environment contained a scoring engine (automated) to allow for performance evaluation based on checking / evaluating the services needing protection by service access attempts. The researchers collected video and audio recordings of the blue and red team to correlate with the any events of interest, (Malviya, Fink, Sego, and Endicott-Popovsky; 2011) .

In this study, measures of SA were based on the timing and performance data collected via simulated "user requests." Four researchers asked SA questions of Blue/Red team members using the SPAM methodology. The authors constructed a taxonomy of questions based on 2 perspectives (defense or threat-related) regarding policies, priorities, and events, organizing questions into 3 sections based on time (past hour, present, or upcoming hour). A list of 48 questions, was sampled every 40 minutes when one question was selected and asked to all 7 teams, for a total of 22 questions asked. Questions were asked to random team members, in order to attempt to sample team SA, and scored according to the SPAM methodology (reaction time was how long before the participant began answering a question).

The experimental results should be viewed under serious scrutiny. For example, while the purpose was to correlate situation awareness with performance, the relationships in the paper may be misinterpreted. They found a *positive* correlation between reaction time and SA accuracy, but this means that with increasing time to answer, there was increased accuracy (they interpret it as increased time, decreased accuracy). When SA accuracy (again subject to the same problems mentioned above) is correlated to the internal performance measure for the team as a whole, the correlation is moderate. Neither of these values was shown here to be significant by any statistic test.

It can also be argued that the measures here is not representative of team SA, and in particular that the SPAM methodology has not be shown to be effective for such uses. In addition, there are scoring concerns uncovered; for example, though 22 questions were asked, the researchers only scored 3 because the others "did not lend themselves to accuracy assessment." It remains unclear why. Many teams did not answer many of the questions, as well. The researchers even state that the metrication specifically took a back seat because of the priority of the exercise and it was possible that no one took the questions seriously.

Overall, while there was great potential in this assessment, the results are uninformative. The contributions are difficult to disentangle and parse for use. Even SPAM questions used were not described in detail, and so cannot inform any future work.

3.6.2 An Experimental Platform and Experiments Conducted

The authors begin by describing a simulation micro-world platform that is designed to support experimentation with cyber tasks. It appears that idsNETS can simulate activity similar to an intrusion detection task.

The design of three different cyber-security task experiments that used the NETS simulation engine were described at a top level. Each was one of the authors' dissertations, (Giacobe, McNeese, Mancuso, and Minotra; 2013), and each is identified and described below as in the paper.

(1) ***"Transactive memory with teamNETS" – Mancuso's dissertation***. 66 teams, with a focus on the collaborative component. 3 functional domains assessed; intrusion and threat analysis, operational and sys admin, and management and policy. Participants in teams of 3 were each given different knowledge, based on whether these roles were separated or integrated - with the idea that solicitation or passing issues would result more often when a single member was an 'expert'. Only minimal impact on performance was found, but teams with more similar skill sets had higher communication and information sharing.

(2) ***"Task prioritization in NETS-DART" – Minotra's dissertation***. 77 participants, with a focus on dual-tasking attention and individual cognition. Attempted to look at the use of a workload preview, in order to help operators schedule out their effort and help to determine when they should correctly switch tasks. Workload preview interacted with increased time pressure to influence performance, but was not shown to be an improvement in the experiment.

(3) ***"Situation awareness in idsNETS" – Giacobe's dissertation***. Looking at how an analyst may work with log file data similar to an intrusion detection system. Giacobe compared a more traditional text-based interface with that of a visual-analytic one; performance was assessed by an in-house capability provided within the simulation for examining accuracy and speed. As discussed in (Levin, Tenney, and Henri; 2001), accuracy and speed may not always be valid measures for assessing visualizations. Nevertheless, the visualization interface improved performance according to this measure. No correlation was found between 2 different measures of SA used (SAGAT and SART), and no differences existed in the SAGAT ratings. The subjective SART ratings showed that ratings of attention division, information quality, and familiarity were higher with the visualization compared to the text-based interface. These are interesting but require additional objective assessments.

While the above is promising, as research is being conducted on SA in cyberspace, there are some caveats. First is that while dissertations are promising, they do not always have the clarity and quality of published work for various reasons. Second, only one of the above dissertations addressed SA and it should be commended for using multiple SA measures. On closer examination, Giacobe's dissertation measures of SAGAT may not always hold up as truly diagnostic of a situation awareness perspective. For example, some of the questions appear to be common knowledge questions, rather than questions about the current scenario, e.g., "Every IDS alert is an indicator of a problem" is most likely false regardless of the current environment and situation. Other domains suffer from this same problem in measurement.

3.6.3 An Examination of the Role of Experience and Visualization on Cyber SA

An experiment was run with 23 participants to determine the effect of a visualization (versus text-based interface, as in Giacobe's dissertation), and of "experience" (novice vs expert) on a cyber defense scenario. This was accomplished using data (vulnerability scans, firewall, Snort, and server) from the 2011 VAST Challenge and approximately 1 days' worth of data to be analyzed, though

server data was excluded. The participants used either a visualization, or a text-based setup to locate data, select a location where a problem is occurring and select the time, and then "file a report" by clicking another button. Events had "correct answers" based on identifying the location and the time. Participants were interrupted to complete one SAGAT questionnaire during performance, and following each of 2 scenarios, completed a SART and NASA-TLX subjective report, Giacobe (2013).

