EAI Endorsed Transactions
on Security and Safety

Research Article

A human-in-the-loop approach to understanding situation awareness in cyber defence analysis

Michael Tyworth[1,] *, Nicklaus A. Giacobe[1], Vincent F. Mancuso[1], Michael D. McNeese[1], David L. Hall[1]

[1]College of Information Sciences & Technology, The Pennsylvania State University, University Park, PA 16802

Abstract

In this paper we argue for a human-in-the-loop approach to the study of situation awareness in computer defence analysis (CDA). The cognitive phenomenon of situation awareness (SA) has received significant attention in cybersecurity/CDA research. Yet little of this work has attended to the cognitive aspects of situation awareness in the CDA context; instead, the human operator has been treated as an abstraction within the larger human-technology system. A more human-centric approach that seeks to understand the socio-cognitive work of human operators as they perform CDA will yield greater insights into the design of tools and interfaces for CDA. As support for this argument, we present our own work employing the Living Lab Framework through which we ground our experimental findings in contextual knowledge of real-world practice.

Keywords: computer defence analysis, cybersecurity, human-in-the-loop, living lab framework, situation awareness

Received on 31 March 2012; accepted on 31 March 2013, published on 03 May 2013

doi: 10.4108/trans.sesa.01-06.2013.e6

1. Introduction

Situation awareness (SA) has received widespread attention in research on cyber defence analysis. In particular, Endsley's (1995a, 1995b) '3 Level' model of SA has received the most attention among cyber-security scholars; primarily serving as a conceptual basis for measuring system performance. Indeed, much of the research on SA in cybersecurity, or cyber-SA, has primarily taken an algorithmic perspective (Orlikowski & Iacono, 2001); focusing primarily on the automation and the development of new defensive tools for protection, detection and response (McMillan & Tyworth, 2012). Examples of this work include data visualizations (D'Amico, A. & Larkin, 2001), data fusion methods for tracking cyber-attacks (Stotz & Sudit, 2007; Yang, Shanchieh J., Stotz, Holsopple, Sudit, & Kuhl, 2009), identification of internal and external threats using intelligent agents (Buford, Lewis, & Jakobson, 2008; Yen et al., 2010), and the use of probabilistic models to assess network vulnerability (Peng, Li, Xinming, Peng, & Levy, 2010; Tadda, G., Salerno, Boulware, Hinman, & Gorton, 2006). Although valuable, this body of work overlooks perhaps the most crucial component of cyber defence analysis: *the human component* (Boyce et al., 2011; Goodall, Lutters, & Komlodi, 2004). Indeed, much of the research on cyber-SA has paid little attention to how operators perform with existing technologies let alone whether or not these new technologies actually improve SA in human operators.

In this paper we argue for a human-in-the-loop perspective on cyber-SA; thus shifting analytical attention away from the development of new technologies towards the socio-cognitive work of human cyber security professionals. We believe that such a shift is critically necessary because human operators are central and critical to any cyber defence system and yet our understanding of how human operators cognitively operate in the unique environment of cyberspace remains poorly understood.

Our argument proceeds as follows. We begin with a discussion of situation awareness theory and its application to the cyber defence domain. Here we draw attention to the strengths and weaknesses of the different

* Corresponding author e-mail: mjt241@smeal.psu.edu

theoretical perspectives on SA as applied to cyber defence analysis. We then proceed to a discussion of the different methodological approaches to the study of cyber-SA with particular attention given to the application of the Living Lab Framework (LLF) to the study of cyber defence analysis. Finally, we conclude the paper with a discussion of our on-going research of cyber defence analysis using LLF. In this section we present our findings from our field research and show how those findings are informing our on-going laboratory-based experiments using scaled-world simulations.

2. Situation Awareness, Cyber Defense Analysis & Cyber-SA

An area of cognitive science that has received significant attention in CDA research is situation awareness. In this section we provide an overview of SA theory and its application to the domain of cyber defence analysis. In doing so we make two main arguments: (1) that extant research on cyber-SA has largely overlooked the human operator; and (2) that a distributed theory of SA may be more ideally suited to the study of cyber-SA than the predominantly used cognitive perspective.

2.1 Situation Awareness

Situation awareness is, in lay terms, the knowing of what is going on around you. In cognitive terms, one's SA is the degree to which one perceives task-salient cues in the environment, correctly understands the meaning of those cues, and is able to take the correct action to achieve a specific future state (Endsley, 1995b). The greater an individual has SA, the more likely the individual is to take the appropriate action; conversely, a lack of SA is a key factor in the commitment of errors by human operators (Endsley, 2000). There are numerous high-profile examples of a lack of SA contributing to catastrophic failures including the crash of Air France Flight 447 as a result of pilots incorrectly reacting to faulty sensor data (Wise, 2012), and the partial meltdown of Reactor 2 at Three Mile Island as a result of plant operators failing to correctly understand the operating state of the reactor tower's cooling system (Perrow, 1999). Scholars have studied situation awareness in numerous domains such as aircraft piloting (Endsley, 1993; Endsley, Farley, Jones, Midkiff, & Hansman, 1998), air traffic control (Endsley & Rodgers, 1994), ship navigation (Lee & Sanquist, 2000), emergency response (Blandford & William Wong, 2004; McGrath & McGrath, 2005), C4i systems (French & Hutchinson, 2002; Salmon, P., Stanton, Walker, & Green, 2006), and surgical teams (Bardram, Hansen, & Soegard, 2006; Hazlehurst, McMullen, & Gorman, 2007).

Table 1 Theoretical perspectives of SA

	Location of SA	Analytical Focus	Treatment of Technology
Cognitive Perspective	Within the individual	Human operator's SA at a point in time	Technology mediates but does not have SA
Technological Perspective	Within the technological artefact	Communication of artefact's SA via information presentation	Technology has its own SA
Distributed Perspective	Distributed across human and technological agents	Interactions between agents and between agents and environment	Both humans and artefacts have SA that may or may not be shared

Stanton, Salmon, Walker & Jenkins (2009) provide an excellent discussion of extant SA theory. They identify three main theoretical perspectives in the literature: (1) the *cognitive perspective*; (2) the *technological (or engineering) perspective*, and (3) the *distributed perspective*.

