Framework for Developing a Brief Interview to Understand Cyber Defense Work: An Experience Report

Miriam E. Armstrong, Keith S. Jones, Akbar Siami Namin
Texas Tech University

Cyber defense is increasingly important for the wellbeing of our economy and our national defense. Universities can help meet our growing cybersecurity needs by training the next generation of cyber defenders, and it is crucial that the curricula for such programs are designed to prepare students for the type of work that is performed in the field. Unfortunately, collecting data about cyber work is hindered in situations where cybersecurity professionals are uncomfortable with traditional human factors work analysis methods. Four potential constraints are 1) no naturalistic observations, 2) anonymity and safety, 3) short data collection time, and 4) no deep process questions. We developed a brief interview technique that allowed us to measure the importance of knowledge, skills, and abilities related to offensive and defensive cyber work. Based on our experience using this technique, it fits within the four potential constraints to cyber research and produces information that is directly applicable to the development of cybersecurity curricula. Our technique could potentially be used for other research purposes and personnel selection and by researchers interested in other high-security populations.

Individuals and organizations communicate, enact financial transactions, and store information digitally. Hackers are incentivized to steal or corrupt such information for monetary gain, and the number of reported cybersecurity incidents continues to grow (Shoemaker, Kohnke, & Sigler, 2016).

Unfortunately, the cybersecurity workforce is not growing fast enough to meet our country's rapidly increasing cyber defense demands. Government agencies and private industries have difficulty filling positions (McGettrick, 2013), with postings for cybersecurity jobs taking longer to fill than those for information technology jobs overall (Restuccia, 2015). The Bureau of Labor Statistics predicts that cyber-related professions will experience 18% growth in 10 years, well above the national average for job growth (BLS, 2015).

Universities must play a role in alleviating the cyber workforce shortage by creating more courses and programs to educate future cybersecurity professionals (Locasto, Ghosh, Jajodia, & Stavrou, 2011; NSA, 2016). Ideally, these courses would prepare students for the cyber job market and decrease the demands of post-graduate training (Vieane et al., 2016). Curricula should be based off of an understanding of cyber work and of which aspects of work are most important for students to know upon graduation.

The human factors field recognizes the need for a better understanding of cybersecurity work (Dodge, Toregas, & Hoffman, 2012; Knott et al., 2013; Lathrop, Trent, & Hoffman, 2016; Madhavan, Cooke, Bass, Meyer, & Roth, 2015; Mancuso et al., 2014; Oltramari, Henshel, Cains, & Hoffman, 2015), and there have been efforts to study aspects of cyber defense work (e.g., D'Amico & Whitley, 2008; Mahoney et al., 2010; Stanard et al., 2004).

Unfortunately, the methodologies often employed by human factors researchers in order to document work are not always feasible in high-security fields such as cyber. Our research team encountered such difficulties when attempting to interview and observe cyber attackers and defenders. From early on, we consulted with a subject matter expert (SME) who has worked closely with cyber professionals for over 5 years and who had conducted research with cyber professionals as participants. Based on our conversations with the SME, we identified four potential constraints on research involving cybersecurity professionals: 1) no naturalistic observations, 2) anonymity and safety, 3) short data collection time, and 4) no deep process questions. Researchers could experience difficulty recruiting cybersecurity attackers or defenders as participants with a procedure that exists outside of one or more of these constraints. Of course, adhering to the four constraints provides its own set of challenges, as most task analysis and cognitive task analysis data collection techniques do not fit within these constraints.

We will describe the four constraints in further detail, but first it is worth acknowledging that our SME's experience within the cybersecurity field is not universal. Prior literature indicates that researchers have successfully gained access to cyber professionals with methodologies that include naturalistic observations, in-depth interviews, or both (D'Amico & Whitney, 2008; Mahoney et al., 2010). That said, we had two reasons for taking the constraints laid out by our SME seriously. First, based on conversations with other human factors researchers, running into one or more of these constraints is a common occurrence. Second, while not all cyber research appears bound by the four constraints, it is unclear in which circumstances one or all of the constraints would apply and in which circumstances they would not apply. For these reasons, it was prudent to proceed with the assumption that our methodology should fit within the four constraints.