At analysis, Giacobe found that situation awareness did not differ between interfaces or experience groups. Nevertheless the visual interface group outperformed text-based. Experience had no effect.

Although the authors use the SAGAT and SART method, the authors do not give examples of questions. One can surmise that it is similar to Giacobe's dissertation measures, which have some issues (outlined in the Giacobe et al., 2013 summary). It therefore remains difficult to determine whether these questions are truly assessing situation awareness or could be useful elsewhere. The task fidelity was increased by re-presenting past real-world data from a VAST cyber challenge, so they succeed in being ecological which is a nice trend in the experiments found by this group and again deserves commendation.

3.6.4 Comparison of SA between Command-Line and Visualization Interfaces

The authors conducted a cyber-tool comparison experiment with 12 student participants (and 2 network engineers), using network intrusions over various cyber defense task sessions. There were 2 intrusion types (a peer-to-peer connection, and a port scan) and a session without an intrusion. The authors sought to compare a visualization, versus a command-line based tool for cyber performance. The implication is that visualizations are supposed to improve SA, and thereby might improve performance, (Thompson, Rantanen, Yurcik, and Baley; 2007).

Participants completed phases of the ID task, including monitoring and analysis (as outlined in Goodall et al. above). Their tasks included receiving an alert to some anomalous activity, determining the cause of the alert by gathering information on the network, and deciding if further investigation was warranted.

The experimenters measured the total time on task to identify the attack, number of commands used (interactions in the visual case, or commands typed in the command line condition), accuracy (attack or not, and type of attack), confidence in their identification, discovery of other problems, and a likeability rating of the interface type. It should be noted that none of these directly addresses the issue of situation awareness, although time to detect an attack may be indicative to the extent that performance and situation awareness are related in the cyber task.

Command line was more effective in terms of speed (about 40% faster), and in determining exactly what the problem was (though they were equal in determining that there *was* a problem). Authors recommended combining them. However, no situation awareness measures were collected. While the main point of the experiment was to evaluate the claims of utility and improvement in SA that are touted by visualization interfaces, these claims were only partially evaluated. In this case, the authors recommended keeping them together because these each provide unique advantages to the task (though the users preferred the textual interface).

3.6.5 Distributed Cognition

This paper described an experiment to study the effects of varying knowledge structures on distributed team cognition. Using the *teamNETS* simulation, integration and differentiated knowledge structures were manipulated by varying the reference materials the participants received during training. While the 2 knowledge structures had no direct effects on team performance, other results

were found in collaborative processes, and team perceptions. Specifically, the results showed that teams with differentiated structures worked more independently of each other, simply coordinated their actions and exhibited minimal communications. Teams with integrated structures worked more interdependently, with a much tighter collaboration and frequent communication, Mancuso and Mcneese (2012).

3.6.6 Visualization is Better!

The author presents a contrast study in which 2 network analysis tools are contrasted in a group of novice network users. The TVN tool is designed in part to educate and facilitate network learning and is the nonstandard tool. In contrast, Ethereal (Wireshark) is a standard networking tool for packet analysis, Goodall (2009).

Participants were evaluated in a series of timed tasks asked to be performed with one of the tools. These tasks were representative of those performed in intrusion detection. There were 2 general types. In the first, there was one possible answer, with 2 component requirements: compare and identify. Comparison required the participant to make a judgment about data size and needed to examine the entire data set to answer correctly. Identification required the participant to "locate and identify an entity" given a set of characteristics, and only required a small subset of data to sift through. In contrast to the single solution type of task, a second type (Exploratory) could elicit open-ended conclusions. Exploratory tasks measured the number of correct insights, unrelated to the well-defined tasks.

The results showed that over all responses, participants were much faster for Wireshark than for the TNV tool; however, these results appear confounded by the apparent interest in the TNV tool. Still, given that learning is one intended purpose of the tool, this should not be looked upon negatively.

There was significantly better accuracy overall all tasks for the TVN tool compared to Wireshark. This effect may stem from an interaction with the type of task, finding a difference between the 2 tools only in the comparison tasks, but not the identification tasks.

TEAMS FOR CYBER DEFENSE

Measured operator cognition in the cyber defense task including measures of situation awareness. Uniquely, as part of their effort they attended the National Security Agency's "CDX" exercise which **used Air Force cadets as participants**. The authors used a cognitive task analysis method to determine 3 areas for focus for their experimentation: overall team organization, team communication, and the negative impact of information overload (Champion, Jariwala, Ward and Cooke; 2012).

The authors then used an emulator, *CyberCog*, to conduct team-based cyber defense research. The emulation contained elements from normal cyberspace defense environments, including alerts, logs, mapping of networks and vulnerabilities, with fictitious users. (CyberCog has transitioned to DEXTAR; Shope, 2013). Participants were grouped into teams of 3. Each team worked to classify alerts into one of 4 classifications: reconnaissance, false positive, failed attack, and attack. Each team also used 2 collaborative displays; one for sharing classification events, and one attack path display to indicate compromised systems and the order the compromises occurred.

Each operator received training on a specific type of security breach not explained to other members, in addition to training on the overall testbed and mission goals. This was similar to Mancuso's experiment above, except that here participants were specialized after initial training. The roles all had the same visualizations across 4 of the 5 available interface tabs, containing information

on intrusion alerts for classification, completed logs, events as they were being classified into intrusion alert groups, and the ability to search the users present in the system. A manipulated fifth tab showed one of 3 displays: either a system vulnerabilities display; a tab with organized information for attack methods in a wiki format; or a network map of the scenario.