The *cognitive perspective*, rooted in Endsley's (1995b, 2000) information-processing model of SA, is the most widely adopted theoretical perspective in SA research. Known as the "3 Level" model, the *cognitive perspective* sees SA as a human internal cognitive state comprised of perception (Level 1), comprehension (Level 2), and projection (Level 3). The analytical focus of the *cognitive perspective* is on the human operator's understanding of the environment at a particular point in time which is assessed using freeze-probe measurement techniques such as the Situation Awareness Global Assessment Technique (SAGAT) (Endsley, 1995a). One strength of the *cognitive perspective* is that it lends itself to the quantitative measurement of SA: for example we can compare a an individual's ability to detect salient cues in the environment against a known ground truth (Salmon, P. M. et al., 2008). A criticism of the *cognitive perspective* is that it does not scale well to other levels of analysis. For example, Stanton *et al.* (2009) argue that team SA is more than just the sum of individual team members' SA as those who have adopted the *cognitive perspective* have suggested. Additionally, the *cognitive perspective* is not well suited to capturing the ways that artefacts offload elements of SA because it assumes SA as a strictly psychological phenomenon.

The *technological perspective* argues that SA is instantiated in the presentation of information by a technological artefact. Implicit in this view is that SA resides within the artefact itself, typically in the form of information (Stanton, Neville A. et al., 2009). For example, the display of route information, travel time, traffic conditions, weather and fuel efficiency in a GPS navigational appliance is considered to be a display of the

SA (information) held by the device. The analytical focus of the *technological perspective* is the design and configuration of information presentations to most effectively convey the SA contained within the device. The utility of the *technological perspective* is that it accounts for cognitive artefacts and how best to display the information they contain. A flaw in the *technological perspective* is the assumption that information in itself is situation awareness, and that by providing access to that information, human SA is necessarily improved. As the Air France disaster makes abundantly clear, even when technological artefacts are conveying their 'awareness' human situation awareness may not improve and may even deteriorate (Wise, 2012).

The *distributed perspective* of SA is a hybrid theory that posits that SA resides in both human and technological agents; distributed throughout a socio-technical system (Stanton, N. A. et al., 2006). Developed relatively recently, distributed situation awareness (DSA) theory seeks to integrate SA with distributed cognition. Within the broader system, different agents may have different SA, and the degree to which agents within the system share SA is a function of the extent to which their goals overlap. The analytical focus of the *distributed perspective* is the socio-cognitive system and interactions between agents within the system and the system, and agents and the environment. A strength of the *distributed perspective* is that it accounts for both human and technological SA; potentially providing a more valid description of the ways in which SA occurs in task environments in which technology is central to the task. Additionally, by moving away from the view that team SA is the aggregate SA of all individual team members, the *distributed perspective* likely gives a more accurate picture of distributed collaborative work. It is not clear, however, how to measure SA distributed across a socio-technical system. Gorman, Cooke and Winner (Gorman, Cooke, & Winner, 2006) have argued for measuring the extent to team members' actions, comments, behaviours and interactions are coordinated. It has also been suggested that performance is a useful proxy measure of DSA on the assumption that greater SA results in greater performance (Salmon, P. et al., 2006; Stanton, N. A. et al., 2006; Walker et al., 2006).

Each theoretical perspective has its strengths and weaknesses, particularly in their application to the study cyber-situation awareness. As we will show below, a review of the extant literature reveals that the *cognitive perspective* is the predominant perspective espoused in cyber-SA research. What will also be shown is that scholars have, in practice, used elements of the *cognitive perspective* to measure the 'SA' of the technological artefact.

2.2 Cyber Defence Analysis & Cyber-SA

A review of the extant literature on situation awareness in CDA reveals that though cyber defence scholars have relied on the cognitive perspective as a basis for their research almost exclusively, the research is, in practice employing the technological perspective of SA (Mathew, Shah, & Upadhyaya, 2005; McMillan & Tyworth, 2012; Okolica, McDonald, Peterson, Mills, & Haas, 2009; Yang, S. J., Byers, Holsopple, Argauer, & Fava, 2008; Yen et al., 2010). The notable exception is Tadda and Salerno's (2010) Situation Awareness Reference Model. The extent to which cyber-SA has been considered beyond the 3-Level model has typically been limited to discussions of what phenomena an analyst or system must perceive (c.f., Barford et al., 2010) or factors that make developing cyber-SA difficult (c.f., Yang, S. J. et al., 2008).

We identify two primary reasons for why SA has been applied to CDA research in this way. One, the levels of the Endsley model of SA closely align to the levels of the JDL Data Fusion Process Model that informs much of the CDA research. The Joint Directors of Laboratories (JDL) Data Fusion Process Model is a model that describes how disparate pieces of data detected by multiple sensors are fused into a coherent picture. In the JDL model, there are four levels of data fusion (Llinas & Hall, 1998). The first three, consisting of Object Refinement (Level 1), Situation Refinement (Level 2), and Threat Refinement (Level 3), are conceptually similar to the three levels of SA respectively (perception, comprehension, and detection). This close conceptual alignment provides a natural point on which to connect the work on data fusion that comprises much of the CDA research to the desired state of situation awareness.

A second reason for why CDA scholars have applied a limited form of SA theory to their research is that their research really is not about SA at all, but about the development of new technologies, techniques, and representations of information. The implicit assumption of this stream of research is that improved sensor/intelligent agent performance or improved presentation of cybersecurity related information will produce improved SA in human operators.

Although valuable, this research is limited in how much it informs our understanding of SA in computer defence analysis for three reasons. First, this stream of research tends to focus on intrusion detection exclusively when computer defence in practice consists of a variety of activities including policy, forensics, remediation, and administration. Second, focusing on technology development does not tell us if and how the cognitive process of forming SA in cyber defence analysis may or may not be different as a result of the unique environmental properties of working in cyberspace. Finally, this stream of research treats the human operator as an abstraction. Little to no attention is given to how cyber defence analysts actually work; make use of

information, and whether the tools they are developing actually improve their situation awareness.