In the following experience report we discuss the potential constraints present in cybersecurity research, how we developed a brief interview technique to work within the constraints, and the effectiveness of our methodology. Our interview was intended to provide a better understanding of the knowledge, skills, and abilities (KSAs) that are most important to include in cybersecurity curricula. Those interested in other aspects of cybersecurity work will likely confront the same four constraints and thus may still benefit from a discussion of our methods and experiences. It is also

possible that researchers interested in other high-security populations will find this experience report helpful.

DEVELOPMENT OF THE BRIEF INTERVIEW

In each of the four proceeding sections, we provide an overview of a potential research constraint and the way in which our brief interview technique accommodates the constraint. When appropriate, alternative means of complying with the constraint are discussed.

First Constraint: No Naturalistic Observations

Cybersecurity professionals are charged with protecting and defending information. Even ethical attackers, such as penetration testers, are concerned with keeping their client's information secure. Our SME cautioned that most cyber professionals would be unwilling to let researchers watch them work even if the researchers were granted security clearances or had signed a non-disclosure agreement. Part of the reason is that the act of performing cybersecurity work is itself part of the information that the cyber professional is guarding. Cyber professionals do not wish to divulge their trade secrets, especially not to anyone who might publish them, for fear that this would expose their company's vulnerabilities or the cyber professional's tactics. Either would put them at a disadvantage in future adversarial situations.

From a practical note, even if some cyber professionals were willing to participate, observations can be challenging because cyber attackers and defenders often work in teams (e.g., Trent, Hoffman, Leota, Frost, & Gonzales, 2016). If one team member does not consent to participation, then researchers cannot observe any work conducted by the team. This would significantly cut down on the amount of and usefulness of any data collected.

Due to this constraint, we elected to forego naturalistic observations and instead focused on collecting data through interviews. An alternative solution would have been to perform observations of training scenarios (e.g., Champion et al., 2012). This solution allows researchers to observe work and problem-solving processes in real-time and can be a close approximation of observing real cyber work. Often the training scenarios are based on real past events. Those participating in training may be cybersecurity novices or, depending on the program, they may be established security professionals attempting to learn a new skill. Ultimately, we did not choose this option because the observation of training scenarios was not appropriate for the types of information about cyber work that we were trying to collect. Observing work that is already part of instruction does not provide insight into the aspects of work that professionals find to be the most important.

Second Constraint: Anonymity and Safety

Cybersecurity professionals prefer to remain anonymous. This can again be to protect their employer or to protect themselves personally. The type of information that is considered too personal can vary depending on the situation. In some cases, the researcher should not ask for demographic information, in other cases they should not ask for personal information, and in most cases the researcher should consider the costs and benefits of collecting data online.

Researchers should be cautious of asking for demographic information such as race or gender when recruiting from small participant pools. This applies, for example, when all potential participants work at one or a few agencies or organizations and especially if it's easy for readers to find out which agencies and organizations the researchers recruited from. In such situations, the recruitment pool is already likely to be small and so answers to demographic questions would provide identifying information.

If recruiting from a larger and more diverse participant pool, such as when recruiting at conferences or through snowballing, participants may be more willing to provide demographic information but may be unwilling to provide personal information such as their name or their workplace. For cyber professionals involved in illegal activity, this anonymity can be a self-preservation measure. But the community as a whole is distrustful of questions about personal information because of concerns about social engineering: being manipulated into sharing compromising personal information through social obligations such as being polite. As such, cyber professionals become suspicious when someone asks for personal information. In cases of strict anonymity, participant compensation can be difficult, if not impossible. Many research institutions require the collection of personal information as part of the compensation process. Doing research without compensation results in its own challenges. For example, in the next section we will discuss the time constraints placed on researchers trying to recruit participants without compensation.

Long-distance data collection, such as through online surveys, is one means of learning about work without observation and has the added benefit of being low-cost and minimally demanding on researchers' time. However, our SME advised us to not use electronic means of data collection because cyber professionals may not trust the security of a digital platform. The nature of cybersecurity research is such that those who are sent an online survey are those with the skills to hack it. This would be a nuisance to the researcher but, in cases where demographic or personal information is collected, could also pose a real risk to participants. For similar reasons, other digital means of recording information, such as video or voice recordings, were discouraged.