Eight teams classified alerts and attack path orders over 2 scenarios. Teams were scored on how well the team classification and attack path ordering matched ground truth. Ratings of subjective SA and confidence were included as part of the completed team report, although the report was necessarily limited in detail on these aspects. On average, teams were subjectively "somewhat" aware of network activity across both practice and test trials (Champion et al., 2012). Most interestingly, however, is that level of confidence in categorization accuracy was far higher than actual performance and correct identification of attack paths. While it may be that confidence and subjective SA were similar here, it **lends more credence and demand for objective SA measures**. Finally, the authors make a point to identify that load (based on moving from the easier training to more difficult, with higher event-frequency test trials) may have explained a decline in awareness.

4. BRIEF FINDINGS

I examined the available literature (ending with literature published up to March 2015) to find relevant experiments which contained the theory and measurement of situation awareness in cyber operators. I attempt to summarize these main findings and gaps in knowledge and this section serves as a conclusion to the current review.

1. Published experiments with human operators in cyber are infrequent although they do point to the potential usefulness of SA measurement, and the potential for useful experiments in this domain.
2. Experiments varied drastically in quality; most suffer in methodological or statistics assessment, do not address measurement issues, lacking a strong basis in any evaluation theory, though exceptions were found.
3. The definition or use of SA metrics is not a focus across the found literature. In the small set in which they appear, metrics are sometimes vague and not tied clearly to the existing human factors theories. A few papers do attempt to replicate the basic methods of SAGAT, but do not provide clear examples of their questions.
4. Visualizations dominate the research understanding of situation awareness for cyber operators, and ideas are abundant. Unfortunately, assessment is extraordinarily rare, and as one potential consequence, many tools are discarded by operators.

4.1 GAPS AND OPPORTUNITIES

A list of 5 gaps and opportunities is provided, and is the summarized content from the current review:

1. I identified that a strong, theory-based approach to measure and analyze human situation awareness in cyberspace would address a critical gap. There are several remaining difficulties.
2. SA for a particular environment or role needs definition. There are several roles of the cyber operator, some of which are much more explored than others (e.g., IDS analysts) which may present opportunities to focus work. It is more likely that SA should be assessed across the entire system, which would be a lot of work but potentially fruitful, in the style of distributed SA.
3. Development of measurement of awareness, and determination of what types are most useful (e.g., subjective? Objective? Implicit?) remain to be understood. There is great promise in the idea of distributed situation awareness, it is simply not as easily performed; yet future proposals for funding an effort may be very fruitful.
4. The influence of cyber SA on cyber *performance* is not clear in the currently found literature. In other words: SA is important but what aspects of cyber defense performance is it affecting? A similar argument could be made for other cognitive facets of the task, such as: the influence of workload, stress or fatigue, and training/expertise on both SA, and cyber defense performance. It is unclear what such performance could be defined as, whether decision making, reaction time, or some combination of communications and informational reporting.

5. Attention switching is crucial and likely implicated in situation awareness. For example, there is repeated mention of a “goal-resumption” problem for cyber analysts, who must juggle multiple hypothesis tests and switch between them. Further understanding how operators task manage in the cyber environment may be of utility; attempting to model this may also provide insight into SA and decision making.

5. NEXT STEPS

5.1 DETERMINING OPERATOR GOALS AND SA MEASURES

As the title of this section suggests, given the close reliance on understanding operator goals in measuring and analyzing situation awareness, determining specific goals is very important. This understanding will help build measures of SA, since they are not easily availed from the existing literature. Methodology exists to support this development (Endsley and Jones, 2012).

In general, a research agenda for human factors in cyberspace will be multifaceted (e.g., Gutzwiller et al., 2015). While SA is one key component, and certainly has not received the research focus it deserves, other aspects of cognition (training, decision making, task switching) remain important as well.

This page is intentionally blank.

6. A NOTE FROM THE AUTHOR FROM THE YEAR 2018:

As may be gleaned from the current document, although it is finally being released in 2018, the original review ended in 2015. There has been much activity following this initial survey of the literature. It has moved from the current form of annotated bibliography, toward integrative review, and I continue to add research not listed here to a more formal assessment of the area. *A full (and updated review) is forthcoming* and I intend to provide more provocative insights therein. For now, it remains clear — situation awareness is key in cyber defense, but it has not received its due research focus as of yet.

In recent years, the research community in human factors has leapt to action since 2015 and the writing of other "call to action" papers (Borghetti, Funke, Pastel, and Gutzwiller, 2017; Gutzwiller, Fugate, Sawyer, and Hancock, 2015; Vieane et al., 2016b) and more experimental and analytical papers on the human factors of cyberspace operations is growing (e.g., Aggarwal, Gonzalez, and Dutt, 2016; 2017; Brynielsson, Franke, and Varga, 2016; Gutzwiller, Ferguson-walter, Fugate, and Rogers, 2018; Gutzwiller, Hunt, and Lange, 2016; Healey, 2017; Vieane et al., 2016a, 2017). My purpose in publishing the current version is simply to offload some of the detail and explanation behind situation awareness concept for cyber defense.