3. A Human-Centric Approach to Cyber Defence Research

In order to better understand how computer defence analysts develop situation awareness it is both important to understand the cognitive processes of cyber-SA *and* the socio-cognitive work of CDA. To do so, we have adopted the Living Laboratory Framework to guide our data collection and analysis. The Living Lab Framework is detailed below; followed by a discussion of findings from our initial round of field work and a description of our on-going experiments using scaled-world simulations.

3.1 The Living Laboratory Framework

We employ the Living Laboratory Framework (LLF) to guide our data collection and analysis. As an ecological approach to the study of cognitive work in multi-operator environments that combines both quantitative and qualitative analysis (McNeese, 1996), the LLF is well suited to the study of the socio-cognitive work of CDA for four reasons. First, the analytical focus of the LLF is on cognitive work by human agents, precisely what we are focusing on in this research. Second, the underlying premise of the framework is cognitive work is situated, jointly determined by agents and environments, and is often 'goal-directed, self-organizing, and intentional (McNeese, Perusich, & Rentsch, 2000).' This premise is an accurate description of CDA in practice. Third, the LLF provides a process for the kind of empirically based theory development of the type we seek to accomplish with this research. Finally, the multi-method structure of the LLF facilitates both methodological and participant triangulation thereby increasing the validity of findings.

There are four data collection components to the LLF: (1) ethnographic observation, (2) knowledge elicitation using established cognitive engineering methodologies, (3) scaled-world simulation, and (4) implementation and evaluation. The components are performed in a mutually-informing, iterative loop.

Ethnographic observation is used to capture cognition and collaboration 'in the wild (McNeese, 1996).' Investigators observe users of systems in their actual work settings in an effort to understand the critical ways in which context impacts the interpretation of work and information (McNeese et al., 2000). Insights gained during ethnographic observation are then used to guide more structured cognitive fieldwork and the development of high-fidelity scenarios for use in the scaled-world simulator.

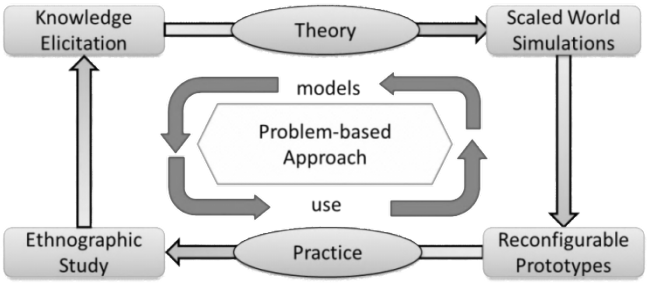

Figure 1 The Living Lab Framework (McNeese, 1996)

Knowledge elicitation – or cognitive fieldwork – supports the modelling and testing of both theory and technology (MacEachren, Cai, McNeese, Sharma, & Fuhrmann, 2006). Collection of contextual data is done using established cognitive field research tools such as concept mapping and cognitive task analyses. Data collected in the field is used to inform both theory development and the development of realistic scenarios for use in scaled-world simulations where participants engage in ambiguous, incompletely understood situations of the type we find in cybersecurity (Tyworth, Giacobe, Mancuso, Dancy, & McMillan, 2012). Findings from experiments using the simulations serve as a basis for the development or modification of information and communication technology (ICT) prototypes, which can then be reintroduced into the field for evaluation.

Scaled-world simulations are a means to quantitatively assess field-based observations in a controlled setting. Because many of the other portions of the Living Lab Framework rely on qualitative methods, it can be difficult to hone in on specific constructs that may be of interest. It is very difficult to use valid field measurements to extrapolate an understanding of what is going on, without interrupting the operational workflow. Moreover, during interviews, one is forced to rely on retrospective accounts and the collection of data in a non-controllable environment in which there are numerous potential confounds that can occur simultaneously. With this in mind, using scaled-world simulations within the living laboratory framework allows researchers to mimic the real environment and control for particular things that they are interested on. Additionally, scaled-world experiments are typically conducted in a laboratory setting, which allow for richer, quantitative, data to be collected via observations, performance measures, surveys and interviews.

Our work to date has primarily been in ethnographic observation, knowledge elicitation, and scaled-world simulation. We have done some initial prototype development in the form of a visual analytic interface for CDA.

Ethnographic Data Collection & Knowledge Elicitation

Our ethnographic work took the form of cognitive ethnography in that it was focused on specific activities, purposive, and verifiable through triangulation of observers, data sources, and methodologies. We conducted ethnographic observations in three contexts: a student team competing in a cyber-defence war game, professional analysts working in a large corporation, and professional analysts working in government. We observed CDA work being done in the form of intrusion detection, forensics, and system administration. During our observation sessions we focused our attention on four areas: (1) development / lack of SA among individuals or the team; (2) collaborative activity among analysts; (3) the use of formal or informal analytical methods; and (4) cognitive breakdowns.

Typically our ethnographic observations were conducted by members of the research team working in pairs. Because of the sensitive nature of CDA work, video and audio recording of what we observed was prohibited. Instead, observers took detailed notes of their observations which were then compared by the investigators during debriefing. Where discrepancies occurred, the observer team discussed the discrepancy until agreement was reached or the discrepancy was discarded. Observations were then coded and categorized using *nVivo* qualitative analysis software.

The insights we gained from the ethnographic work were used to structure our knowledge elicitation activities. Knowledge elicitation consisted of semi-structured interviews of CDA subject matter experts (SMEs) in industry, government, education, and the military. We chose semi-structured interviews as our method of knowledge elicitation because the semi-structured interview format gave us enough structure to facilitate directed inquiry while allowing us the flexibility to explore emergent topics that were interesting (Spradley, 1979). Each semi-structured interview was approximately one-hour in duration and consisted of 20 to 25 predesigned questions. The predesigned questions were designed to elicit data about four areas of inquiry: routine work activity; cognitive processes associated with the development of cyber-SA; data, information, and information-processing tools used by analysts; and the influence of organizational variables such as policy, culture, and work environment on the development of cyber-SA.

