The Human Factors of Cyber Network Defense

Robert S. Gutzwiller[1], Sunny Fugate[1], Benjamin D. Sawyer[2], & P. A. Hancock[2]
[1]Space and Naval Warfare Systems Center Pacific
[2] University of Central Florida

Technology's role in the fight against malicious cyber-attacks is critical to the increasingly networked world of today. Yet, technology does not exist in isolation: the human factor is an aspect of cyber-defense operations with increasingly recognized importance. Thus, the human factors community has a unique responsibility to help create and validate cyber defense systems according to basic principles and design philosophy. Concurrently, the collective science must advance. These goals are not mutually exclusive pursuits: therefore, toward both these ends, this research provides cyber-cognitive links between cyber defense challenges and major human factors and ergonomics (HFE) research areas that offer solutions and instructive paths forward. In each area, there exist cyber research opportunities and realms of core HFE science for exploration. We raise the cyber defense domain up to the HFE community at-large as a sprawling area for scientific discovery and contribution.

INTRODUCTION

Cyberspace was first coined as a portmanteau of the terms *cybernetics* and *space* in the short story "Burning Chrome" (Gibson, 2003), inspired by Norbert Wiener's seminal work on cybernetics (Wiener, 1965). Colloquial meaning of the term cyberspace encompasses communication between computing devices (networks), the devices themselves and the interconnection of machines and networks to the physical world as both sensors and actuators (cyber-physical systems). Within the last decade, attacks in cyberspace have disrupted entire countries, disturbed critical infrastructure, and deeply affected information-based cultures and economies (Singer & Friedman, 2014). Cyber-attacks continue to increase both their rate (Garnaeva et al., 2014), and in the importance of their targets, as illustrated by over 200 attacks on major industrial control systems in 2013 alone (DHS, 2013). As technology proliferates, these vulnerabilities multiply. For example, in the increase of medical device software and hardware, the patient is now vulnerable to increasingly numerous and diverse exploit vectors (Sametinger et al., 2015).

To stem the flow of malicious attacks, human cyber analysts must monitor and protect cyberspace – i.e., *cyber defense*. Although to understand and act in the physical world is something humans are well-equipped to achieve, in the ethereal cyberspace, interfaces are the single point of connection used to extend human perception and action into the dense world of the network (e.g., Hancock, 2009). As such, humans in this domain are limited. Interfaces must compensate for most network activity occurring under the threshold for human reaction and response time and decision-making. The early techno-centric rush to bolster defensive capabilities only led to software creation, with almost no validation and only sparse *cognitive understanding* of cyber-defense performance. However, a growing recognition in industry is that the cognition of the operator must be studied, along with design to support human capabilities, rather than developing software or hardware alone (as it clearly does not meet security needs; Line et al., 2014).

The above suggests two thrusts for cyber human factors researchers. The first is to improve the community's understanding of cyberspace, in order to apply our research to solve defense challenges (increasing *breadth*). The second, parallel thrust is to push the boundaries of core HFE science, building on prior known areas (increasing *depth*). Cyber defense offers a rich research environment for HFE that will allow growth in both breadth and depth. However, identifying focus areas can be a challenge in such emergent domains.

HFE work is emerging for cyberspace; panel sessions at recent meetings (McNeese et al., 2012; Knott et al., 2013; Mancuso et al., 2014), and a special issue of the *Journal of Cognitive Engineering and Decision Making* (June, 2015), reflects cyberspace as a growing focus. Theory and frameworks for considering cyber-attacks are emerging. Some theory and frameworks are developed from the top-down (Mancuso et al., 2014), while others are created from the bottom up, building upon cognitive task analyses (CTAs) of cyber analyst roles and demands on cognition (D'Amico et al., 2005; Mahoney et al., 2010). The role of cyber-teams and team cognition has also been captured (Champion et al., 2012; Rajivan et al., 2013), and helps clarify the collaborations among different analysts and organizations.