A final note of concern is worthy of attention here. Cybersecurity as a conceptual area has many different definitions and variations on a theme. It is often viewed a la carte, without a particular targeted user in mind (e.g., "cyber-attacks" can be targeted toward a single user, a particular service, a web domain, a company, a country, various infrastructures, etc.). So, it is important to begin decomposing the different places where humans comprise the system. As attempted here and elsewhere, I note that cyber defense relies on humans at many different levels of the networked systems we live within. End-user privacy and security, the kind that makes the news and encourages (rightly!) you to strengthen your passwords is important, but it is not the only aspect and the watch floor for cyber incident centers is just as critical. Similarly, analysts are rarely in the roles many of us associate with computers and security at the office or job site — that of the administrator, and therefore are not simply patching or installing software all day. These simple distinctions may seem trivial, but they are hugely important when attempting to work with analysts, for all the same reasons that you would not call all police employees "detectives" or all of those in the armed forces "Major." The nuances are emerging and the best source for understanding the complexity is in flux, though I can suggest the NIST documentation on roles and goals in cyberspace as a starting point.

Robert Gutzwiller, PhD

This page is intentionally blank.

REFERENCES

Aggarwal, P. 2017. “Modeling the Effects of Amount and Timing of Deception in Simulated Network Scenarios Modeling the Effects of Amount and Timing of Deception in Simulated Network Scenarios,” (May).

Aggarwal, P., C. Gonzalez, and V. Dutt. 2016. “Looking from the hacker’s perspective: Role of deceptive strategies in cyber security,” *2016 International Conference on Cyber Situational Awareness, Data Analytics and Assessment, CyberSA 2016*, (June). ” Available online at: http://doi.org/10.1109/CyberSA.2016.7503288

Alberts, C., A. Dorofee, G. Killcrece, R. Ruefle, and M. Zajicek. 2004. “Management Processes for CSIRTs: A Work in Progress,” *Carnegie Mellon Technical Report CMU/SEI-2004-TR-015*.

Altmann, E. M., and J. G. Trafton. 2002. “Memory for goals: an activation-based model. *Cognitive Science,” 26*(1)39–83. Available online at: http://doi.org/10.1207/s15516709cog2601_2

Barford, P., M. Dacier, T. G. Dietterich, M. Fredrikson, J. Giffin, S. Jajodia, J. … Yen, 2010. “Cyber SA: Situational awareness for cyber defense. In *Cyber Situational Awareness*,” (pp. 3–14).

Bass, T. 2000. “Intrusion detection systems and multisensor data fusion: Creating Cyberspace Situational Awareness. *Communications of the ACM*,” 43(4), 99–105.

Best, D. M., A. Endert, and D. Kidwell. 2014. “7 key challenges for visualization in cyber network defense. *Proceedings of the Eleventh Workshop on Visualization for Cyber Security*,” (pp.33–40). Available online at: http://doi.org/10.1145/2671491.2671497.

Borghetti, B., G. Funke, R. Pastel, and R. Gutzwiller. 2017. “Cyber Human Research from the Cyber Operator’s View. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*,” 61(1), (pp. 350–350). Available online at: http://doi.org/10.1177/1541931213601569.

Brickman, B., L. Hettinger, M. Roe, D. Stautberg, M. Vidulich, M. Haas, and R. Shaw. 1999. “The global implicit measurement of situation awareness: Implications for design and adaptive interface technologies.” In M. Scerbo and M. Mouloua (Eds.), *Automation Technology and Human Performance: Current Research and Trends* (pp. 60–64). Mahwah, NJ: Lawrence Erlbaum.

Brynielsson, J., U. Franke, and S.Varga. 2016. “Cyber Situational Awareness Testing.” In B. Akhgar and B. Brewster (Eds.), *Combatting Cybercrime and Cyberterrorism Challenges, Trends and Priorities* (Vol. 2, pp. 209–233). Springer International Publishing, Switzerland. Available online at: http://doi.org/10.1007/978-3-319-38930-1.

Cain, A., and D. Schuster. 2014. “Measurement of situation awareness among diverse agents in cyber security.” *IEEE International Inter-Disciplinary Conference on Cognitive Methods in Situation Awareness and Decision Support (CogSIMA)*, (pp. 124–129).

Champion, M., S. Jariwala, , P. Ward, , and N. J. Cooke. 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, 310–314. http://doi.org/10.1177/1541931214581064

Champion, M., P. Rajivan, N. J. Cooke, and S. Jariwala. 2012. “Team-based cyber defense analysis. *Cognitive Methods in Situation Awareness and Decision Support (CogSIMA)*,” (pp. 218–221).

Conti, G., Ahamad, M., and Stasko, J. 2005. “Attacking information visualization system usability overloading and deceiving the human.” *Proceedings of the 2005 Symposium on Usable Privacy and Security - SOUPS ’05, 89–100*. Available online at: https://doi.org/10.1145/1073001.1073010

Conti, G., J. Nelson, and D. Raymond. 2013. “Towards a cyber common operating picture.” In K. Podins, J. Stinissen, and M. Maybaum (Eds.), *International Conference on Cyber Conflict* (pp. 1–17). Tallinn: NATO CCD COE Publications.

D’Amico, A. D., J. R. Goodall, D. R. Tesone, and J. K. Kopylec. 2007. “Visual Discovery in computer network defense. IEEE Computer Graphics and Applications,” (October), (pp. 20–27).

