

Identifying Expertise Gaps in Cyber Incident Response: Cyber Defender Needs vs. Technological Development

Megan Nyre-Yu, PhD.
Sandia National Laboratories
mnyreyu@sandia.gov

Abstract

Incident response is an area within cyber defense that is responsible for detecting, mitigating, and preventing threats within a given network. Like other areas of cyber security, incident response is experiencing a shortage of qualified workers which has led to technological development aimed at alleviating labor-related pressures on organizations. A cognitive task analysis was conducted with incident response experts to capture expertise requirements and used an existing construct to help prioritize development of new technology. Findings indicated that current software development incorporates factors such as analyst efficiency and consistency. Gaps were identified regarding communication and team navigation that are inherent to dynamic team environments. This research identified which expertise areas are needed at lower-tier levels of incident response and which of those areas current automation platforms are addressing. These gaps help focus future studies by bridging expertise research to development efforts.

1. Introduction

Incident response (IR) is an area within cyber defense that focuses on detecting, mitigating, and preventing threats for a given network. IR is often considered the first line of defense that includes human analysts working closely with technology to process network data. One issue within this area is hiring and maintaining adequate numbers of human analysts who are qualified to work in these roles.

In order to alleviate the pressure on organizations to continuously hire and train these individuals, technological approaches are being pursued to reduce overall workload on analysts [1-4]. Expertise requirements for incident responders have been collected and published [5], but may not accurately depict the full scope of expertise needs due to the

variety of tasks required in the analyst role, including communication and collaboration.

Technological development is rapid in response to the urgency from industry to mitigate burnout and increase coverage of incident responders. However, it is unclear if current technology adequately addresses needs of analysts, especially given the depth and breadth of knowledge, skills, and abilities (KSAs) needed to do the job effectively. Adoption rates of automation and orchestration platforms remain low compared to other incident response technologies [6] likely due to time and cost investments for deployment. Though business literature hails these platforms as the ultimate solution for struggling firms, little evidence is available that suggests that they have made an impact in operational and organizational performance.

This paper presents a holistic approach to qualifying expertise needs in IR teams, particularly at the lower tiers of analysts. The study presented in this paper includes novel work based on the author's own research. The author conducted cognitive task analysis interviews to collect expertise data that could be analyzed and compared to recent automation-focused technology. The goal was to identify gaps and opportunities in development. The study was part of a larger methodology in the author's dissertation that explored information sharing in IR processes.

2. Background

Security IR teams are comprised of analysts who are often structured into tiers. Tiers represent different levels of investigation and response, and typically correlate with breadth of expertise across incident types and depth of expertise within specific skill areas. Incidents typically start at the lowest tier where a novice analyst will conduct triage investigation, collecting information and compiling it into a ticket. If resolution of an incident is outside the scope of an analyst's role or expertise, the incident is escalated to the next tier. This cycle repeats, adding richness and

depth to the investigation as the analyst tier increases, until the incident is resolved.

2.1. Expertise in cyber incident response

Incident response in cyber security requires applying a wide range KSAs in a decision-making and problem-solving capacity, often under temporal pressure [7]. Classification of KSAs to work roles represents the bulk of building foundational understanding of expertise in this domain. It is also important to understand how expertise is distributed and shared in dynamic work settings across multiple analysts. Previous studies have worked on mapping expertise to tasks or job requirements. Within cyber security specifically, Chen et al. [8] performed different task analyses and linkage analysis to understand expertise gaps in IR. The National Initiative for Cybersecurity Education (NICE) framework also details a comprehensive inventory of knowledge required for different positions in cyber security [5]. Another study produced knowledge maps from surveys to understand how well-equipped incident responders were to perform daily tasks [9]. These examples illustrate how correlating functions with expertise can provide an understanding of knowledge gaps in groups.

Trends across business literature show that one response to a shortage of qualified workers has been developing technology that can augment expertise or eliminate human-in-the-loop tasks with the goal of reducing investigation times and improving consistency [1, 10, 11]. However, there is not strong evidence that work conducted to understand expertise in IR has impacted the development of this new technology.

One potentially useful construct within expertise literature uses a dimensional approach to capture other aspects of expertise beyond the traditional definition of subject matter expertise [12]. The “six dimensions of expertise” construct offers a theoretical foundation for categorizing additional expertise elements (i.e. communication, interface/tool, etc.), but lacks a path to put into practice, particularly for collecting, analyzing, and synthesizing data toward a development goal. The discussion section of this paper presents opportunities for improving the construct and suggests some areas to focus development efforts and help augment targeted expertise in IR.

2.2. Technological trends: helping the analyst

In order to identify new directions in technological development for IR, it is necessary to identify the

current state of the market and how technology is addressing analyst needs at present. The current market for IR software is flooded with potential options that vary in scope, depth, sophistication, and cost. Within this pool, a specific subset of developers has focused on providing platforms with capabilities to automate (perform autonomously) and orchestrate (guide coordinated activities) processes in security organizations.