Experiments in analyst decision making processes (e.g., Dutt, Ahn, & Gonzalez, 2012) represent excellent steps forward in our understanding of cyber defense. In particular, modeling provides high value, as operators are in high demand, and undergraduates may not be ideal experimental participants. Familiar territory exists in cyber defense, in that the job of cyber operators may be in part a vigilance task, thus leveraging another core HFE concept (Hancock, 2013; Sawyer et al., 2014). Although visualizations for network defense are not new, designs which capitalize on HFE principles (e.g., Bennett, 2014) and cognitive task analyses (Goodall & Sowul, 2009) are now emerging, and point to an enlightened way forward for these technologies.

The goal of this paper is to highlight a series of research areas likely to benefit cyber-defense operations, and to inspire advances in the core science. The additions presented in part echo Boyce et al. (2011), but we have also linked cyber challenges with cognitive areas of interest (e.g.,. "cyber-cognitive") and provided a much more specific research agenda for each in cyber defense. In particular, we highlight training and feedback, cognitive biases, situation awareness and interface design, multi-tasking, vigilance, and automation interaction, expanding cyber-HFE breadth and depth.

Training, Feedback and Cognitive Biases. Cyber-defense operators require extensive training, and accordingly the primary evidence-based training recommendation specific to cyber defense is to make sure operators have a wealth of experience observing threats (Dutt et al., 2013). Since the demand for these professionals is high but the available population is low, determining the best ways of rapidly training them is crucial. Professionals (but especially novices) face continuing and repeated training challenges because of the ever-changing cyber defense landscape.

A few of the most pressing problems are how to tailor training to benefit defense against specific kinds of attacks, and how to promote rapid knowledge and skill acquisition. Given the difficulty of the domain, trainers may be attracted to methods that have the operator practice "pieces" of the cyber task, such as in part-task training (Wightman & Lintern, 1985) or training with automation (Gutzwiller et al., 2013). These parts and automation could result in training by tool, by threat detection task, or even by event-escalation process stages. However, undeniably, operators often switch between multiple, disparate cyber-defense tasks, and between real-world tasks outside of cyberspace. Training should reflect these demands. Some methods of splitting tasks may be less effective: for example, the basic fractionation approach of training individual "pieces" precludes timesharing skill development. On the other hand, reducing cognitive demands during training (i.e. simplification, perhaps by lowering event rates) or variable priority training, may still allow timesharing skill to develop. Scaffolding and adaptive training are also approaches to address difficulty and rapid acquisition (Wickens et al., 2012).

Further complicating the operators role, systems may divorce operator action from observable feedback, violating the human-action cycle (Norman, 1992). Cyber systems are not devoid of all feedback, but rarely report an operators' *impact* (Roth, in McNeese et al., 2012). Additionally, as analysts make decisions, the rare feedback they receive possesses ambiguity and may be delayed enough that it is subject to cognitive biases and illusions of validity (Einhorn & Hogarth, 1978). For example, selecting potentially malicious events on the network for escalation and seeing which ones are valid as true attacks, still fails to take into account potentially numerous attacks *missed*, and may instill a false sense of confidence. Due to these types of biases, cyber-defense operations could serve as a promising test ground for emergent new techniques in bias mitigation (Clegg et al., 2014), and is also a nascent opportunity to explore the role of cognitive biases at the intersection of humans and technology.

Cyber-Cognitive Situation Awareness. At the core of complex, dynamic task performance exists cognition about the general state of the task and information critical to goal completion. Endsley's definition of situation awareness (SA; Endsley, 1995) as perception of critical elements of information, their comprehension, and projection of the environment into the future, is applicable to many operational realms including cyber defense. (In addition, see Champion et al., 2012, for a team SA perspective). It is necessary to first resolve the lexicon between HFE and the surrounding literatures in network security which have a different representation of cyber SA (c.f., Bass, 2000, as cyber SA is inexorably linked with data fusion and can be possessed by systems). Therefore, the new label "cyber-cognitive SA" (CCSA) differentiates *human* situation awareness of cyberspace operations. Validating how, and what to measure is an especially underdeveloped area for cyber (Tenney & Pew, 2006), as it can be for other domains (Smith & Hancock, 1995). It requires derivation of operators' goal-driven behaviors. Existing cognitive task analyses provide starting points (e.g., Mahoney et al., 2010). However, these analyses often lack granular measures of CCSA, or validation in actual cyber-defense performance (but see Champion et al. 2012).