D’Amico, A. D., and M. Kocka. 2005. “Information assurance visualizations for specific stages of situational awareness and intended uses: lessons learned. IEEE Workshop on Visualization for Computer Security (VizSEC),” (pp. 107–112). Available online at: http://doi.org/10.1109/VIZSEC.2005.1532072

D’Amico, A. D., and S. Salas. 2003. “Visualization as an aid for assessing the mission impact of information security breaches.” Proceedings of the DARPA Information Survivability Conference and Exposition, (2)190–195.

D’Amico, A. D., and K. Whitley. 2008. “The real work of computer network defense analysts: The analysis roles and processes that transform network data into security situation awareness.” In J. Goodall, G. Conti, and K. Ma (Eds.), Proceedings of the Workshop on Visualization for Computer Security. Springer Berlin Heidelberg.

D’Amico, A., K. Whitley, D. Tesone, B. O’Brien, and E. Roth. 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)229–233. Retrieved from http://pro.sagepub.com/content/49/3/229.shortDekker, S. W. A. 2012. “On the epistemology and ethics of communicating a Cartesian consciousness.” *Safety Science*. Available online at: http://doi.org/10.1016/j.ssci.2012.05.028

Dressler, J., C. L. Bowen, W. Moody, and J. Koepke. 2014. “Operational data classes for establishing situational awareness in cyberspace.” In P. Brangetto, M. Maybaum, and J. Stinissen (Eds.), *International Conference on Cyber Conflict* (pp. 175–186). Tallinn: NATO CCD COE Publications.

Drury, J. L., J. Scholtz, and H. Yanco, 2003. “Awareness in human-robot interactions.” Proceedings of the *IEEE International Conference on Systems, Man, and Cybernetics*, (pp. 111–119).

Durso, F. T., M. K. Bleckley, and A. R. Dattel, 2006. “Does Situation Awareness Add to the Validity of Cognitive Tests? Human Factors,” *48*(4)721–733.

Durso, F. T., C. A. Hackworth, T. R. Truitt, J. Crutchfield, D. Nikolic, and C. A. Manning, 1998. “Situation awareness as a predictor of performance for en route air traffic controllers. Air Traffic Control Quarterly,” *6*(1)1–20.

Durso, F. T., K., Rawson, and S. Girotto, 2007. “Comprehension and situation awareness.” In F. Durso, R. Nickerson, S. Dumais, S. Lewandowsky, and T. Perfect (Eds.), *Handbook of Applied Cognition* (pp. 163–193). West Sussex: John Wiley and Sons, LTD.

Dutt, V., Y.-S. Ahn, and C. Gonzalez, 2013. “Cyber situation awareness: Modeling detection of cyber attacks with instance-based learning theory.” *Human Factors*, 55(3), 605–618. Available online at: http://doi.org/10.1177/0018720812464045.

Endsley, M. 1993. “Situation awareness and workload- Flip sides of the same coin.” *Proceedings of the 7th International Symposium on Aviation Psychology*.

Endsley, M. R. 1995a. “Measurement of situation awareness in dynamic systems.” *Human Factors*, 37(1)65–84.

Endsley, M. R. 1995b. “Toward a theory of situation awareness in dynamic systems.” *Human Factors*, *37*(1)32–64.

Endsley, M. R. 2000. “Direct measurement of situation awareness: Validity and use of SAGAT.” In M. Endsley and D. Garland (Eds.), *Situation awareness: Analysis and measurement* (pp. 147–174). Mahwah, NJ: Lawrence Erlbaum Associates, Inc.

Endsley, M. R., and D. Jones, 2012. “*Designing for situation awareness: An approach to human-centered design (2nd ed.)*.” New York: CRC Press.

Endsley, M. R., and E. O. Kiris, 1995. “The out-of-the-loop performance problem and level of control in automation.” Human Factors, 37(2)381–394.

Erbacher, R. F., D. A. Frincke, P. C. Wong, S. Moody, and G. Fink, 2010. “Cognitive task analysis of network analysts and managers for network situational awareness.” *IS&T/SPIE Electronic Imaging, 75300H–12*. Available online at: http://doi.org/10.1117/12.845488

Ericsson, K., and W. Kintsch, 1995. “Long-term working memory.” *Psychological Review*, (102)211–245.

Fink, G. A., C. L. North, A. Endert, and S. Rose, 2009. “Visualizing cyber security: Usable workspaces.” *International Workshop on Visualization for Cyber Security*, 45–56.

Franke, U., and J. Brynielsson, 2014. “Cyber situational awareness – a systematic review of the literature. Computers and Security,” (46)18–31.

Gaba, D. M., and S. K. Howard, 1995. “Situation awareness in anesthesiology.” *Human Factors*, *37*(1)20–31.

Giacobe, N. A. 2013. “A picture is worth a thousand alerts.” *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, *57*(1)172–176. Available online at: http://doi.org/10.1177/1541931213571039

Giacobe, N. A., M. D. McNeese, V. F. Mancuso, and D. Minotra, (2013). “Capturing human cognition in cyber-security simulations with NETS.” *IEEE ISI Conference*, (pp. 284–288).

Goodall, J., W. Lutters, and A. Komlodi, 2004b. “The work of intrusion detection: Rethinking the role of security analysts.” AMCIS, 1421–1427.

Goodall, J. R. 2009. “Visualization is better! A comparative evaluation.” *International Workshop on Visualization for Cyber Security*, 57–68. Available online at: http://doi.org/10.1109/VIZSEC.2009.5375543

Goodall, J. R., W. Lutters, and A. Komlodi. 2004a. “I know my network: collaboration and expertise in intrusion detection.” *Proceedings of the ACM Conference on Computer Supported Cooperative Work*, (pp. 342–345).