Table 2 Breakdown of Interviews

Domain	**# of Interviews**
Military	14
Government	4
Education	5

As with our ethnographic observations, the sensitive nature of CDA work prohibited us from making visual or audio recordings of the interview session. Investigators instead worked in pairs or groups of three, took detailed notes during the interview, and then debriefed after each interview. Interviews were transcribed, checked for inter-observer reliability, and then coded using *nVivo* qualitative research software using both *a priori* codes and codes generated from the data. For example, our *a priori codes* included codes related to SA (e.g., perception, comprehension, and recognition), work (information flows, collaboration, breakdowns, tasks), and social structures (e.g., policy, norms, values).

Our ethnographic and knowledge elicitation data revealed two findings that inform our on-going experimental work and prototype development. First, cyber-SA is distributed across both human operators and ICT artefacts in a complex socio-technical system spanning multiple operational domains. The domains we identified were intrusion detection, policy, operations or administration, and strategic analysis. Though an organization's cybersecurity posture is an aggregate of all these domains, operators often have limited awareness of the environmental state of other domains.

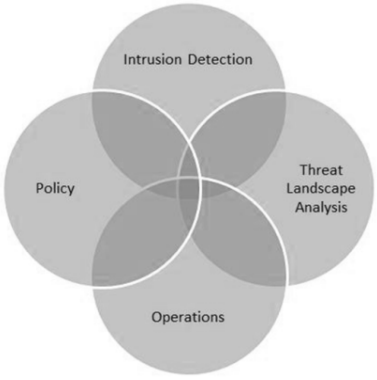

Figure 2 Domains of CDA

For example, individuals working in the policy and operations domains are often unaware of the nature and volume of malicious activity directed towards their network because it is caught and disposed of by the network intrusion analysts without comment. This can be problematic for system administrators and in policymakers in particular, because they have incomplete understanding of the severity of the malicious activity to which their assets are exposed. This dynamic is

analogous to a battlefield commander not being aware that the enemy is probing his line because the forward operating base did not bother reporting the squad of enemy they engaged and destroyed/repulsed.

In talking to intrusion detection system (IDS analysts, it became clear that the reason this dynamic occurs is because their sole focus is on identifying and blocking malicious activity. IDS analysts are not concerned about the larger threat landscape, and, in practice, the cognitive effort it takes to do that work effectively prohibits them from being able to, even if they wanted to. As one subject noted, this can be a problem because traffic may get missed for lack of a broader perspective.

The inverse of this dynamic also occurs. Individuals working in the operations and threat landscape analysis domains often do not provide mission-salient information to IDS analysts resulting in wasted effort and limited understanding. For example, it is not uncommon for an IDS analyst to diagnose traffic as suspicious only to find out upon contacting the monitored network administrators, that no, that traffic is actually approved. Similarly, a comment that repeatedly came up during interviews was that operators working in threat analysis domain would regularly ask for information on traffic from a particular device, or directed at a particular device, without providing further contextual information as to why. As a result, the IDS specialists were monitoring traffic without knowing why it was important or worthy of special attention.

The second finding is that boundaries in both physical and virtual form impair the development of system-wide SA. The boundaries separating the functional domains are opaque; and task-salient information is only able to only partially pass through the boundaries. As a result, individuals' domain-specific cyber-SA is degraded as a result of lacking key information or knowledge from other domains. The operational domains are, in practice, heterogeneous communities of practice with their own processes, and operational languages. This finding suggests that the CDA technologies that are likely to be most effective at facilitating cyber-SA are those that function as effective boundary objects.

The boundary object is a concept that comes from the sociology work that has received widespread attention in the study of collaborative work. A boundary object is simultaneously understood by multiple communities of practice and uniquely understood by individual communities of practice (Star, 1989). Take, for example, the system log of a workstation computer. A system administrator, forensic analyst, and a policymaker all understand that the log is a file containing a record of activity on the machine, but each is capable of understanding the log in ways unique to their community of practice. A system administrator may see that the user rights were not properly secured, the forensic analyst may see the point where malware was installed on the system, and the policymaker may see where the organization is exposed to liability as a result of a data breach. Our informants repeatedly identified an interface that would function in this boundary-spanning role as something that would immediately enhance their ability to establish and maintain their understanding of what was happening in their cyber-environment.

On-going Experimental Work

We are currently conducting experiments using scaled-world simulations. One set of experiments examines transactive memory and CDA. To conduct these experiments we have updated NeoCITIES scaled-world simulation (c.f., Jones, McNeese, Connors, Jefferson, & Hall, 2004; McNeese et al., 2005) to better support the dynamic and rich nature of the cyber security environment. The new simulation, the NeoCITIES Experimental Task Simulation (NETS), has been extended to support richer scenarios and complex decision making. The current implementation of NETS (referred to as idsNETS) has been implemented using intrusion detection data to mimic the role of an intrusion detection analyst. We have plans to extend the NETS functionality to be able to simulate scenarios from the other operational domains we identify in the future.

For our own research, we are addressing the issue of the formation and maintenance of transactive memory systems in synchronous distributed collaborations. To study this, a new version of the *NETS* simulation was designed (*teamNETS*) to simulate collaborative problem solving tasks within a cyber-environment. This version of the simulation was extended with numerous enhancements to better support our research questions and transactive memory research at large. Within the study, each team member is assigned a particular specialty, and in order to achieve high performance, it is necessary that they communicate and share relevant information to solve different types of events. From this study we hope to gain an understanding of how these transactive memory systems are formed in distributed collaborations, and how new systems can be designed to better support this process.