Security Orchestration, Automation and Response (SOAR) technologies are software platforms that are designed, built, and marketed to increase capacity and efficiency in IR organizations. They focus on integrating existing software within an organization and developing new capabilities to reduce an analyst’s time on an incident. IR analysts must monitor, use, and pivot between a variety of programs and appliances. SOAR technology aims to reduce pivoting by integrating signals from different tools into a single interface. Moreover, the industry has recognized the shortage of analysts, which has resulted in an aim to reduce hands-on time of analysts on menial tasks by automating low-level activities and allowing more time for applying expertise in more difficult tasks.

There are at least a dozen different platforms currently available on the market, though more software companies are aiming to add SOAR capabilities to their existing products and services to compete. The expected growth of SOAR solutions in practice is 15% by 2020 [6], up from 1% in 2018 [13]. This market validation indicates that more firms are recognizing the potential benefits of integration, automation, and orchestration in their security organizations, as well as the need to address labor shortages, data deluge, and disparate tools.

Gartner, Inc. conducted a detailed analysis of SOAR capabilities [10] that identified requirements of what platforms should be able to do to meet industry needs. These include integration across security software solutions, process and workflow guidance, journaling support, case management, and reporting capabilities, to name a few. Many of the recommendations aim to address organizational needs in computer security. Employing new approaches like the one presented in this paper may also be useful in developing system requirements from the users’ perspectives and may help address some underlying causes of the issues felt in the field. This approach evaluated the features developed towards Gartner’s requirements against what IR experts believe is needed to effectively work at the lower tiers of IR.

2.3. Research goals

The goals of the study described in this paper were to identify development gaps in building new technology for cyber defenders based on data collected by the author. First, the author aimed to show breadth of expertise needed in lower tiers of IR by using a dimensional construct to capture non-traditional aspects of expertise from Cognitive Task Analysis (CTA) data. The results act as guidance for what types of training or technology might be useful for this level of IR.

Second, the author compared expertise requirements of incident responders to available technological capabilities that could address those expertise requirements. Specifically, automation and orchestration platforms include features and capabilities that claim to help analysts; the author classified those features using the dimensional expertise construct to identify where technology is focusing development. Results summarize overlaps and gaps in how the latest technologies are addressing expertise requirements of cyber defenders in incident response.

3. Methods

This study included two main methods for data collection and analysis. The author conducted CTA interviews with subject matter experts in IR to identify expertise that helps incident responders perform their daily tasks. A market analysis was conducted to catalog technology being developed to automate IR tasks. Both sets of data were coded and categorized within the dimensional expertise construct, then compared to capture general gaps in alignment between analyst needs and technological solutions.

3.1. Knowledge elicitation from experts

CTA is a subset of methods used to assess the knowledge and cognitive activities needed to perform a particular set of tasks based on subject matter expert experience [14]. Literature indicates that CTA is the best suited for expertise-aimed studies interested in understanding what knowledge is required and how it relates to the overall task structure [14]. CTA is especially appropriate for developing technological solutions to support cognitive processes, and has been performed in cyber security to understand and improve team effectiveness [8], situation awareness [15] and system design [16]. This study used the Applied Cognitive Task Analysis methodology (ACTA) [17] to explore expertise needs in IR teams. ACTA provided a well-structured protocol split into a task analysis, a knowledge audit, and a simulation

interview. ACTA also includes interview prompts for novice researchers and produces an organized set of findings for comparison and additional analysis.

The study included five participants ($N = 5$) with five or more years of experience in IR. While small in sample size qualitative research, including CTA methods, often has a lower number of participants, but produces rich data with high cost-benefit ratio [14, 18-21]. Determination of sample size occurred dynamically based on data saturation after each participant was interviewed. Participants had diverse backgrounds across different sectors (academia, government, industry). The author also notes that this population was extraordinarily difficult to study, especially given security concerns around interacting with individuals outside their respective organizations.

Each interview followed the ACTA format [17] and lasted approximately 90-120 minutes. As CTA methods are most effective with a specific task on which to focus, the author had previously determined information sharing functions (i.e. escalations) as a critical but understudied step in the larger IR process [22]. These functions acted as the main task for investigation using ACTA.

The main output of the ACTA method was a table that compiles the full range of interview responses from participants into a digestible and usable format for informing design; this table is called a cognitive demands table (CDT). Each transcript was used to populate the details of a CDT per participant.

3.2. Dimensions of expertise in IR

The “six dimensions of expertise” construct [12] includes subject matter, communication, information flow path, expert identification, interface/tool, and situational context as unique categories of expertise. This construct was not developed beyond definition of and justification for each dimension at the time of this study. However, it is a generalizable construct across different knowledge-based tasks. Additionally, the construct offered a common ground on which domains or needs can be compared.

The author aimed to apply and further define this construct in the context of security IR. Accordingly, two additional “dimensions” (*policy* and *self-awareness*) were added based on existing IR literature [23-26] and a preceding ethnographic study conducted by the author that investigated incident response organizations [22]. The additions incorporated elements of lower-tier analyst work that require an individual to evaluate organization-based rules and own performance in order to determine next steps in an investigation. A summary of the dimensions of

expertise applied in this research are described in Table 1; the ‘codes’ were used for qualitative analysis.

3.2.1. Interview data analysis. Analysis of interview data employed a top-down qualitative coding scheme based on the dimensions of expertise described in Table 1. In order to increase the trustworthiness of research findings, analysis included two raters (including the author) with a background in qualitative research methods. After a 45-minute training exercise with the codebook, raters independently coded the participant CDTs.