Goodall, J. R., and M. Sowul. 2009. “VIAssist: Visual analytics for cyber defense.” *IEEE Conference on Technologies for Homeland Security*, (pp. 143–150).

Gronlund, S. D., M. R. P. Dougherty, D. D. Ohrt, G. L. Thomson, M. K. Bleckley, D. L. Bain, … C. A. Manning. 1997. “*The role of memory in air traffic control*.” Washington, DC.

Gugerty, L. J. 1997. “Situation awareness during driving: Explicit and implicit knowledge in dynamic spatial memory.*” Journal of Experimental Psychology: Applied*, *3*(1)42–66.

Gutzwiller, R. S., and B. A. Clegg. 2013. “The role of working memory in levels of situation awareness.” *Journal of Cognitive Engineering and Decision Making*, *7*(2)141–154.

Gutzwiller, R. S., K. Ferguson-walter, S. Fugate, and A. Rogers. 2018. “Oh, look, a butterfly!" A framework for distracting attackers to improve cyber defense.” *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 62(1), 272-276.

Gutzwiller, R. S., Fugate, S. Sawyer, B. D. and P. A. Hancock. 2015. “The human factors of cyber network defense.” *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, *59*(1)322–326.

Gutzwiller, R. S., S. M. Hunt, and D. S. Lange. 2016. “A Task Analysis toward Characterizing Cyber-Cognitive Situation Awareness (CCSA) in Cyber Defense Analysts.*” IEEE CogSIMA*, (pp. 14–20).

Healey, C. G. 2017. “Theory and Models for Cyber Situation Awareness, 10030 (October).” http://doi.org/10.1007/978-3-319-61152-5.

Holsopple, J., M. Sudit, M. Nusinov, D. Liu, H. Du, and S. Yang. 2010. “Enhancing situation awareness via automated situation assessment.” *IEEE Communications Magazine*, (March), (pp. 146–152).

Jones, D. 2000. Subjective measures of situation awareness. In M. Endsley and D. Garland (Eds.), “*Situation awareness: Analysis and measurement*.” (pp. 113–128). Mahwah, NJ: Lawrence Erlbaum Associates, Inc.

Jones, D., and M. Endsley. 1996. “Sources of situation awareness errors in aviation.” *Aviation, Space, and Environmental Medicine*, (67)507–512.

Killcrece, G., K. P. Kossakowski, R. Ruefle, and M. Zajicek. 2003. “State of the practice of computer security incident response teams (CSIRTS).” *Carnegie Mellon Technical Report CMU/SEI-2003-TR-001*.

Klein, G. 1997. “Developing Expertise in Decision Making. Thinking and Reasoning,” *3*(4)337–352.

Klein, G., Ruckert, C., Kleiber, M., Jahnke, M., & Toelle, J. 2010. “Towards a model-based cyber defense situational awareness visualization environment.” *Proceedings of the RTO Workshop “Visualising Networks: Coping with Chance and Uncertainty*,” 1–11.

Komlodi, A., J. R. Goodall, and W. Lutters, 2004. “An information visualization framework for intrusion detection.” CHI, 1743–1746.

Komlodi, A., P. Rheingans, U. Ayachit, J. R. Goodall, and A. Joshi, 2005. “A user-centered look at glyph-based security visualization workshop on visualization for computer security.” *Workshop on Visualization for Computer Security*, 21–28.

Kott, A., Buchler, N. and Schaefer, K. E. 2014. “Kinetic and cyber.” In A. Kott, C. Wang, and R. F. Erbacher (Eds.), *Cyber Defense and Situational Awareness* (Vol. 62, pp. 29–45). Cham: Springer International Publishing.

Lemay, A., and S. Leblanc. 2018. “Cognitive Biases in Cyber Decision-Making.” *Proceedings of the 13th International Conference on Cyber Warfare and Security*, (p. 395).

Levin, D., Tenney, Y. and H. Henri. 2001. “Issues in human interaction for cyber command and control.” *DARPA Information Survivability Conference*, (1)141–151.

Li, J., Ou, X. and R. Rajagopalan. 2009. “Uncertainty and Risk Management in Cyber Situational Awareness Abstract.” *ARO Workshop on Cyber Situational Awareness*.

Lützhöft, M. H., and S. W. A. Dekker. 2002. “On Your Watch: Automation on the Bridge.” *Journal of Navigation*, 55(01), 83–96. Available online at: http://doi.org/10.1017/S0373463301001588

Ma, R., and D. B. Kaber. 2005. “Situation awareness and workload in driving while using adaptive cruise control and a cell phone.” *International Journal of Industrial Ergonomics*, *35*(10)939–953.

Mahoney, S., E. Roth, K. Steinke, J. Pfautz, C. Wu, and M. Farry. 2010. “A cognitive task analysis for cyber situational awareness.” *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, (54)279–283.

Malviya, A., G. A. Fink, L. Sego, and B. Endicott-Popovsky. 2011. “Situational awareness as a measure of performance in cyber security collaborative work.” *IEEE International Conference on Information Technology: New Generations*, (pp. 937–942). Available online at: http://doi.org/10.1109/ITNG.2011.161.

Mancuso, V. F., and M. D. McNeese. 2012. “Effects of Integrated and Differentiated Team Knowledge Structures on Distributed Team Cognition.” *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, *56*(1)388–392. Available online at: http://doi.org/10.1177/1071181312561088.