Transactive Memory was first conceptualized by Wegner (1985) as an "interpersonal awareness of others' knowledge" and can be conceptualized as a specialized form of Cyber Situation Awareness, where rather than focusing on, or being aware of, aspects within the cyber environment, your awareness is grounded in the cyber knowledge, activities and behaviours of your collaborators. An effective Transactive Memory System can give a human quick and coordinated access to another person's specialized expertise (Lewis, 2004). Numerous studies have shown a positive link between a team's Transactive Memory System and its performance in collaborative tasks (c.f., Ellis, 2006; Moreland & Myaskovsky, 2000; Pearsall & Ellis, 2006).

Whereas Transactive Memory is an important thread within team research it is mainly approached from a management or organization psychology lens, often only considering the humans. Since its inception, technology and information have evolved dramatically, though

Transactive Memory has remained fairly constant. Research has focused primarily on exploring its effect in new domains, and extending the concept as a research tool, but no one has examined how new technologies have changed how we, as humans use this transactive memory. In order to bring Transactive Memory into the 21st century, it is imperative that we understand how transactive memory has changed with synchronous distributed collaboration systems, social networks, and crowd-sourced knowledge repositories, to name a few.

A second set of experiments is being conducted to look at the impact of task load on the ability of participants to establish and maintain cyber-SA and prioritize tasks. Maintaining cyber SA is, in part, dependent on the ability to prioritize attention. Cyber defence analysts must attend to alerts associated to potential threats and respond to them within time constraints, requiring a prioritization of events in accordance to their threat level. However, high levels of cognitive workload may limit the ability of analysts to focus their attention on priority tasks. For example, unexpected surges in threat level in some events may not get noticed in time. An interface that provides information on anticipated threat level could facilitate analysts' ability to attend to unexpected surges.

In this set of experiments we explore the effect of a workload-preview on performance in a dual-task cyber-security event monitoring context using our NETS-DART scaled-world simulation. The simulation provides a dual-task environment. The primary and secondary tasks represent internal and external networks in an organization. All participants are presented with two types of scenarios – regular scenarios and surge scenarios. The difference between the two is that surge scenarios consist of secondary-task events that grow in threat-level and exceed that of concurrent primary-task events. Experimental results are expected to provide insight on the effect that workload previews have on attention-allocation, task management and cyber-SA in multi-task cyber-security contexts.

Prototype Development

We have developed a limited prototype of a visual analytic tool for the purpose of assessing its impact on SA. A large number of technical solutions attempt to address problems in cyber situation awareness. Typically, these solutions can be described as either data fusion technologies or visual analytic techniques. Though these two general areas of improvement will likely increase situation awareness, their impacts are seldom tested and proven in a systematic way. This part of our research seeks to answer the question of how to measure improvements in the human analyst's cyber SA. To measure improvements in cyber SA due to a data fusion or visual analytic artefact, a theoretically-grounded measurement technique must be developed specific to the cyber domain. This measurement technique must also be able to differentiate between increased cyber SA due to the knowledge and experience of the analyst from increased SA due to enhancements of the interface.

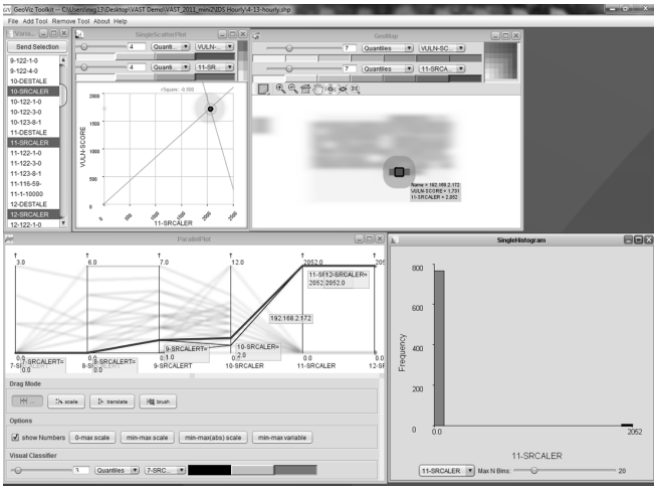

Figure 3 Visual Analytics Toolkit Prototype

As we develop the measurement technique, the simulated environment in which it is used must remain ecologically valid. The simulation that we have developed is relatively high fidelity, providing several diverse sources of cyber security data. We rely on the simulated data provided in the 2011 IEEE Visual Analytics Science and Technology (VAST) Mini-Challenge 2 including some from firewalls, intrusion detection systems, server logs and vulnerability scanners for a 3-day period over the same network (Grinstein, Whitting, Liggett, & Nebesh, 2011). These four sources of data provide a much better representation of a true cyber security environment than the single-sensor datasets previously published (Lippmann, R. P., Fried, et al., 2000; Lippmann, R., Haines, Fried, Korba, & Das, 2000; Sangster et al., 2009).

To measure SA improvement we rely heavily on SAGAT (Endsley, 1988) and its well-accepted theoretical model (Endsley, 1995b). In this project, we develop a set of freeze-probe queries to use in the simulated environment. Level 1 questions, such as which IP addresses are inside or outside of the network (D'Amico, Anita, Whitley, Tesone, O'Brien, & Roth, 2005), identify the participant's understanding of specific elements in the environment. Higher level questions probe at the memory constructs that should be present if the expected knowledge exists in the participant's working and long-term memory.

SAGAT alone, however, is unable to distinguish between SA based on knowledge and experience or whether the interface and underlying technologies provided the support for the insight. A combination of several SA measurement techniques to include the Situational Awareness Rating Technique (SART) (Taylor, 1990), National Aeronautics and Space Administration Task Load Index (NASA-TLX) (Hart & Staveland, 1988) and the Human Performance Scoring Model (Hamilton et al., 2010; Wellens & Ergener, 1988) in conjunction with a domain-specific cyber version of SAGAT should provide

sufficient measurement fidelity to be able to differentiate. A 2x2 between-subjects experiment execution should provide comparison of measures between experts and novices when presented with either high or low perceived workload interfaces.