In general, the evaluations needed to verify the usefulness of tools, interfaces, and choices made for their design are sparse at best. A recent review found the majority of publications on cyber SA are techno-centric (Franke & Brynielsson, 2014), and a further review of cyberspace visualizations found that less than half of the 130 papers surveyed included an evaluation of the interface – and only 3 included actual user performance! (Staheli et al., 2014). Clearly, the wide variety of cyber-defense interfaces have yet to fulfill their promise, and this unrealized potential explains why so many cyber tools continuously populate a ghostly menagerie, rather than serving as useful cognitive orthotics (Ford et al., 1997) for the cyber network defense operator.

Cyber defense also relies on the interface as the sole means to develop CCSA. There is no standard interface; many operators cobble together collections of open-source, off-the shelf software and custom-built scripts to do their jobs. Silva et al. (2014) observed 75 unique software tools in use. In part, these ad-hoc creations allow for individual flexibility, while commercial suites often lack functionality *and* require extensive training to use. Operator performance strategies often differ as a result of customization (Hao et al., 2013), further clouding empirical assessment of CCSA. Nevertheless, evaluating even these amalgamations for CCSA could be useful as a baseline.

Many existing cyber-defense tools fail to link information with the goals of operators (e.g., threat determination), making them unlikely to enhance CCSA or performance. A "cyber common operating picture", for example, does not guarantee CCSA just because it presents useful information. Computer system awareness does not mean the decision-maker will have the same understanding. A multitude of factors comes to bear on the incorporation of any display information into operators' situational awareness: if the operator is distracted, information is not tied to current or future goals, the interface is difficult to navigate, or the operator is fatigued and overload, such a cyber operating picture may not provide any value at all. Further, design principles must be consulted (e.g., display, ecological, and proximity compatibility), as they repeatedly show improved benefits under stressful situations, such as those expected for cyber operations. Following these empirically-based principles should improve usability (Wickens et al., 2013; Hancock, Sawyer & Stafford, 2015; and see Bennett, 2014). However, a reliance on *visual* interfaces, requiring visual attentional resources for every task, may greatly

diminish the ability of an analyst to build CCSA and creates incipient overload (Thompson et al., 2007). The development of multi-modal interfaces for cyber are materializing (Ballora et al., 2011) and create new design requirements, such as using ranges of sounds that are useful without being fatiguing or annoying (Vickers et al., 2014).

Finally, we note that interfaces are yet another attack vector, with unique exposure as the central point of human interaction with cyber-defense systems, and an emerging challenge of good design will be creating them to be relatively impervious to perceptual manipulations brought on by malicious data transmissions (e.g., Conti et al., 2005). Designing on principle and then empirically validating the results are two key aspects likely to provide improvement in cyber-defense interfaces, and CCSA – but testing against simulated malicious attacks designed to disrupt or trick interface presentations must be undertaken.

Multi-Tasking and Attention Modeling. Cyber defense analysis is, in effect, a sequence of small "experiments", with each one testing a hypothesis about computer systems, information flows, and attackers' intent. Analysts often search large spaces of data (sometimes exhaustively), each search taking minutes up to hours, with the returns occurring at random time intervals due to the inherent unpredictable nature of the data being examined. These interruptions result in additional needs for cognitive effort, and effective cues to resume the interrupted task (Altmann & Trafton, 2002; 2004). Interruptions also create prospective memory demands, e.g. remembering what to do next for a given task (Dodhia & Dismukes, 2009). The study of task management provides a starting point for understanding how operators balance monitoring the network, taking investigative and reporting actions, and communicating information to others.