The CDTs were organized in a table format; within each cell, text was broken down into smaller segments. Each segment was a statement copied or summarized from the interviews and acted as a unit for coding. The raters labeled these units using the codebook as a guideline. Units could be coded with more than one category (that is, segments could have any number of codes that applied). Each rater labeled segments using the designation in the codebook (C1-C6) for the original six dimensions of expertise and the two additional codes previously described (C7-C8). When complete, coding results were compiled into a spreadsheet format.

Table 1. Definitions of dimensions of expertise in IR

Dimension of Expertise	Code	Definition
Subject matter	C1	Expertise in a given subject matter area; Usually related to a specific area but can also be general; Pertaining to domain knowledge
Communication	C2	The style used to communicate with someone; tactics for how analysts are approached; vocabulary used to communicate something; using different styles for different people; being receptive of communication
Information flow path	C3	Concerning the method used to contact someone; Knowing which path is the most appropriate for a given person
Expert identification	C4	Knowing who to go to when you need additional knowledge or expertise in a given area; Knowing who to send something to, or who should address a given issue
Interface/tool	C5	User skill in manipulating technological systems; Familiarity with tools and navigating interfaces
Situational context	C6	Knowing the environmental and situational context and how each affects the outcome of an incident
Policy	C7	Institutionalized knowledge regarding security posture; Driven by rules or procedure developed at upper management / company official level

Self-awareness	C8	Driven by understanding of self, including limitations and self-evaluation; meta-cognition
----------------	----	--

3.2.2. Inter-rater reliability (IRR). In qualitative research, one way to define reliability is the extent to which a set of scores is random [27], or how much of the variance is due to variability in participants being scored. In addition to understanding dimensions of expertise in IR, the author evaluated the original expertise construct in practice to identify how it could be strengthened as a tool. To establish trustworthiness in applying the six dimensions of expertise construct, the author included a second rater (in addition to herself) for data analysis and assessed IRR between both raters for *the original six dimensions only*. The author used Cohen’s κ [28], which is a reliability coefficient designed for fully crossed design with exactly two raters. Cohen’s κ includes probability of agreement by chance in addition to rates of agreement.

To compute κ , a contingency table was calculated to compare ratings by code and by rater (Table 2). As shown, the contingency table is a 6 x 6 table in which full agreements between raters were tallied in the diagonal and disagreements were tallied by rater and by code. The κ coefficient for this dataset between two raters was $\kappa = 0.51$, or “fair agreement” [29], but suggests that additional work is needed to understand potential overlap in certain dimensions, particularly subject matter (C1) and situational context (C6). This IRR outcome is further discussed in Section 4.1.

Table 2. Evaluating the original construct: Contingency table of agreement between raters

		Rater 1					
		C1	C2	C3	C4	C5	C6
Rater 2	C1	61	1	1	8	5	26
	C2	0	49	1	0	0	3
	C3	3	3	20	2	0	5
	C4	4	7	1	60	2	4
	C5	6	0	0	3	30	7
	C6	2	3	0	5	0	56

3.3. Technology market analysis

The author used the dimensions of expertise to categorize technological capabilities advertised in the market in 2019. Analysis included SOAR platforms to capture the latest commercially available automation capabilities, which mainly market to security operations organizations and include tiered incident response teams.

SOAR platforms included in this study were selected using a combination of two different

techniques. First, the author used a well-cited report on SOAR technologies [10] to identify some of the platforms to be included in the analysis. The report highlighted 16 different SOAR vendors including in the in-depth analysis of SOAR capabilities. Second, a generic Internet search was conducted for “SOAR, technology, cyber security” to identify other prominent tools that might not have existed at the time or were not included in the report. From these two techniques, nine (9) platforms were chosen for analysis based on feature data availability as not all platform websites offered insights into their features and capabilities. The platforms for this study included Cybersponse, Demisto, Siemplify, Swimlane, Phantom, D3 Soar, LogRhythm, Syncurity, and Resilient.

In order to gain information about each platform, each associated website was evaluated as the main source of data and included technical reports, white papers, and sales information. Collection included a line-by-line capability assessment of each platform; each capability or feature described was recorded. Each feature was then evaluated against the dimensions of expertise to identify which dimension(s) the feature could potentially augment for a human user. For example, if a platform advertised a capability of “codeless playbook creation”, the author tallied the relevant dimensions of expertise as subject matter and interface/tool expertise, as the feature alleviates need to express functions in a specific computer language and overcomes the need to interact with a specialized tool. The author notes that interpretation of the capability may be dependent upon the rater’s familiarity with domain-specific tools and terms; this was a key limitation of the approach and the main reason for the use of a single rater for this activity.

Tallied information was recorded in a large matrix-style table. Sums of tallies for each dimension were calculated across all SOAR platforms to understand total platform capabilities compared to data from experts regarding dimensions of expertise needed in IR.

4. Results

4.1. Lower tier IR expertise requirements

Expertise requirements were elicited from CTA with IR experts about information sharing tasks at lower tier response. Sums of tallies for each code, including policy and self-awareness, are depicted in Figure 1. These counts are indicators of frequency in the expert interviews as interpreted by the two raters

and include partial and full agreement between raters. Due to the nature and scope of interview, the author cannot definitively conclude that frequency indicates “ground truth” importance of dimensions of expertise within IR. However, the frequency may suggest *perception* of importance amongst experts who have deep experience in the field.