Michel, M. C. K., N. P. Helmick, and L. M. Mayron. 2011. “Cognitive cyber situational awareness using virtual worlds.” *IEEE International Inter-Disciplinary Conference on Cognitive Methods in Situation Awareness and Decision Support (CogSIMA)*, (pp. 179–182).

Mogford, R. H. 1997. “Mental models and situation awareness in air traffic control.” *The International Journal of Aviation Psychology*, *7*(4)331–341.

Neisser, U. 1976. “*Cognition and reality: Principles and implications of cognitive psychology*.” San Francisco, CA: W.H Freeman.

O’Regan, J. K. 1992. “Solving the “real” mysteries of visual perception: the world as an outside memory.” *Canadian Journal of Psychology*, *46*(3)461–488.

Okolica, J., J. McDonald, G. Peterson, R. F. Mills, and M. Haas. 2009. “Developing systems for cyber situational awareness.” In J. Gourd, V. V Phoha, and S. S. Iyengar (Eds.), *Proceedings of the Cyberspace Research Workshop* (pp. 46–56). Shreveport, LA.

Parasuraman, R., T. Sheridan, and C. D. Wickens. 2008. “Situation awareness, mental workload, and trust in automation: Viable empirically supported cognitive engineering constructs.” *Journal of Cognitive Engineering and Decision Making*, 2(2)140–160.

Paul, C. 2014. “Human-centered study of a network operations center: Experience report and lessons learned.” *Proceedings of the ACM Workshop on Security Information Workers*, (pp. 39–42).

Paul, C. L., Rohrer, R. Sponaugle, P. Huston, J. and Nebesh, B. 2013. “CyberSAVI : A Cyber Situation Awareness Visual Interface for Mission-Level Network Situation Awareness.” *Proceedings of the Workshop on Visualization for Computer Security*, (pp. 1–2).

Paul, C., and K. Whitley. 2013. “A taxonomy of cyber awareness questions for the user-centered design of cyber situation awareness.” In L. Marinos and I. Askoxylakis (Eds.), *Lecture Notes on Computer Science: HAS/HCII* 2013 (pp. 145–154). Springer-Verlag Berlin Heidelberg.

Pfleeger, S., and D. Caputo. 2012. “Leveraging behavioral science to mitigate cyber security risk.” *Computers and Security*, *31*(4)597–611.

Pierce, R. S. 2012. “The Effect of SPAM Administration During a Dynamic Simulation.” *Human Factors*, *54*(5)838–848.

Pritchett, A., and R. Hansman. 2000. “Use of testable responses for performance-based measurement of situation awareness.” In M. Endsley and D. Garland (Eds.), *Situation awareness: Analysis and Measurement* (pp. 189–209). Mahwah, NJ: Lawrence Erlbaum.

Reed, T., R. G. Abbott, B. Anderson, K. Nauer, and C. Forsythe. 2014. “Simulation of workflow and threat characteristics for cyber security incident response teams.” *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, *58*(1)427–431.

Shiravi, H., Shiravi, A. and A. A. Ghorbani. 2012. “A survey of visualization systems for network security. *IEEE Transactions on Visualization and Computer Graphics*,” *18*(8)1313–1329.

Shope, S. 2013. “Effective cyber situation awareness (CSA) assessment and training final report.” *US Army Final Report #W911NF-13-C-0060*, 1–99.

Smith, K., and P. A. Hancock. 1995. “Situation awareness is adaptive, externally directed consciousness.” *Human Factors*,” *37*(1)137–148.

Sohn, Y. W., and S. M. Doane. 2004. “Memory processes of flight situation awareness: Interactive roles of working memory capacity, long-term working memory, and expertise.” *Human Factors*, *46*(3)461–475.

Staheli, D., Yu, T. R. Crouser, S. Damodaran, K. Nam, D. O’Gwynn, … L. Harrison. 2014. “Visualization evaluation for cyber security: trends and future directions.” *Proceedings of the 11th Workshop on Visualization for Cyber Security*, (pp. 49–56).

Stanton, N. A. 2016. “Distributed situation awareness. Theoretical Issues in Ergonomics Science,” *17*(1)1–7. Available online at: http://doi.org/10.1080/1463922X.2015.1106615

Stanton, N. A., P. M. Salmon, and G. H. Walker. 2015. “Let the Reader Decide: A Paradigm Shift for Situation Awareness in Sociotechnical Systems.” *Journal of Cognitive Engineering and Decision Making*, *9*(1)44–50.

Stanton, N. A., R. Stewart, D. Harris, R. J. Houghton, C. Baber, McMaster, R. … D. Green. 2006. “Distributed situation awareness in dynamic systems: theoretical development and application of an ergonomics methodology.” *Ergonomics*, *49*(12–13)1288–1311.

Sulistyawati, K., C. D. Wickens, and Y. P. Chui. 2009. “Exploring the concept of team situation awareness in a simulated air combat environment.” *Journal of Cognitive Engineering and Decision Making*, *3*(4)309–330.

Sulistyawati, K., C. D. Wickens, and Y. P. Chui. 2011. “Prediction in situation awareness: Confidence bias and underlying cognitive abilities.” *The International Journal of Aviation Psychology*, *21*(2)153–174.

Thompson, R., E. Rantanen, and W. Yurcik. 2006. “Network intrusion detection cognitive task analysis: Textual and visual tool usage and recommendations.” *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, (50)669–673.