Clearly, the demand for exists in cyber defense to improve these attention allocation problems (Fink et al., 2009). Recent advances have provided models which track and predict attention allocation in terms of the "eyeball", by understanding the salience, effort, expectancy and value of information (SEEV; Wickens, 2014) in the visual environment. Progress is also being made in understanding how attention allocation *choices* are made when operators are overloaded (Gutzwiller et al., 2014; and see Hancock & Warm, 1989): in other words, modeling aspects of the "mindball" (Wickens et al., 2015). Once validated for cyber defense, these models could lead to rapid understanding and testing of how various cyber-defense interfaces would influence or alter operator behaviors. For example, such models can help determine whether operators may "tunnel" into certain cyber-defense tasks, or perhaps even specific cyber-defense displays (Wickens & Alexander, 2009).

Hedonomic Design and Cyber Vigilance. The work in cyber defense resolves issues through patching vulnerabilities, preventing intrusions, monitoring the network for attackers, and determining who they are and their intentions. These activities catch some attacks and successfully thwart others, but occasionally attacks *are* missed. The reward for completing each of these effortful security tasks in all cases is simply *more tasks to do* – more vulnerabilities to patch, more issues to resolve, and more intrusions to monitor. This negative performance metric – "how did I/we fail *this* time?" - appears to be a significant source of input to operators' day-to-day experience and work.

Cyber interfaces tend to exacerbate this negative interaction. Upon making a decision, for example, all signs of an event disappear, and there is no providence of effective feedback. Furthermore, analysts may feel disconnected from their job, and disconcerted that their prior decisions seem of no measurable value or impact. This psychological burden of joyless operation (and resulting frustration) is contributing to turnover and burnout (see Hancock & Warm, 1989).

Once again, the focal point becomes improving the interface for cyber defense. Hedonomic design approaches suggest that once an interface facilitates safe, effective and usable performance, further design and experimentation should determine how to make these interactions *pleasurable*. Efforts to establish cyber doctrine, infrastructure, interfaces - and the longevity with which the cyber domain will be relevant - suggest designing for cyber defense presents a prescient opportunity to incorporate these hedonomics principles (Hancock, Pepe, & Murphy, 2005). These principles should then be tested and refined, both of which will contribute to cyber defense and design science.

Negative cycles of interaction observed in cyber defense are also remarkably similar to those observed in other domains. Lengthy, repetitive work with little or no positive feedback, a rare signal (attacks) within large problem spaces, and help from automated systems that can sometimes be as overwhelming as non-automated performance are hallmarks of a *vigilance* task**,** a construct well studied in air traffic control and medical device monitoring.

The vigilance problem may be iatrogenic in nature, a result of the *artificiality* of the display and visualizations implemented (see Hancock, 2013). It is perhaps not surprising then, that cyber defense has this characterization (Sawyer et al., 2014; 2015). The fact that vigilance issues exist in cyber defense is a fortunate reality, given amassed knowledge about causes and mitigations of vigilance decrements. Clear design recommendations follow from each of these basic tenants, and such interventions have a history of leading not only to greater efficiency, but also to healthier, happier operators (Sawyer et al., 2015). This well-founded optimism regarding mitigation does not belie the fact that this domain is novel, and much is still unknown. Cyber-vigilance is perhaps unique among prior vigilance tasks in the complexity of the signal, specifically in the technologically and philosophically diverse delivery methods. It is also unique regarding the uncertainty surrounding ground truth, as a well-executed cyber-attack need leave no trace.

Human Automation Interactions. Incentives for cyber automation are high, particularly because of the benefits related to millisecond-level reaction time and decision-making and the massive scale of the domain. History predicts engineers will attempt to "automate everything that can be automated" (Bainbridge, 1983). Yet automated decision making is one particularly well-known way to create problems for operators (Onnasch et al., 2014). A short-term solution

may be to avoid automating decision-making, and only explore other automation schemas.