To work under these constraints, we elected to not collect demographic information on gender or race and administer and record interviews on paper. This resulted in a methodology that could be used for any cyber participant pool. We decided to maximize the number of participants we could recruit by recruiting at well-populated hacking conferences. For this reason, we forewent participant compensation. Researchers recruiting from a small participant pool may be required to compensate participants, especially if collecting data during work hours. In this case, it would be necessary to collect personal information.

We did feel it necessary to ask some type of demographic questions in order to gauge the expertise of the participants. Our SME recommended two questions that would not produce identifiable information: "How many years have you been

interested in cyber?" and "How many capture-the-flag events have you participated in?" We also asked participants what was their highest level of education and, when applicable, what was their major.

Third Constraint: Short Data Collection Time

A cybersecurity professional's time is valuable, both in the sense that the work that they do is important and in the sense that their pay is lucrative (median income of $43.33/hour, BLS, 2015). For that reason, cyber professionals may prefer to not be interrupted from their job for extended periods of time.

In cases when researchers are not planning to compensate participants, potential participants are being asked to give up their free time. As with most populations, it is easier to ask cyber professionals to volunteer for shorter rather than longer amounts of time.

To work within this constraint, we built an interview intended to last 15 minutes (20 maximum). Although we were building an interview primarily around the time constraints associated with a large, anonymous participant pool, our SME indicated that cyber professionals from smaller pools would also be unlikely to participate if they expected that the researcher would take over 15 minutes of their time. Researchers who recruit from small participant pools would be expected to compensate participants at least as much as they would make for the same amount of time spent at their job.

Fourth Constraint: No Deep Process Questions

In-depth interviews can pose the same risks to cybersecurity professionals as do naturalistic observations. Namely, participants could be put at a tactical disadvantage if the researcher published too much information about how the cyber professionals perform their work.

Therefore, we were presented with the task of constructing shallow, non-threatening questions about cyber work that would provide enough useful information to help guide curriculum development. We focused on the knowledge, skills, and abilities (KSAs) necessary in the current cyber workforce. KSAs can be collected through a variety of methods including past literature, training manuals, and SME input (Salas et al., 1999). We relied on the KSAs found within the National Initiative for Cybersecurity Education (NICE) Cybersecurity Workforce Framework (Shoemaker et al., 2016). The Framework identifies over 1000 KSAs used by cyber professionals and 79 KSAs within the protect and defend knowledge area which is the knowledge area that best covers the cybersecurity attacking and defending work of interest in this project. Because we wanted our survey to last for 15 minutes or less, we decided against including all 79 KSAs. Some of the KSAs were listed as necessary for multiple specialty areas within the protect and defend knowledge area. For example, the KSA "Knowledge of basic system administration, network, and operating system hardening techniques" appeared under both the Computer Network Defense Analysis and Incident Response specialty areas. There were 32 such repeated KSAs within the NICE Framework protect and defend knowledge area, and we determined that this amount was appropriate for the 15-minute timeframe we wanted. An additional benefit of using the 32 KSAs for the survey was that these were likely to be considered more important to all participants than the KSAs that were only relevant to one specialty area. This was appropriate for our project because we were interested in determining the KSAs that were most important to cyber attacker and defenders generally in order to make suggestions for cyber curricula.

We asked two questions relating to each of the 32 KSAs. The first question was, "How important is [the KSA] to your job?" Participants responded by choosing a number on a Likert scale: 1 to 6 with 1 being "not important at all" and 6 being "very important." The second question was, "Where did you learn [the KSA]?" This was an open-ended question; if participants provided multiple answers they were asked to specify where they had learned the most about the KSA. The answers provided fell into 5 categories: KSA was learned at work/on the job (job); KSA was learned at school (school); KSA was learned through self-study (self); KSA was learned someplace else including through government training, certification programs, and friends (other); or the KSA was not learned at all (n/a).

These two questions were chosen to help curriculum developers prioritize the KSAs. KSAs rated as being more important should be emphasized more so than those rated as less important. Where the KSAs were learned would indicate whether current training programs were meeting the demands to the cyber workforce and, if not, which KSAs educators could prioritize to fill the gaps.