The high perceived workload interface is what would be currently available to analysts. In this interface data is correlated by IP address, but generally individual element records of the four cyber security data sources are presented in list form. In the low perceived workload interface, we present the same level of correlation, but provide the data in a visual analytic interface. This allows for the individual cyber security data records to be displayed graphically using a geographic metaphor. Host system data is placed on a "geographic map" of the network with workstations physically separated from servers and the outside Internet visually. Coordinated views in the GeoViz Toolkit (Hardisty & Robinson, 2011) provide this functionality as well as a number of powerful visual analytic representations (Giacobe & Xu, 2011).

4 Conclusion

To conclude, we argue for a more human-centric approach to the study of situation awareness in computer defence analysis in order to yield greater insight into the socio-cognitive challenges of CDA work. Though valuable, much of the work done to date on situation awareness in CDA has done little to further our understanding of SA as either a cognitive state or process or empirically assessed the extent to which new technologies actually improve SA.

Our own work, which we present here, employs the Living Lab Framework to study CDA work and gain insight into both human cognitive processes related to CDA and the broader socio-technical context within which that work is done. Our findings from our field work indicate that CDA work is distributed across human actors and technological agents operating in different functional domains such as intrusion detection, forensics, and strategic analysis. We are currently engaged in multiple experiments using our scaled-world simulation – *NETS* – to examine questions related to transactive memory and CDA, the impact of task load on SA in CDA work, and the impact of a prototype visual analytic tool on SA in CDA work.

Acknowledgements

This work was partially supported by U.S. Army Research Office (ARO) MURI Grant "Computer Aided Human Centric Cyber Situation Awareness" W911-NF-09-1-0525, and by the U.S. Department of Homeland Security under Award #: 2009-ST-061-CI0001

References

Bardram, Jakob E., Hansen, Thomas R., & Soegard, Mds. (2006). *AwareMedia: a shared interactive dsplay supporting social, temporal, and spatial awareness in surgery*. Paper presented at the 20th Anniversary Conference on Computer Supported Cooperative Work (CSCW '06), Banff, Alberta, Canda.

Barford, P., Dacier, M., Dietterich, T.G., Fredrikson, J., Griffin, S., Jajodia, S., . . . Yen, John. (2010). Cyber SA: Situational Awareness for Cyber Defense - Issues and Research. In S. Jajodia, P. Liu, V. Swarup & C. Wang (Eds.), *Cyber Situational Awareness* (pp. 3-14). New York: Springer.

Blandford, Ann, & William Wong, B. L. (2004). Situation awareness in emergency medical dispatch. *International Journal of Human-Computer Studies, 61*(4), 421-452. doi: 10.1016/j.ijhcs.2003.12.012

Boyce, Michael W., Duma, Katherine M., Hettinger, Lawrence J., Malone, Thomas B., Wilson, Darren P., & Lockett-Reynolds, J. (2011). *Human Performance in Cybersecurity: A Research Agenda*. Paper presented at the Human Factors and Ergonomics Society 55th Annual Meeting, Las Vegas, NV.

Buford, J. F., Lewis, L., & Jakobson, G. (2008, June 30 2008-July 3 2008). *Insider threat detection using situation-aware MAS*. Paper presented at the 11th International Conference on Information Fusion.

D'Amico, A., & Larkin, M. (2001). *Methods of visualizing temporal patterns in and mission impact of computer security breaches*. Paper presented at the DARPA Information Survivability Conference & Exposition II, 2001 (DISCEX '01).

D'Amico, Anita, Whitley, Kirsten, Tesone, Daniel, O'Brien, Brianne, & Roth, Emilee. (2005). *Achieving Cyber Defense Situational Awareness: A Cognitive Task Analysis of Information Assurance Analysts*. Paper presented at the Human Factors and Ergonomics Society 49th Annual Meeting, Orlando, FL.

Ellis, A. P. J. (2006). System breakdown: The role of mental models and transactive memory in the relationship between acute stress and team performance. *Academy of Management Journal, 49*(3), 576-589.

Endsley, Mica R. (1988, 23-27 May 1988). *Situation awareness global assessment technique (SAGAT)*. Paper presented at the Aerospace and Electronics Conference, 1988 (NAECON 1988).

Endsley, Mica R. (1993). A Survey of Situation Awareness Requirements in Air-to-Air Combat Fighters. *The International Journal of Aviation Psychology, 3*(2), 157-168. doi: 10.1207/s15327108ijap0302_5

Endsley, Mica R. (1995a). Measurement of Situation Awareness in Dynamic Systems. *Human Factors, 37*(1), 65-84.

Endsley, Mica R. (1995b). Toward a Theory of Situation Awareness in Dynamic Systems. *Human Factors: The Journal of the Human Factors and Ergonomics Society, 37*(1), 32-64. doi: 10.1518/001872095779049543

Endsley, Mica R. (2000). Theoretical Underpinnings of Situational Awareness: A Critical Review. In M. R. Endsley & D. J. Garland (Eds.), *Situation Awareness Analysis and Measurement* (pp. 3-30). Mahwah, NJ: Lawrence Earlbaum Associates.

Endsley, Mica R., Farley, Todd C., Jones, William M., Midkiff, Alan H., & Hansman, R. John. (1998). Situation Awareness

Information Requirements for Commercial Airline Pilots. *DSpace@MIT*. http://hdl.handle.net/1721.1/35929

Endsley, Mica R., & Rodgers, Mark D. (1994). Situation Awareness Information Requirements Analysis for En Route Air Traffic Control. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting, 38*(1), 71-75. doi: 10.1177/154193129403800113

French, HT, & Hutchinson, A. (2002). *Measurement of situation awareness in a C4ISR experiment.* Paper presented at the 2002 Command and Control Research and Technology Symposium.

Giacobe, N., & Xu, S. (2011). *Geovisual Analytics for Cyber Security: Adopting GeoViz Toolkit*. Paper presented at the IEEE Symposium on Visual Analytics Science and Technology (VAST), Providence, RI.

Goodall, John R., Lutters, Wayne G., & Komlodi, Anita. (2004). *I Know My Network: Collaboration and Expertise in Intrusion Detection.* Paper presented at the ACM Conference on Computer-Supported Cooperative Work (CSCW), Chicago, IL.