Long term, examining the notion of adaptive automation in this new domain may prove fruitful, as it mitigates many interaction issues by matching the dynamism found in the real world and in the operators themselves with a commensurate level of assistance (Kaber, 2013). Integrating *working agreements* (de Greef et al., 2010), which articulate and constrain the operations of the system, and user expectancies for human-automation collaborations, would be expected to bring benefits to both. However, particular solutions must be subjected to experimental jeopardy in the cyber defense environment (and in others, see Gutzwiller et al., 2015).

Perennial questions will arise again concerning issues of transparency and trust (see Hancock et al., 2011) as cyber operations begin using intelligent aiding. Distrust of automation and hidden mechanisms seems implicated in the refusal of operators in using novel interfaces. Within the data and algorithms, there is a perceived lack of definition. As a result, trust in newer tools may be inappropriate. Similarly, trust accumulated in more automated tools may end up being misguided, creating complacency and automation bias (Parasuraman & Manzey, 2010). The "old guard" ardently disregards newer tools. But, it is interesting to note that tools are likely to have the most success implanting when significant operator turnover occurs, suggesting a sort of *inertia* in their use, and harkening to the training challenges mentioned in the sections above.

Finally, notifications through alarms and alerts are a particular concern of intrusion detection systems in cyber defense. Operators must contend with alerts that may or may not be "hits" (Champion et al., 2012). Setting up the conditions for the alerts themselves could reap benefits from existing work regarding alarms in supervisory control (Stanton et al., 1992; Woods, 1995). For example, if excessively triggered alarms may not reflect priority, require interruption (or cause it), or trigger without respect to the operators' current cognitive context - they are more likely to *degrade* task performance, not improve it (Woods, 1995). Similar issues should be examined behaviorally within cyber defense performance.

DISCUSSION

The future of HFE in cyber defense is one of prime importance to life, as cyber threats target more aspects of our existence, and with increasing precision and impact (e.g., the recent OPM breach, which released secret information about millions of federal employees). Fundamental human capacities for information processing limit the ability of operators to defend this element of our lives. Thus, HFE serves as an originating point for considering and augmenting cognition in cyberspace. The cyber operational domain is still in its relative infancy, but this newness is a fortuitous state. The HFE science in many of the areas we highlight is mature and can provide utility to the cyber domain (but also reaps benefits from collateral exploration).

Much of this paper points out where HFE is not yet exercising full influence over the course of cyberspace defense development. However, this is not a negative pronouncement, despite several ongoing challenges. For example, access to the domain may be limited to real-world environments, which poses terminal challenges (Paul, 2014). *Confidentiality* further hinders research in cyberspace by placing limits on research *and* the publishing process. Rapid evolution imposes the difficulty of experimenting in an always-changing environment, such that unless it touches on a common cognitive core, or a principle of the task domain, many evaluations and experiments could become obsolete. However, these efforts are still important for establishing HFE within cyber research.

The field will also continue to struggle with the tradeoffs between maintaining internal and external validity. Toward that end, conducting human-in-the-loop cyber-defense research in the laboratory setting appears difficult, because the platform and simulation capabilities are still being developed (see CyberCog - Champion et al., 2012, which is being built into DEXTAR - Shope, 2013; and the idsNETs testbed - Mancuso et al., 2012). Note that this struggle is fundamentally similar to the recent problems overcome in experimentation in multiple unmanned vehicle control. The equivalent platforms for research are required for studying both defensive and offensive cyber-operations. These particular challenges are opportunities to strengthen HFE ties to computer science, between universities, and between Department of Defense laboratories to help build and make these testbeds available. Thus, HFE can then become welcome apostles to these communities, strengthening the inter-disciplinary value added.