The NICE Framework does not list any KSAs related to specific tools, programming languages and scripts, or non-technical KSAs. We asked the questions "What programming languages and scripts are most important for your job?" "What soft skills are most important for your job?" and "What skills and topics that we haven't covered so far are most important for your job?" These three open-ended questions were intended to capture additional KSAs that may not be included in the NICE Framework but should still be considered when building curricula. Additionally, we asked, "Was there anything you had learned on the job that you wish you had learned in school?" This question was intended to highlight the KSAs that professionals felt they could have benefitted from as novice or that were particularly difficult to grasp when introduced outside of a classroom setting.

To ensure that no participants quit the study because they felt that any of the questions were too invasive, we told participants that they could skip any questions that they didn't want to answer.

IMPLEMENTATION OF THE BRIEF INTERVIEW

To implement our technique, we interviewed cyber professionals attending the conferences Black Hat USA 2016 and DEF CON 24. Researchers approached fellow conference attendees and asked them if they would be willing to participate in an interview to "help develop better cyber training programs." Researchers read the questions to the participants and then recorded their responses on paper.

Forty-four cybersecurity professionals participated. About half had jobs related to offensive types of cyber defense work such as penetration testing, and the remainder performed more defensive work such as security analysis. On average, participants had been in cyber for 13.79 years (SD = 8.83; range: 1-34 years) and had participated in 3.95 capture-the-flag events (SD = 6.26; range: 0-30). Years in cyber and number of capture-the-flag events were not correlated, $r(42) = -.04$, $p = .775$. Twenty-five participants (57%) had college degrees in a computer science related field, 9 (20%) had degrees in a field unrelated to computer science, and 10 (23%) had no college degrees.

To analyze which of the 32 NICE KSAs were the most important for current cyber work, we performed a series of *t*-test comparing the mean importance rating of each KSA to 3.5 (the middle value of our 1 to 6 Likert scale). KSAs that had a mean importance rating that was significantly higher than 3.5 after Bonferroni correction ($\alpha = .05/32 = .002$) were considered to be of greater-than-neutral importance.

For each NICE KSA, we calculated the percentage of respondents who learned the KSA at each of the 5 locations (job, school, self, other, n/a).

We determined the importance of KSAs that were not included in the NICE framework by counting how frequently they were given as answers to the open-ended questions.

Our results are not detailed here given that this paper focused on the development of our brief interview technique. Our complete results will be described in Jones, Namin, and Armstrong (2017).

LESSONS LEARNED

The success of our brief interview methodology depends on two factors: whether our methodology fit within the constraints laid out by our SME, and whether it provided useful information about cyber work.

Interactions with our participants were generally positive and did not seem to overstep any of the four constraints it was designed to accommodate (no naturalistic observations, anonymity and safety, short data-collection time, and no process questions). Of course, we did not perform nor ask to perform observations, so there was no danger of stepping outside the boundaries of the first constraint.

The demographic questions that we asked our participants fit within our constraint of anonymity. No participants elected to skip the questions about their years in cyber, the number of capture-the-flag events that they participated in, or their educational background. There was indication that our concerns were valid; many participants appeared standoffish about participating until they were assured that we would not ask personal information such as their name. This was particularly true at DEF CON, which is a cash-only conference that does not take names at registration in order to protect anonymity (How much is admission DEF CON, n.d.). A few participants told researchers as an aside that they had been afraid that our study was a social engineering ploy. The fact that we identified ourselves as part of a university research team and did not ask for identifying information mollified their concerns.

The concern that the interviews should be kept short was also validated. When we did experience missing data, it was because participants needed to end the interview early in order to attend a presentation or meet up with other people at the conference. Our interview regularly went over the intended 15-20 minutes, so creating a shorter interview could have cut down on missing data. That said, a large proportion of the participants were willing to spend a great deal of time talking to the researchers, even providing unprompted stories and explanations for their responses throughout the interview. However, it is unclear whether they would have been equally generous in another environment or if we had asked other questions.

Other than some participants needing to end the interview early due to time constraints, we experienced very little missing data. All of the NICE KSA questions were answered by at least 41 of the 44 participants. When participants did skip a question, it was usually because they didn't understand the way the KSAs were worded. Overall, this indicates that participants were comfortable with the depth of the questions.

For the reasons listed above, we found our methodology to be successful in that it allowed us to comply with the four potential research constraints.