Gorman, Jamie C., Cooke, Nancy J., & Winner, Jennifer L. (2006). Measuring team situation awareness in decentralized command and control environments. *Ergonomics, 49*(12-13), 1312-1325. doi: 10.1080/00140130600612788

Grinstein, Georges, Whitting, Mark, Liggett, Kristin, & Nebesh, Danko. (2011). IEEE VAST Challenge 2011. Retrieved from http://hcil.cs.umd.edu/localphp/hcil/vast11/

Hamilton, Katherine, Mancuso, Vincent, Minotra, Dev, Hoult, Rachel, Mohammed, Susan, Parr, Alissa, . . . McNeese, Michael. (2010). Using the Neocities 3.1 Simulation to Study and Measure Team Cognition. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting, 54*(4), 433-437. doi: 10.1177/154193121005400434

Hardisty, Frank, & Robinson, Anthony. (2011). The geoviz toolkit: using component-oriented coordination methods for geographic visualization and analysis. *International Journal of Geographical Information Science, 25*(2), 191-210. doi: 10.1080/13658810903214203

Hart, S.G., & Staveland, L.E. (1988). Development of NASA-TLX (Task Load Index): Results of empirical and theoretical research. *Human mental workload, 1*, 139-183.

Hazlehurst, Brian, McMullen, Carmit K., & Gorman, Paul N. (2007). Distributed cognition in the heart room: How situation awareness arises from coordinated communications during cardiac surgery. *Journal of Biomedical Informatics, 40*(5), 539-551. doi: 10.1016/j.jbi.2007.02.001

Jones, Rashaad E. T., McNeese, Michael D., Connors, Erik S., Jefferson, Tyrone, & Hall, David L. (2004). A Distributed Cognition Simulation Involving Homeland Security and Defense: The Development of Neocities. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting, 48*(3), 631-634. doi: 10.1177/154193120404800376

Lee, J. D., & Sanquist, T. F. (2000). Augmenting the operator function model with cognitive operations: assessing the cognitive demands of technological innovation in ship navigation. *IEEE Transactions on Systems, Man and Cybernetics, Part A: Systems and Humans, 30*(3), 273-285. doi: 10.1109/3468.844353

Lewis, K. (2004). Knowledge and performance in knowledge-worker teams: A longitudinal study of transactive memory systems. *Management science*, 1519-1533.

Lippmann, R. P., Fried, D. J., Graf, I., Haines, J. W., Kendall, K. R., McClung, D., . . . Zissman, M. A. (2000). *Evaluating intrusion detection systems: the 1998 DARPA off-line intrusion detection evaluation*. Paper presented at the DARPA Information Survivability Conference and Exposition 2000 (DISCEX '00), Lexington, MA.

Lippmann, Richard, Haines, Joshua W., Fried, David J., Korba, Jonathan, & Das, Kumar. (2000). The 1999 DARPA off-line intrusion detection evaluation. *Computer Networks, 34*(4), 579-595. doi: 10.1016/S1389-1286(00)00139-0

Llinas, J., & Hall, D. L. (1998). *An introduction to multi-sensor data fusion*. Paper presented at the 1998 IEEE International Symposium on Circuits and Systems (ISCAS '98), Monterey, CA. 10.1109/ISCAS.1998.705329

MacEachren, Alan M., Cai, Guoray, McNeese, Michael, Sharma, Rajeeve, & Fuhrmann, Sven. (2006). *GeoCollaborative Crisis Management: Designing Technologies to Meet Real-World Needs*. Paper presented at the 7th Annual National Conference on Digital Government Research (dg.o. 2005), San Diego, CA.

Mathew, S., Shah, C., & Upadhyaya, S. (2005, March 23-24). *An alert fusion framework for situation awareness of coordinated multistage attacks.* Paper presented at the Third IEEE International Workshop on Information Assurance, College Park, Maryland.

McGrath, Dennis, & McGrath, Susan P. (2005). *Simulation and Network-Centrick Emergency Response*. Paper presented at the The Interservice/Industry Training, Simulation & Education Conference (I/IT SEC), Hanover, NH.

McMillan, Eric, & Tyworth, Michael. (2012). An Alternative Framework for Research on Situational Awareness in Computer Network Defense. In C. Onwubiko & T. Owens (Eds.), *Situational Awareness in Computer Network Defense: Principles, Methods and Applications* (pp. 71-85). New York: IGI Global.

McNeese, Michael D. (1996). An Ecological Perspective applied to Multi-Operator Systems. In O. Brown Jr. & H. W. Hendrick (Eds.), *Human Factors in Organizational Design and Management - V* (pp. 365-370). Amsterdam: Elsevier Science B.V.

McNeese, Michael D., Bains, Priya, Brewer, Isaac, Brown, Cliff, Connors, Erik S., Jefferson, Tyrone, . . . Terrell, Ivanna. (2005). The Neocities Simulation: Understanding the Design and Experimental Methodology Used to Develop a Team Emergency Management Simulation. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting, 49*(3), 591-594. doi: 10.1177/154193120504900380

McNeese, Michael D., Perusich, Karl, & Rentsch, Joan R. (2000). Advancing Socio-Technical Systems Design Via the Living Laboratory. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting, 44*(12), 2-610-612-613. doi: 10.1177/154193120004401245

Moreland, R. L., & Myaskovsky, L. (2000). Exploring the performance benefits of group training: Transactive memory or improved communication? *Organizational Behavior and Human Decision Processes, 82*(1), 117-133.

Okolica, James, McDonald, J. Todd, Peterson, Gilbert L., Mills, Robert F., & Haas, Michael W. (2009). Developing Systems for Cyber Situational Awareness. *Proceedings of the 2nd. Cyberspace Research Workshiop*, 46-56. http://www.csc.latech.edu/Web%20Attachments/2009-crw-proceedings.pdf#page=52

Orlikowski, Wanda J., & Iacono, C. Suzanne. (2001). Research commentary: Desperately seeking "IT" in IT research - A call to theorizing the IT artifact. *Information Systems Research, 12*(2), 121-134.