In conclusion, a well-defended cyber environment will almost certainly rely on humans. It is positive news that so much of the cyberspace defense domain is ripe for study, with mutual benefits to the defense tasks, and to HFE science. However much work remains to be completed and in a rapid manner to address the pressing challenges to security, which increase daily.

Note: The views and opinions expressed in this article are solely those of the authors, and do not reflect the official policy or position of any agency of the U.S. government. We also thank Phillip Verbancsics for his insightful commentary.

References

Altmann, E. M., & Trafton, J. G. (2002). Memory for goals: an activation-based model. *Cognitive Science*, *26*(1), 39–83.

Altmann, E. M., & Trafton, J. G. (2004). Task interruption: Resumption lag and the role of cues. *Michigan State University.*

Bainbridge, L. (1983). Ironies of automation. *Automatica*, *19*(6), 775–779.

Ballora, M., Giacobe, N. A., & Hall, D. L. (2011). Songs of cyberspace: An update on sonifications of network traffic to support situational awareness. *SPIE Defense, Security, & Sensing*, *8064*, 80640P–6.

Bass, T. (2000). Intrusion detection systems and multisensor data fusion: creating cyberspace situational awareness. *ACM, 43*(4), 99-105.

Bennett, K. B. (2014). Veils: an ecological interface for computer network defense. *Proc Hum Fctrs Ergnmcs Soc Ann Mtg*, *58* (1), 1233-37.

Boyce, M., Duma, K., Hettinger, L., Malone, T., Wilson, D., & Lockett-Reynolds, J. (2011). Human performance in cybersecurity: a research agenda. *Proc of the Hum Fctrs Ergnmcs Soc Ann Mtg*, *55*, 1115–1119.

Champion, M., Rajivan, P., Cooke, N., & Jariwala, S. (2012). Team-based cyber defense analysis. *CogSIMA*, 218–221.

Clegg, B., Martey, R., Stromer-Galley, J., Kenski, K., Saulnier, T., Folkestad, J., … Tomek, S. (2014). Game-based training to mitigate three forms of cognitive bias. *I/ITSEC*, (14180), 1–12.

Conti, G., Ahamad, M., & Stasko, J. (2005). Attacking information visualization system usability overloading and deceiving the human. *Symposium on Usable Privacy and Security - SOUPS '05*, 89–100.

D'Amico, A., Whitley, K., Tesone, D., O'Brien, B., & Roth, E. (2005). Achieving cyber defense situational awareness: a cognitive task analysis of information assurance analysts. *Proc of the Hum Fctrs Ergnmcs Soc Ann Mtg*, *49*, 229–233.

De Greef, T., Arciszewski, H., & Neerincx, M. (2010). Adaptive automation based on an object-oriented task model: Implementation and evaluation in a realistic C2 environment. *JCEDM*, *4*(2), 152–182.

DHS. (2013). Incident Response Activity: Trends in incident response in 2013. *ICT-CERT Monitor*, (December), 1–14.

Dodhia, R., & Dismukes, R. (2009). Interruptions create prospective memory tasks. *Applied Cognitive Psychology*, *89*, 73–89.

Dutt, V., Ahn, Y.-S. Y.-S., & Gonzalez, C. (2013). Cyber situation awareness: modeling detection of cyber attacks with instance-based learning theory. *Human Factors*, *55*(3), 605–618.

Einhorn, H. J., & Hogarth, R. M. (1978). Confidence in judgment: Persistence of the illusion of validity. *Psychological Review*, *85*(5), 395–416.

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

Fink, G. A., North, C. L., Endert, A., & Rose, S. (2009). Visualizing cyber security: Usable workspaces. *Visualization for Cyber Security*, 45–56.

Ford, K., Glymour, C., & Hayes, P. (1997). Cognitive prostheses. *AI Magazine*, *18*(3), 1997.

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

Garnaeva, M., Chebyshev, V., Makrushin, D., Unuchek, R., & Ivanov, A. (2014). Kaspersky security bulletin 2014. *Kaspersky Lab*, 1–31.