Thompson, R., E. Rantanen, W. Yurcik, and Bailey, B. 2007. “Command line or pretty lines?: comparing textual and visual interfaces for intrusion detection.” *Proceedings of the ACM CHI*.

Tyworth, M., N. A. Giacobe, and V. Mancuso, F. 2012. “Cyber situation awareness as distributed socio-cognitive work.” Cyber Sensing - *Proceedings of SPIE, 8404*. Available online at: http://doi.org/10.1117/12.919338

Tyworth, M., N. A. Giacobe, V. F. Mancuso, and C. Dancy, 2012. “The distributed nature of cyber situation awareness.” *IEEE International Inter-Disciplinary Conference on Cognitive Methods in Situation Awareness and Decision Support (CogSIMA)*, (pp. 174–178).

Vieane, A., G. Funke, E. Greenlee, V. Mancuso, B. Borghetti, B. Miller, … D. Boehm-Davis. 2017. “Task interruptions undermine cyber defense.” *Proceedings of the Human Factors and Ergonomics Society*, (61), 375–379.

Vieane, A., G. Funke, V. Mancuso, E. Greenlee, G. Dye, B. Borghetti, … R. Brown. 2016a. “Coordinated displays to assist cyber defenders.” *Proceedings of the Human Factors and Ergonomics Society*,” (60), 344–348.

Vieane, A. Z., G. J. Funke, R. S. Gutzwiller, V. F. Mancuso, B. D. Sawyer, and C. D. Wickens. 2016b. Addressing human factors gaps in cyber defense. *Proceedings of the Human Factors and Ergonomics Society*,” (60)770–773.

Williams, F. C. B., W. J. Faithfull, and J. C. Roberts, 2012. “SitaVis - Interactive Situation Awareness Visualization of large datasets.” *IEEE Symposium on Visual Analytics Science and Technology*,” (5)273–274.

Yanco, H. A., and J. L. Drury. 2004. “Where am I?” acquiring situation awareness using a remote robot platform.” *IEEE International Conference on Systems, Man and Cybernetics*, (3)2835–2840.

Yurcik W., J. Barlow, and J. Rosendale. 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.

This page is intentionally blank.

INITIAL DISTRIBUTION

84300	Library	(1)
85300	Archive/Stock	(1)
71000	R. Gutzwiller	(1)

Defense Technical Information Center
Fort Belvoir, VA 22060–6218 (1)

This page is intentionally blank.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-01-0188

PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

1. REPORT DATE *(DD-MM-YYYY)*	2. REPORT TYPE	3. DATES COVERED *(From - To)*
June 2019	Final	

4. TITLE AND SUBTITLE

Situation Awareness in Defensive Cyberspace Operations:
An Annotated Bibliographic Assessment Through 2015

5a. CONTRACT NUMBER

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

6. AUTHORS

Robert Gutzwiller

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

NIWC Pacific
53560 Hull Street
San Diego, CA 92152–5001

8. PERFORMING ORGANIZATION REPORT NUMBER

TR-3184

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

Office of Naval Research Naval Innovative Science and Engineering
53560 Hull Street
San Diego, CA 92152–5001

10. SPONSOR/MONITOR'S ACRONYM(S)

NISE

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT

DISTRIBUTION STATEMENT A: Approved for public release.

13. SUPPLEMENTARY NOTES

14. ABSTRACT

Situation awareness (SA) is a buzzword concept, but unlike most buzzwords, it has a robust and scientific research focus. Grounded in cognitive psychology and human factors research findings, SA is essentially characterization of what a person knows about their current and future environment, in the context of their current goals. Being aware of critical information, comprehending it, and even projecting the situation into the near future is a highly useful skill in many different domains, such as aviation, driving, and healthcare.

Cyberspace operators share many of the same concerns as more traditional roles of the pilot and the driver. In particular, cyber network defense (CND) operators must remain aware of different types of activity, comprehend what various sources of information mean and how they are pieced together, and project what effects on their network will be in order to stop or prevent threats.

For future Warfighter performance in cyberspace, substantial benefits are derived from improving defender SA, as we have in other domains. Improving SA is permanently connected with understanding how to measure awareness in the cyber environment. In support of this concept, a review of the available literature in cyber SA was conducted to determine how to begin the process of measurement and improvement, and to derive key points for further research.

This report finds that the utility of SA analysis and measurement has yet to be realized in cyberspace. With a few exceptions, almost no experimental work was found on measuring or characterizing the process or product in developing cyber situational awareness. In those few found, there are still occasional methodological or analytical deficits and concerns, which preclude any strong conclusions. The development of specific measurement techniques must be explored elsewhere. It was additionally found that almost no research was found measuring SA in the CND environment, many reports claim that a new or unique interface could improve it. The willingness and rapacity of these claims in the literature mirrors claims by industrial software solutions, which many cyber professionals abhor and abandon in favor of Excel spreadsheets and command line interfaces. Together this suggests that demand is present for SA improvement, but capability has not risen to match it.

15. SUBJECT TERMS

Cyber Situation Awareness; Endsley's Situation Awareness; subjective measures; cyber network defense;

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			Karl Van Orden
U	U	U	U	58	19B. TELEPHONE NUMBER *(Include area code)* 619-553-8015

This page is intentionally blank.

This page is intentionally blank.

NIWC Pacific
San Diego, CA 92152-5001