Pearsall, Matthew J., & Ellis, Aleksander P. J. (2006). The Effects of Critical Team Member Assertiveness on Team Performance and Satisfaction. *Journal of Management, 32*(4), 575-594. doi: 10.1177/0149206306289099

Peng, Xie, Li, J. H., Xinming, Ou, Peng, Liu, & Levy, R. (2010, June 28 2010-July 1 2010). *Using Bayesian networks for cyber security analysis.* Paper presented at the IEEE/IFIP International Conference on Dependable Systems and Networks (DSN).

Perrow, Charles. (1999). *Normal Accidents: Living with High-Risk Technologies*. Princeton, NJ: Princeton University Press.

Salmon, Paul M., Stanton, Neville A., Walker, Guy H., Baber, Chris, Jenkins, Daniel P., McMaster, Richard, & Young, Mark S. (2008). What really is going on? Review of situation awareness models for individuals and teams. *Theoretical Issues in Ergonomics Science, 9*(4), 297 - 323.

Salmon, Paul, Stanton, Neville, Walker, Guy, & Green, Damian. (2006). Situation awareness measurement: A review of applicability for C4i environments. *Applied Ergonomics, 37*(2), 225-238. doi: DOI: 10.1016/j.apergo.2005.02.001

Sangster, Benjamin, O'Connor, T. J., Cook, Thomas, Fanelli, Robert, Dean, Erik, Adams, William J., . . . Conti, Gregory. (2009). *Toward instrumenting network warfare competitions to generate labeled datasets.* Paper presented at the 2nd Workshop on Cyber Security Experimentation and Test (CSET '09), Montreal, Canada.

Spradley, James P. (1979). *The ethnographic interview*. New York: Holt, Rinehart and Winston.

Stanton, N. A., Stewart, R., Harris, D., Houghton, R. J., Baber, C., McMaster, R., . . . Green, D. (2006). Distributed situation awareness in dynamic systems: theoretical development and application of an ergonomics methodology. *Ergonomics, 49*(12-13), 1288-1311. doi: 10.1080/00140130600612762

Stanton, Neville A., Salmon, Paul M., Walker, Guy H., & Jenkins, Daniel P. (2009). Is situation awareness all in the mind? *Theoretical Issues in Ergonomics Science, 11*(1-2), 29-40. doi: 10.1080/14639220903009938

Star, Susan Leigh. (1989). The Structure of Ill-Structured Solutions: Boundary Objects and Heterogeneous Distributed Problem Solving. In M. Kuhn & L. Gasser (Eds.), *Readings in Distributed Artificial Intelligence* (pp. 36-54). Menlo Park, CA: Morgan Kaufman.

Stotz, A., & Sudit, M. (2007, 9-12 July 2007). *Information fusion engine for real-time decision-making (INFERD): A perceptual system for cyber attack tracking.* Paper presented at the 10th International Conference on Information Fusion.

Tadda, George P., & Salerno, John S. (2010). Overview of Cyber Situation Awareness. In S. Jajodia, P. Liu, V. Swarup & C. Wang (Eds.), *Cyber Situational Awareness* (pp. 15-35): Springer US.

Tadda, George, Salerno, John J., Boulware, Douglas, Hinman, Michael, & Gorton, Samuel. (2006, April). *Realizing situation awareness within a cyber environment.* Paper presented at the SPIE Conference on Multisensor, Multisource Information Fusion: Architectures, Algorithms, and Applications, Orlando (Kissimmee), FL, USA.

Taylor, R.M. (1990). Situational awareness rating technique (SART): The development of a tool for aircrew systems design *Situational Awareness in Aerospace Operations (AGARD-CP-478)* (pp. 3/1-3/17). Neuilly Sur Seine, France: NAOT - AGARD.

Tyworth, Michael, Giacobe, N., Mancuso, Vincent, Dancy, Christopher, & McMillan, Eric. (2012). *Cyber Situation Awareness as Distributed Socio-Cognitive Work*. Paper presented at the SPIE Conference on Defense, Security & Sensing 2012, Baltimore, MD.

Walker, Guy H., Gibson, Huw, Stanton, Neville A., Baber, Chris, Salmon, Paul, & Green, Damian. (2006). Event analysis of systemic teamwork (EAST): a novel integration of ergonomics methods to analyse C4i activity. *Ergonomics, 49*(12-13), 1345-1369. doi: 10.1080/00140130600612846

Wegner, D. M., Giuliano, T., & Hertel, P. T. (1985). Cognitive interdependence in close relationships. In W. J. Ickes (Ed.), *Compatible and incompatible relationships* (pp. 253-276): Springer-Verlag.

Wellens, A.R., & Ergener, D. (1988). The CITIES Game. *Simulation & Gaming, 19*(3), 304.

Wise, Jeff. (2012). What Really Happened Aboard Air France 447. *Popular Mechanics*. Retrieved from Popular Mechanics website: http://www.popularmechanics.com/technology/aviation/crashes/what-really-happened-aboard-air-france-447-6611877

Yang, S. J., Byers, S., Holsopple, J., Argauer, B., & Fava, D. (2008). *Intrusion activity projection for cyber situational awareness*. Paper presented at the IEEE International Conference on Intelligence and Security Informatics (ISI 2008). Tapei, Taiwan. 10.1109/ISI.2008.4565048

Yang, Shanchieh J., Stotz, Adam, Holsopple, Jared, Sudit, Moises, & Kuhl, Michael. (2009). High level information fusion for tracking and projection of multistage cyber attacks. *Information Fusion, 10*(1), 107-121. doi: 10.1016/j.inffus.2007.06.002

Yen, John, McNeese, Michael D., Mullen, Tracy, Hall, D. L., Fan, Xiaocong, & Liu, P. (2010). RPD-based Hypothesis Reasoning for Cyber Situation Awareness. In S. Jajodia, P. Liu, G. Swarup & C. Wang (Eds.), *Cyber Situational Awareness: Issues and Research* (pp. 39-49). New York: Springer.