Gibson, W. (2003). *Burning Chrome* (p. 179). New York, NY: HarperCollins.

Goodall, J. R., & Sowul, M. (2009). VIAssist: Visual analytics for cyber defense. *IEEE Conference on Tech for Homeland Security*, 143–150.

Gutzwiller, R. S., Clegg, B. A., & Blitch, J. G. (2013). Part-task training in the context of automation: Current and future directions. *American Journal of Psychology, 126*(4), 417-432.

Gutzwiller, R. S., Lange, D. S., Reeder, J., Morris, R. L., & Rodas, O. (2015). Human-computer collaboration in adaptive supervisory control and function allocation of autonomous system teams. *HCI International*.

Gutzwiller, R. S., Wickens, C. D., & Clegg., B. A. (2014). Workload overload modeling: An experiment with MATB II to inform a computational model of task management. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting, 58*, 849-853

Hancock, P. A. (2009). *Mind, machine and morality: Toward a philosophy of human-technology symbiosis* (pp. 1–184). Ashgate Publishing, Ltd.

Hancock, P. A, (2013). In search of vigilance: the problem of iatrogenically created psychological phenomena. *American Psychologist, 68*(2), 97.

Hancock, P. A., Billings, D. R., Schaefer, K. E., Chen, J. Y. C., de Visser, E. J., & Parasuraman, R. (2011). A meta-analysis of factors affecting trust in human-robot interaction. *Human Factors*, *53*(5), 517–527.

Hancock, P. A., Sawyer, B. D., & Stafford, S. (2015). The effects of display size on performance. *Ergonomics, 58*(3), 337-354.

Hancock, P. A., Pepe, A. A., & Murphy, L. L. (2005). Hedonomics: The power of positive and pleasurabe ergonomics. *Ergonomics in Design*, *13*(1), 8–14.

Hancock, P. A., & Warm, J. (1989). A dynamic model of stress and sustained attention. *Human Factors, 31*(5), 519-537.

Hao, L., Healey, C. G., & Hutchinson, S. E. (2013). Flexible web visualization for alert-based network security analytics. *VizSec*, 33–40.

Kaber, D. B. (2013). Adaptive automation. In J. Lee & A. Kirlik (Eds.), *Oxford Handbook of Cognitive Engineering*. Oxford University Press.

Knott, B., Mancuso, V., Bennett, K., Finomore, V., McNeese, M., McKneely, J., et al. (2013). Human factors in cyber warfare: Alternative perspectives. *Proc of Hum Fctrs Ergnmcs Soc Ann Mtg*, *57*, 399–403.

Line, M., Zand, A., Stringhini, G., & Kemmerer, R. (2014). Targeted attacks against industrial control systems: is the power industry prepared? *ACM Workshop on Smart Energy Grid Security*, 13–22.

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

Mancuso, V. F., Christensen, J. C., Cowley, J., Finomore, V., Gonzalez, C., & Knott, B. (2014). Human factors in cyber warfare II: Emerging perspectives. *Proc of the Hum Fctrs Ergnmcs Soc Ann Mtg*, *58*, 415-8.

Mancuso, V. F., Minotra, D., Giacobe, N., McNeese, M., & Tyworth, M. (2012). idsNETS: An experimental platform to study situation awareness for intrusion detection analysts. *CogSIMA*, 73–79.

Mancuso, V. F., Strang, A. J. A., Funke, G. J., & Finomore, V. S. (2014). Human factors of cyber attacks: A framework for human-centered research. *Proc of the Hum Fctrs Ergnmcs Soc Ann Mtg*, *58*, 437-41.

McNeese, M., Cooke, N. J., D'Amico, A., Endsley, M. R., Gonzalez, C., Roth, E., & Salas, E. (2012). Perspectives on the role of cognition in cyber security. *Proc of Hum Fctrs Ergnmcs Soc Ann Mtg*, *56*, 268-71.