Figure 1 shows sums of tallies for both raters across the eight (8) codes used in data analysis. Subject matter expertise (169 tallies) and situational context expertise (161 tallies) were the top two dimensions, followed by expert identification expertise (126 tallies) and communication expertise (91 tallies). Interface or tool expertise (75 tallies), self-awareness (74 tallies), policy (65 tallies), and information flow path (59 tallies) were relatively close in frequency.

Subject matter and situational context expertise were the top two dimensions of expertise indicated by experts as necessary for successfully conducting lower tier incident response activities, especially regarding information sharing tasks. The co-occurrence of these two dimensions indicates that knowing *what signals mean* is but one aspect of IR; knowing *the context* in which that signal occurs is also important to help determine if the signal can be ignored or if it requires action. This is a critical decision point for lower tier analysts who must often act as a filter for determining if something is an incident that requires investigation. The IRR matrix (Table 2) also suggests that further definition and development may be needed in the construct itself to achieve higher IRR. These two dimensions may have significant overlap or even dependency, which could be explored in future studies.

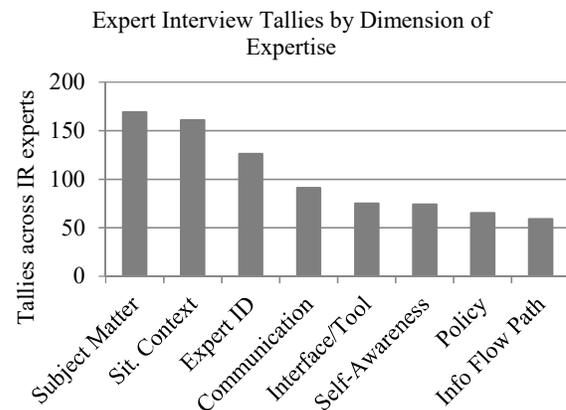


Figure 1. Tallies of interview data for dimensions of expertise across two independent raters

Expert identification and communication expertise were commonly identified by experts as necessary for successfully conducting information sharing tasks, such as handoffs and escalations. These two

dimensions support knowing to whom an incident should be sent for additional response activities, which is often dependent on the recipient's skills and knowledge areas. Accordingly, communication helps effectively transfer pertinent information related to that incident with appropriate level of detail and urgency, which can impact the recipient's available time to respond and prioritization of actions.

Prior research has identified a strong teaming component in IR [7-9, 30, 31] despite the tiered structure seen in many organizations. Several dimensions of expertise support team activities, including communication, expert identification, and self-awareness, all of which are represented in results from the CTA. These results suggest that the software and environment of IR analysts should support these activities, especially for lower tier analysts who may not have developed deep expertise in any given dimension.

4.2. Expertise augmentation by technology

Technological focus of automation and expertise augmentation was elicited through a market analysis of SOAR platforms. Figure 2 shows summarized data from this analysis, depicting sums of tallies for each code. These sums are indicators of frequency in SOAR marketing materials as interpreted by the author. The order of presentation used for Figure 1 was also used in Figure 2 to show ordered alignment of the categorical data between expert interviews and SOAR features.

Among the dimensions of expertise listed in Table 1, the most apparent dimension in SOAR features was situational context expertise. IR analysts need to pivot continuously between screens and platforms in order to gain context about an incident. Thus, developing technology to address these inefficiencies may take higher priority. According to the data, many SOAR platforms focus on bringing the context to the analyst by fetching data from different appliances and displaying them to the user, effectively reducing the need for the analyst to manually retrieve and assemble all relevant information to make a decision.

The next dimension of expertise most apparent in SOAR features was policy. Decisions in IR often rely on knowledge of how an analyst's firm wants to handle particular incidents. That is, not all companies want to respond to every incident in a uniform way, and incident responders at the lowest tier of the organization must know how to act based on that posture. However, experts indicated that interpreting policy can be difficult or even inappropriate for lower tier responders because of inherently low policy expertise. Accordingly, SOAR platforms may try to augment this dimension by having preprogrammed

guidance to provide to the analyst when making decisions about different incidents.

Expert identification was the next most represented dimension of expertise in SOAR features. Analysts are often organized into tiers, which correlate with expertise. Incidents enter the workflow at the lowest level in which analysts begin collecting and compiling relevant information from a suite of tools to support investigation and outcomes. However, not all incidents can be completely resolved at the lowest tier. An analyst might reach a point when they cannot go further, at which point they escalate the incident, pushing the compiled information as well as responsibility for resolution to the next level of responders.