We also found that this methodology was reasonably successful at capturing the types of data that would help educators create cyber defense curricula. When deciding what to include in their curricula, educators should prioritize the KSAs that are the most important for cybersecurity professionals. This would help prepare students for the cyber workforce thereby minimizing post-graduate training. With the brief interview technique, we were determined which KSAs are the most important to cyber professionals by analyzing the mean importance ratings for each of the 32 NICE KSAs and by notating which responses to the open-ended questions were given most frequently. In the case of the open-ended responses, further validation of the answers is needed in order to determine their importance in relation to that of the NICE KSAs.

It would additionally be helpful for university educators to know whether students were already receiving instruction on each KSA from school. If they were, curricula developers should not feel obligated to emphasize the KSAs further. We attempted to answer this question by asking participants where they had learned each of the NICE KSAs. The results were difficult to interpret. We found that the NICE KSAs were primarily learned on the job and through self-study, not from school. This could indicate that there is a large need for more higher education programs and classes that focus on these key cyber skills. It is also possible that participants learned about the KSAs at school but then gained a deeper comprehension of them later. This ambiguity is due to how the question was framed; participants were asked where they had learned the *most* about each skill. In addition, many of our participants had been out of school for an extended period of time and/or had not studied in a field related to computer science during their education. Therefore, it was difficult to determine exactly why a KSA was not learned in school and how educators should apply this information.

We propose that in the future the question "Where did you learn [the KSA]?" should be replaced with "How difficult

was it to learn [the KSA]?" This question would be answered with a Likert scale (with 1 indicating "not difficult at all" and 6 indicating "very difficult"). Presumably, cybersecurity educators would want to spend more time in class on the KSAs that were more difficult to learn and to delegate less time for easy-to-learn KSAs. When and whether participants received a computer science education should not affect the difficulty ratings; a KSA that is difficult to learn in a classroom setting should also be difficult to learn through self-study or while on the job.

Overall, we are satisfied with our brief interview technique and with the information it allowed us to measure. This methodology and the data collected with it could be used to develop objectives for courses, group exercises (Patriciu & Furtuna, 2009), and simulation training (Pastor, Diaz, & Castro, 2010) within cybersecurity. Because there is little data to guide personnel selection for cybersecurity jobs (Forsythe, Silva, Stevens-Adams, & Bradshaw, 2013; Mancuso et al., 2014), hiring managers may wish to interview their top performers to determine the KSAs they should prioritize during hiring decisions. The brief interview technique could be customized to fit other research goals in and outside of the cybersecurity field.

A limitation of the brief interview is that it relies on self-report. As such, there may be a mismatch between a KSA's perceived importance and its actual role in cybersecurity work. Ideally, data collected with the brief interview technique will be validated by objective measures of cybersecurity work performance, though it may be some time before such measures are available (Mancuso et al., 2014).

One of the benefits of this methodology is that it requires relatively little time to put together and to execute. Long, extensive work analyses are undoubtedly valuable to the overall effort to meet the challenges of cybersecurity, but they can take even longer given the challenges unique to cyber such as gaining access to and becoming familiar with the cyber domain (Lathrop et al., 2016). This can be prohibitive to researchers who are not already participating in cybersecurity research and certainly slows down the pace of data collection. This brief interview methodology provides a faster turnaround in order to address pressing questions like what to include in cybersecurity curricula.

ACKNOWLEDGMENTS

This work was supported by the National Science Foundation under award number DGE: 1516636.

REFERENCES

Bureau of Labor Statistics, U.S. Department of Labor (2015). *Information Security Analysts*, retrieved from https://www.bls.gov/ooh/computer-and-information-technology/information-security-analysts.htm on March 7, 2017.

Champion, M. A., Rajivan, P., Cooke, N. J., & Jariwala, S. (2012). Team-based cyber defense analysis. *Proceedings of the 2012 IEEE International Multi-Disciplinary Conference on Cognitive Methods in Situation Awareness and Decision Support*, 218-221.

D'Amico, A. & Whitley, K. (2008). The real work of computer network defense analysts. In J. R. Goodall, G. Conti, & K. Ma (Eds): *VizSEC 2007* (pp. 19-37). Springer Berlin Heidelberg.