Norman, D. (1992). Design principles for cognitive artifacts. *Research in Engineering Design*, *4*, 43–50.

Onnasch, L., Wickens, C. D., Li, H., & Manzey, D. (2014). Human performance consequences of stages and levels of automation: an integrated meta-analysis. *Human Factors*, *56*(3), 476–488.

Parasuraman, R., & Manzey, D. H. (2010). Complacency and bias in human use of automation: An attentional integration. *Human Factors*, *52*(3), 381–410.

Paul, C. (2014). Human-centered study of a network operations center: experience report and lessons learned. *ACM Workshop*, 39–42.

Rajivan, P., Janssen, M., & Cooke, N. (2013). Agent-based model of a cyber security defense analyst team. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting, 57*, 314-318.

Sametinger, J., Rozenblit, J., Lysecky, R., & Ott, P. (2015). Security challenges for medical devices. *Comm of the ACM*, *58*(4), 74–82.

Sawyer, B. D., Finomore, V. S., Funke, G., Warm, J. S., Matthews, G, Hancock, P.A. (2015). Cyber vigilance: the human factor. *Ergonomics*

Sawyer, B. D., Finomore, V. S., Funke, G., Mancuso, V., Funke, M.,...Warm, J. S. (2014). Cyber vigilance: effects of signal probability and event rate. *Proc of the Hum Fctrs Ergnmcs Soc Ann Mtg*, *58*(1), 1771–1775.

Shope, S. (2013). Effective cyber situation awareness (CSA) assessment and training. *US Army Final Report W911NF-13-C-0060.*

Singer, P., & Friedman, A. (2014). *Cybersecurity: What everyone needs to know* (pp. 1–224). USA: OUP.

Staheli, D., Yu, T., Crouser, R., Damodaran, S., Nam, K., O'Gwynn, … Harrison, L. (2014). Visualization evaluation for cyber security: Trends and future directions. *VizSec*, 49-56.

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

Stanton, N., Booth, R., & Stammers, R. (1992). Alarms in human supervisory control: A human factors perspective. *International Journal of Computer Integrated Manufacturing*, *5*(2), 81–93.

Tenney, Y., & Pew, R. (2006). Situation awareness catches on: What? So what? Now what? *Reviews of Human Factors and Ergonomics*.

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

Vickers, P., Laing, C., Debashi, M., & Fairfax, T. (2014). Sonification aesthetics and listening for network situational awareness. *SoniHED.*

Wickens, C. D. (2014). Noticing events in the visual workplace: The SEEV and NSEEV models. In R. Hoffman & R. Parasuraman (Eds.), *Handbook of Applied Perception*. Cambridge University Press.

Wickens, C. D., & Alexander, A. L. (2009). Attentional tunneling and task management in synthetic vision displays. *International Journal of Aviation Psychology*, *19*(2), 182–199.

Wickens, C. D., Gutzwiller, R. S., & Santamaria, A. (2015). Discrete task switching in overload: A meta-analysis and a model. *International Journal of Human-Computer Systems, 79*, 79-84.

Wickens, C. D., Hollands, J., Banbury, S., & Parasuraman, R. (2013). *Engineering psychology and human performance* (4th ed.). Upper Saddle River, NJ: Pearson.

Wickens, C. D., Hutchins, S., Carolan, T., & Cumming, J. (2012). Attention and cognitive resource load in training strategies. A. F. Healy & L. E. Bourne (Eds.), *Training Cognition* (pp. 67–88). NY:Psychology Press.

Wiener, N. (1965). *Cybernetics or Control and Communication in the Animal and the Machine (Vol. 25)*. MIT Press.

Wightman, D. C., & Lintern, G. (1985). Part-task training for tracking and manual control. *Human Factors*, *27*(3), 267–283.

Woods, D. D. (1995). The alarm problem and directed attention in dynamic fault management. *Ergonomics*, *38*(11), 2371–2393.