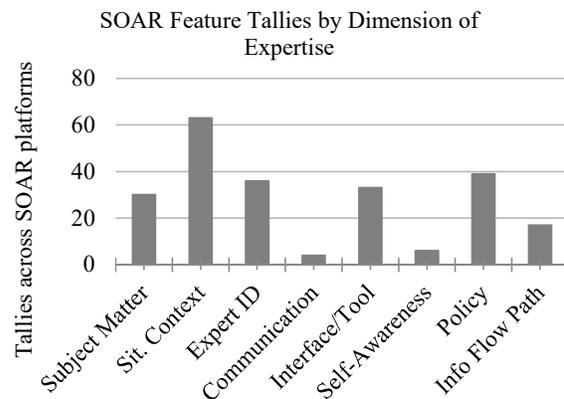


Figure 2. Tallies for dimensions of expertise across orchestration and automation features

One key finding from previous studies [7, 9, 32] and from expert interviews in this paper was that knowing who to send the ticket to, or who to ask for help, was a key piece of navigation within an IR team. Yet, it is not always obvious or reinforced within a given team. SOAR platforms, which are highly customized per organization, claim to provide guidance in support of this dimension of expertise. Augmenting expert identification expertise includes helping analysts determine where information or knowledge might exist, whether it is a person or some other non-human source. Some platforms, in bringing the context to the user, automate this dimension of expertise altogether. Others claim to provide recommendations about who might be able to help with a particular ticket. One trend observed from this dataset was the feature of playbooks or runbooks [4, 10, 11], which create predefined paths for incidents based on different indicators and may include to whom the incident should be escalated.

Another dimension of expertise that was well-represented by both experts and by evidence in SOAR features was interface and tool expertise. As

mentioned, IR analysts must often navigate a suite of tools and even perform sufficiently in multiple scripting languages [33-35]. Each additional tool presents another interface and set of rules that must be mastered to perform the job effectively. SOAR platforms aim to overcome the learning curve of individual tools to allow the analyst to work seamlessly between them with the goal of reducing overall time to respond. This effort is both practical and needed, especially considering the complexity of the environment and software.

Finally, subject matter expertise was represented in materials about SOAR technologies. In fact, SOAR platforms advertised that this dimension of expertise could be embedded into the rules of SOAR protocols [1, 11], essentially modeling decisions from experts. Expertise can come from the SOAR developer itself or from the purchasing firm, effectively using their own expertise to shape the protocols in their customized SOAR platform. This trend of pulling expertise from experts and embedding it into a system is not unlike the development of 'expert systems' [36]. Some of the lessons learned from the evolution of expert systems in other fields may be useful in helping guide SOAR development.

Information flow path expertise was not strongly represented in expert interviews. However, data from SOAR platforms provide evidence that new technologies are augmenting this expertise anyways through playbooks. These playbooks claim to automate information dissemination in the background while providing explicit guidance to the analyst regarding escalating decisions.

5. Discussion

5.1. Dimensional expertise in lower tier IR

The six dimensions of expertise proved to be a useful construct for expanding discussion around understanding and transfer of expertise in cyber security, especially as it related to non-traditional 'dimensions' that capture dynamic and collaborative tasks like IR. Applying and expanding the construct allowed the author to qualitatively evaluate expertise representation within cyber security incident response. The CTA results provide a basic level of understanding of the breadth of expertise needed in this domain beyond widely accepted and expanding subject matter expertise requirements.

Further development of the theoretical construct presents an opportunity for future research to further validate findings and provide a basis for quantitative measurement of expertise. This would especially be

useful if future work compared and integrated other existing knowledge and expertise theories with emphasis on content and construct validity. Integrating existing measurement instruments or developing new instruments that account for correlation or dependencies, such as the co-occurrence of subject matter and situational context expertise, would also be valuable.

The results of this study indicate that the construct was useful as a categorization tool for qualitative comparison. However, even this limited application was not without limitations. The sums of instances per dimension as mentioned by an expert or targeted by a developer is not a robust measure of frequency. The sample sizes for both sets of data are small for a full quantitative comparison, and the SOAR data were rated by only the author due to domain specific knowledge barriers. Content analysis of more platforms and raw interview data is recommended to provide a more representative quantitative outcome.

This study was able to identify a gap between what IR experts believe is a needed dimension of expertise for lower tier responders and what features automation platforms are targeting for development. The gap is evident in the misalignment between expert and technological emphasis around communication expertise and self-awareness, which will be further discussed in the next section.

Considering the current labor shortage in IR, the findings of this study indicate that there are potential avenues for augmenting certain dimensions of expertise to reduce the burden on lower tier analysts. However, the findings also highlight that some dimensions, such as communication expertise, self-awareness, and expert identification, might also be addressed through other methods, such as training and team development. Research to further explore this angle could better inform training, development, and retention strategies for incident responders [37].

5.2. Implications for technology development

Using knowledge elicitation methods to identify needs in design is not a new concept. CTA methods are often used to help determine requirements for new systems [14]. This paper explored how to expand the application of CTA by comparing outputs (from experts) to what the market currently offers in IR systems. This section discusses trends observed in SOAR platforms as well as opportunities for development to address analyst needs.

One common trend catalogued from SOAR features was integration of technologies and showing analysts a unified presentation of data sources and potential paths forward; these were tallied as

'situational context expertise' as they support knowledge pertaining to context and respective potential courses of action. The goal of many features relating to this dimension was to achieve higher analyst efficiency by effectively reducing the time (and interface/tool expertise) needed to retrieve relevant data in support of an investigation. However, evaluation of analyst efficiency with and without these platforms was not feasible due to the low adoption rate of the technology. Furthermore, there was little information available in the marketing and sales information indicating that human-centered methods were used to develop and support development efforts. Situation awareness (SA) literature specifically discusses how systems can support multiple levels of SA in relation to expertise, which can act as guidelines for system development and design [38]. Additional cyber SA literature and ongoing research [39-42] also may help guide application design specifically for cyber security.