Dodge, R., Toregas, C., & Hoffman, L. J. (2012). Cybersecurity workforce development directions. In *HAISA*, 1-12.

How much is admission DEF CON, and do you take credit cards? (n.d.). In *Frequently asked questions about DEF CON*. Retrieved from https://www.defcon.org/html/links/dc-faq/dc-fa.html on March 7, 2017.

Forsythe, C., Silva, A., Stevens-Adams, S., & Bradshaw, J. (2013). Human dimension in cyber operations research and development priorities. In *International Conference on Augmented Cognition*, 418-422.

Jones, K. S., Namin, A. S., & Armstrong, M. E. (2017). The Core Cyber-Defense Knowledge, Skills and Abilities (KSAs) That Cybersecurity Students Should Learn in School: Results from Interviews with Cyber Security Professionals. Manuscript in preparation.

Knott, B. A., Mancuso, V. F., Bennett, K., Finomore, V., McNeese, M., McKneely, J. A., & Beecher, M. (2013). Human factors in cyber warfare: Alternative perspectives. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, *57*, 399-403.

Locasto, M. E., Ghosh, A. K., Jajodia, S., & Stavrou, A. (2011). The ephemeral legion: Producing an expert cyber-security work force from thin air. *Communications of the ACM*, *54*, 129-131.

Lathrop, S. D., Trent, S., & Hoffman, R. (2016). Applying human factors research towards cyberspace operations: A practitioner's perspective. *Advances in Human Factors in Cybersecurity*, 281-293.

Madhavan, P., Cooke, N. J., Bass, E. J., Meyer, J., & Roth, E. M. (2015). The National Academies Board on Human Systems Integration panel applying human-systems integration to cyber security. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 405-408.

Mahoney, S., Roth, E., Steinke, K., Pfautz, J., Wu, C., & Farry, M. (2010). A cognitive task analysis for cyber situational awareness. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 279-83

Mancuso, V. F., Christensen, J. C., Cowley, J., Finomore, V., Gonzalez, C., & Knott, B. (2014). Human factors in cyber warfare II: Emerging perspectives. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, *58*, 415-418.

McGettrick, A. (2013). Toward curricular guidelines for cybersecurity: Report of a workshop on cybersecurity education and training. Retrieved from http://www.acm.org/education/TowardCurricularGuidelinesCybersec.pdf on March 7, 2017

National Security Agency (2016). *National Centers of Academic Excellence in Cyber Defense*, retrieved from https://www.nsa.gov/resources/educators/centers-academic-excellence/cyber-defense/ on June 12, 2017.

Oltramari, A., Henshel, D., Cains, M., & Hoffman, B. (2015). Towards a human factors ontology for cyber security. In *STIDS*, 26-33.

Pastor, V., Díaz, G., & Castro, M. (2010). State-of-the-art simulation systems for information security education, training and awareness. In *Education Engineering (EDUCON)*, 1907-1916.

Patriciu, V. V., & Furtuna, A. C. (2009). Guide for designing cyber security exercises. In *Proceedings of the 8th World Scientific and Engineering Academy and Society International Conference on E-Activities and Information Security and Privacy*, 72-177.

Restuccia, D. (2015). Job market intelligence: Cybersecurity jobs, 2015.

Salas, E., Prince, C., Bowers, C. A., Stout, R. J., Oser, R. L., & Cannon-Bowers, J. A. (1999). A methodology for enhancing crew resource management training. Human Factors, 41, 161-172.

Shoemaker, D., Kohnke, A., & Sigler, K. (2016). *A Guide to the National Initiative for Cybersecurity Education (NICE) Cybersecurity Workforce Framework (2.0).* Boca Raton, FL: Taylor & Francis.

Stanard, T., Lewis, W. R., Cox, D. A., Malek, D. A., Klein, J., & Matz, R. (2004). An exploratory qualitative study of computer network attacker cognition. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, *48*, 401-405.

Trent, L. S., Hoffman, R., Leota, T., Frost, C. R., & Gonzalez, M. D. (2016). Cyberspace operations and the people who perform them. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 216-217.

Vieane, A., Funke, G., Gutzwiller, R., Mancuso, V., Sawyer, B., & Wickens, C. (2016). Addressing human factors gaps in cyber defense. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 770-773.