A dimension of expertise that was identified as important for lower tier analysts was communication expertise, also identified as interactional expertise within expertise literature [43, 44]. Many SOAR platforms advertised that they could help overcome communication barriers by providing chat features within the tool, as well as the ability to share other artifacts and documents. This aims to not only facilitate collaboration, but also to document the process for an auditable record. However, the definition of communication expertise extends beyond the mode of communication (expressed as information flow path in the original construct) and audibility of analyst interactions. Instead, the CTA results suggested that everyday communication skills (i.e. knowing how to talk to people and interact with them in different situations) are critical to analyst success at lower tier tasks in IR. Results provided little evidence of true augmentation of communication expertise in SOAR platforms; this is one potential avenue for technological exploration.

Results from the CTA showed that self-awareness was a concept identified by experts as important but was relatively unaddressed by SOAR platforms. Self-regulation and reflectiveness are both important aspects of learning [45, 46]. Thus, self-awareness in this context may be an important underlying aspect of building expertise in general and progressing with personal development. This is especially helpful in IR tasks in which an individual must know his or her boundaries, observe and evaluate their own performance, and adjust as they learn. Within the ACTA methodology, provided prompts identify aspects of self-awareness in relation to expertise [47-49]. The author also asserts that, while self-awareness

may be an individual trait, external feedback from the analysts' environments (including peers, systems, and platforms) may contribute to increasing performance indicators signals and trigger opportunities to self-reflect.

The goal of SOAR platforms is to reduce repetitive tasks for humans through automation and increase consistency by guiding response activities through orchestration [1-4]. The desired effect of achieving this goal is decreased labor shortages and some additional level of protection to companies inundated with data, false alarms, and a complicated array of software. However, this strategy is based on improving current operational stability. While the added SOAR capability offers some level of solution to immediate problems, the next steps of the field should progress towards long-term development of cyber security professionals [37, 50], recognizing that the traditional path to becoming an expert in security has fundamentally changed due to the introduction of automation. Furthermore, system developers should consider the role of menial (but fundamental) tasks in analyst development, and the potential effects of system failure coupled with incomplete system understanding. These 'ironies of automation' [51] are critical to identify and mitigate early in design cycles to prevent potentially catastrophic outcomes. The findings presented here indicate that there is no explicit need for an entirely new platform, but expanding design considerations to build expertise, in addition to augmenting it, is warranted.

6. Conclusions

Cyber security incident response is one area within the cyber defense domain currently struggling with labor shortage issues at different levels of experience. Technology development is currently focusing on alleviating this pressure by automating low-level repetitive tasks and providing additional guidance, or augmenting expertise, in lower-tier response activities. However, much of the development has focused on eliminating inefficiencies; available materials about automation in incident response does not support a deep understanding of expertise-driven development.

Expertise has traditionally been associated with deep knowledge within a particular subject matter domain. However, expanding the scope of more holistic expertise constructs can help identify additional areas to focus research, education and development efforts. This paper demonstrates how such a construct can be applied within the cyber incident response domain to capture gaps in how

technology is addressing expertise shortages using automation.

Future work in this area could focus on a variety of areas, including an empirical study of how SOAR platforms minimize expertise gaps while also improving analyst effectiveness. Additional work is also needed to connect findings from studies such as this to existing frameworks, such as NICE, and to educational curricula.

7. Acknowledgments

The author would like to acknowledge her advisor, Dr. Barrett S. Caldwell, for his support in complete this research.

8. References

- [1] C. Brooks, "Security Orchestration , Automation and Response (SOAR) - The Pinnacle For Cognitive Cybersecurity," in *Security Essentials*, ed: AlienVault, 2018.
- [2] Siemplyfy, "The Business Case for SOAR," New York, NY, 2018. [Online]. Available: <https://www.siemplyfy.co/resources/the-business-case-for-soar/>
- [3] VIB and Demisto, "The State of SOAR Report, 2018," 2018. [Online]. Available: <https://go.demisto.com/hubfs/Resources/2018%20SOAR%20Report/SOAR%20Report%202018.pdf>
- [4] C. Bedell, "Definitive Guide to SOAR," CyberEdge Press, Annapolis, MD, 9781948939027, 2019.
- [5] W. Newhouse, S. Keith, B. Scribner, and G. Witte, "National Initiative for Cybersecurity Education (NICE) Cybersecurity Workforce Framework," 2017.
- [6] R. Bhargava, "Gartner SOAR Adoption Rate Prediction: From 1% to 15% by 2020 - Why Should You Care?," in *InfoSec Island*, ed: Wired Business Media, 2018.
- [7] R. Ruefle, A. Dorofee, D. Mundie, A. D. Householder, M. Murray, and S. J. Perl, "Computer Security Incident Response Team Development and Evolution," *IEEE Security & Privacy*, vol. 12, pp. 16-26, 2014.
- [8] T. R. Chen, D. B. Shore, S. J. Zaccaro, R. S. Dalal, L. E. Tetrick, and A. K. Gorab, "An organizational psychology perspective to examining computer security incident response teams," *IEEE Security & Privacy*, vol. 12, pp. 61-67, 2014.
- [9] J. Steinke *et al.*, "Improving Cybersecurity Incident Response Team Effectiveness Using Teams-Based Research," *IEEE Security and Privacy*, vol. 13, pp. 20-29, 2015.
- [10] C. Neiva, C. Lawson, T. Bussa, and G. Sadowski, "Innovation Insight for Security Orchestration, Automation and Response (SOAR)," ed. 2017.
- [11] J. Oltsik, "The evolution of security operations, automation and orchestration," in *CSO by IDG Communications*. New York, NY, 2018.
- [12] S. K. Garrett, B. S. Caldwell, E. C. Harris, and M. C. Gonzalez, "Six dimensions of expertise: a more comprehensive definition of cognitive expertise for team coordination," *Theoretical Issues in Ergonomics Science*, vol. 10, pp. 93-105, 2009.
- [13] A. Chuvakin and A. Barros, "Preparing Your Security Operations for Orchestration and Automation Tools," Gartner, Inc., G00325580, February, 2018. [Online]. Available: <https://www.gartner.com/en/documents/3860563>
- [14] B. Crandall, G. Klein, and R. R. Hoffman, "Working Minds: A Practitioner's Guide to Cognitive Task Analysis," 2006.
- [15] A. D'Amico, B. O'Brien, K. Whitley, D. Tesone, and E. Roth, "Achieving Cyber Defense Situational Awareness: a Cognitive Task Analysis of Information Assurance Analysts," in *Proceedings of the Human Factors and Ergonomics Society 49th Annual Meeting—2005*, 2005, pp. 229-233.
- [16] R. R. Hoffman and S. V. Deal, "Influencing versus Informing Design, Part 1: A Gap Analysis," *Intelligent Systems, IEEE*, vol. 23, 2008.
- [17] L. G. Militello, R. J. B. Hutton, R. M. Pliske, B. J. Knight, G. Klein, and J. Randel, "ACTA Methodology," *Navy Personnel Research and Development Center*, pp. 1-59, 1997.
- [18] W. Zachary, C. M. Z. H. Technologies, B. Crandall, T. Miller, and C. Nemeth, "'Rapidized' Cognitive Task Analysis," *IEEE Intelligent Systems*, vol. 27, pp. 61-66, 2012.
- [19] K. Yates, D. Ed, M. Sullivan, D. Ph, R. Clark, and D. Ed, "Integrated studies on the use of cognitive task analysis to capture surgical expertise for central venous catheter placement and open cricothyrotomy," *AJS*, vol. 203, pp. 76-80, 2012.
- [20] R. Roberts, R. Flin, and J. Cleland, "How to recognise a kick: A cognitive task analysis of drillers' situation awareness during well operations," *Journal of Loss Prevention in the Process Industries*, vol. 43, pp. 503-513, 2016.
- [21] J. Nielsen and T. K. Landauer, "A Mathematical Model of the Finding of Usability Problems," in *Proceedings of the INTERACT '93 and CHI '93 Conference on Human Factors in Computing Systems*, New York, NY, USA, 1993: ACM, pp. 206-213.
- [22] M. M. Nyre-Yu, "Determining System Requirements for Human-Machine Integration in Cyber Security Incident Response," Doctor of Philosophy, School of Industrial Engineering, Purdue University, 2019.
- [23] M. Bishop *et al.*, "Cybersecurity Curricular Guidelines," in *Information Security Education for a Global Digital Society: 10th IFIP WG 11.8 World Conference, WISE 10, Rome, Italy, May 29-31, 2017, Proceedings*, M. Bishop, L. Futcher, N. Miloslavskaya, and M. Theocharidou, Eds., ed. Cham: Springer International Publishing, 2017, pp. 3-13.
- [24] S. Cobb, "Mind This Gap: Criminal Hacking and the Global Cybersecurity Skills Shortage, A Critical Analysis," presented at the Virus Bulletin Conference, 2016.

- [25] S. E. Freed, "Examination of Personality Characteristics among Cybersecurity and Information Technology Professionals," Masters of Science, Psychology, University of Tennessee at Chattanooga, 2014.
- [26] Y. Lee and S.-J. Lee, "An Exploratory Investigation of Factors Affecting Computer Security Incident Response Team Performance," in *AMCIS 2004 Proceedings*, 2004: AIS Electronic Library, pp. 4402-4406.
- [27] B. B. Frey, *There's a stat for that!: what to do & when to do it*. SAGE Publications, Inc., 2016.
- [28] J. Cohen, "A Coefficient of Agreement for Nominal Scales.," *Educational and Psychological Measurement*, vol. 20, 1960.
- [29] J. R. Landis and G. G. Koch, "The Measurement of Observer Agreement for Categorical Data," *Biometrics*, vol. 33, pp. 159-174, 1977.
- [30] L. E. Tetrick *et al.*, "Improving Social Maturity of Cybersecurity Incident Response Teams," p. 298, 2016.
- [31] P. Rajivan and N. J. Cooke, "Information-Pooling Bias in Collaborative Security Incident Correlation Analysis," *Human Factors*, vol. 60, pp. 626-639, 2018.
- [32] G. White and N. Granado, "Developing a Community Cyber Security Incident Response Capability," in *2009 42nd Hawaii International Conference on System Sciences*, 2009: IEEE, pp. 1-9.
- [33] I. A. Tøndel, M. B. Line, and M. G. Jaatun, "Information security incident management: Current practice as reported in the literature," *Computers & Security*, vol. 45, pp. 42-57, 2014.
- [34] A. Silva, G. Emmanuel, J. T. McClain, L. Matzen, and C. Forsythe, "Measuring Expert and Novice Performance Within Computer Security Incident Response Teams," in *Foundations of Augmented Cognition: 9th International Conference, AC 2015, Held as Part of HCI International 2015, Los Angeles, CA, USA, August 2-7, 2015, Proceedings*, D. D. Schmorrow and C. M. Fidopiastis, Eds., ed. Cham: Springer International Publishing, 2015, pp. 144-152.
- [35] R. Werlinger, K. Muldner, K. Hawkey, and K. Beznosov, "Preparation, detection, and analysis: the diagnostic work of IT security incident response," *Information Management & Computer Security*, vol. 18, pp. 26-42, 2010.
- [36] B. G. Buchanan, R. Davis, R. G. Smith, and E. A. Feigenbaum, "Expert systems: A perspective from computer science," in *Cambridge Handbook of Expertise and Expert Performance*, K. A. Ericsson, R. R. Hoffman, and A. Kozbelt, Eds., 2nd ed. Cambridge, UK: Cambridge University Press, 2018.
- [37] M. Nyre-Yu, K. A. Sprehn, and B. S. Caldwell, "Informing Hybrid System Design in Cyber Security Incident Response," Cham, 2019: Springer International Publishing, in *HCI for Cybersecurity, Privacy and Trust*, pp. 325-338.
- [38] M. R. Endsley, "Expertise and Situation Awareness," in *Cambridge Handbook of Expertise and Expert Performance*, K. A. Ericsson, R. R. Hoffman, and A. Kozbelt, Eds., 2nd ed. Cambridge, UK: Cambridge University Press, 2018.
- [39] M. Albanese, H. Cam, and S. Jajodia, "Automated Cyber Situation Awareness Tools and Models for Improving Analyst Performance," in *Cybersecurity Systems for Human Cognition Augmentation*, R. E. Pino, A. Kott, and M. Shevenell, Eds., ed. Cham: Springer International Publishing, 2014, pp. 47-60.
- [40] V. Mancuso, D. Minotra, N. Giacobe, M. McNeese, and M. Tyworth, "idsNETS: An experimental platform to study situation awareness for intrusion detection analysts," in *2012 IEEE International Multi-Disciplinary Conference on Cognitive Methods in Situation Awareness and Decision Support*, 2012, pp. 73-79.
- [41] V. Mancuso, D. Staheli, M. J. Leahy, and M. M. Kalke, "Cloudbreak: Answering the Challenges of Cyber Command and Control," *Lincoln Laboratory Journal*, vol. 22, 2016.
- [42] M. Tyworth, N. A. Giacobe, and V. Mancuso, "Cyber situation awareness as distributed socio-cognitive work," in *SPIE Proceedings*, Baltimore, Maryland, 2012, vol. 8408, in *Cyber Sensing 2012*, pp. 84080F1 - 84080F9.
- [43] H. Collins and R. Evans, "A sociological/philosophical perspective on expertise: The acquisition of expertise through socialization," in *Cambridge Handbook of Expertise and Expert Performance*, K. A. Ericsson, R. R. Hoffman, and A. Kozbelt, Eds., 2nd ed. Cambridge, UK: Cambridge University Press, 2018.
- [44] H. Collins and R. Evans, "The Third Wave of Science Studies: Studies of Expertise and Experience," *Social Studies of Science*, vol. 32, pp. 235-296, 2002.
- [45] B. Breed, "The reflective abilities of expert and novice learners in computer programming," presented at the British Educational Research Association Annual Conference, Edinburgh, UK, 2003.
- [46] P. Ertmer and T. Newby, "The expert learner: Strategic, self-regulated, and reflective," *Instructional Science*, vol. 24, pp. 1-24, 1996.
- [47] M. S. Cohen, J. T. Freeman, and S. Wolf, "Meta-recognition in time-stressed decision making: {Recognizing}, critiquing, and correcting," *Human Factors*, vol. 38, pp. 206-219, 1996.
- [48] R. Glaser and M. T. H. Chi, *The nature of expertise*. Hillsdale, NJ: Lawrence Erlbaum Associates, 1988, pp. xv-xxvii.
- [49] G. A. Klein and R. R. Hoffman, "Seeing the invisible: Perceptual/cognitive aspects of expertise," in *Cognitive science foundations of instruction*, M. Rabinowitz, Ed., ed. Hillsdale, NJ: Lawrence Erlbaum Associates, 1993, pp. 203-226.
- [50] M. Nyre-Yu, R. S. Gutzwiller, and B. S. Caldwell, "Observing Cyber Security Incident Response: Qualitative Themes From Field Research," *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, vol. 63, no. 1, pp. 437-441, 2019.
- [51] L. Bainbridge, "Ironies of automation," *Automatica*, vol. 19, no. 6, pp. 775-779, 1983.