

## ABSTRACT

The Analyst Mindset: A Cognitive Skills Assessment of Digital Forensic Analysts

Chris Sanders, Ed.D.

Mentor: Sandi Cooper, Ph.D.

Despite significant investment in cyber security, the industry is unable to stem the tide of damaging attacks against computer networks. This unfortunate situation is, in part, because cyber security exists in a state of cognitive crisis defined by tacit knowledge and poorly understood processes. At the heart of the crisis are digital forensic analysts that identify and investigate intrusions. Unfortunately, even skilled analysts in these roles are often unable to explain how they go about the process of finding intruders and assessing their foothold on a network. Without this knowledge, professional and academic educators are unable to build a standardized industry-accepted curriculum for the identification and training of new analysts. While there have been some attempts to inventory the skills, processes, and knowledge required to serve in the digital forensic analyst role, no current efforts provide a thorough, research-backed accounting of the profession with consideration for cognitive skill elements.

This problem of practice study details a cognitive skills assessment of the digital forensic analyst profession by leveraging two Cognitive Task Analysis (CTA) research methods. The Simplified Precursor, Action, Result, Interpretation (PARI) method

provided a framework for eliciting procedural skills, and the Critical Decision Method (CDM) supported the discovery of decision-making skills. Using these techniques, interviews conducted with expert analyst practitioners revealed four unique procedural skill categories, characteristics of two significant facets of analyst decision making, and numerous subcategory elements that describe additional dimensions of expert analyst performance. The results converged on a model of diagnostic inquiry that represents the relationships between how analysts formed investigative questions, interpreted evidence, assessed the disposition of events, and chose their next investigative actions. These findings establish explicit knowledge that provides a foundational understanding of how skilled analysts perform investigations. They also lay new groundwork for cyber security's emergence from its cognitive crisis, with implications for educators and practitioners alike.

Copyright © 2021 by Chris Sanders

All rights reserved

The Analyst Mindset: A Cognitive Skills Assessment of Digital Forensic Analysts

by

Chris Sanders, B.S., M.S.

A Dissertation

Approved by the Department of Curriculum and Instruction

---

Brooke Blevins, Ph.D., Chairperson

Submitted to the Graduate Faculty of  
Baylor University in Partial Fulfillment of the  
Requirements for the Degree  
of  
Doctor of Education

Approved by the Dissertation Committee

---

Sandi Cooper, Ph.D., Chairperson

---

Sandra Talbert, Ed.D.

---

Jess Smith, Ph.D.

Accepted by the Graduate School

December 2021

---

J. Larry Lyon, Ph.D., Dean

## TABLE OF CONTENTS

LIST OF FIGURES .....	viii
LIST OF TABLES .....	ix
ACKNOWLEDGMENTS .....	x
CHAPTER ONE .....	1
Introduction to the Problem of Practice .....	1
Introduction .....	1
Statement of the Problem .....	2
Purpose of the Study .....	4
Philosophical Assumptions .....	5
Research Design .....	6
Definition of Key Terms .....	7
Conclusion .....	11
CHAPTER TWO .....	12
Literature Review .....	12
Introduction .....	12
Dependence and Vulnerability .....	12
The Role of the Digital Forensic Analyst .....	15
The Analyst Skills Gap .....	20
Prior Work Towards Investigative Skills Assessment .....	23
Digital Forensics as a Unique Domain .....	24
Federally Developed Frameworks .....	25
Academic Curriculum Development and Research .....	29
Industry-Led Attempts .....	32
Conclusion: Common Themes Among Existing Skill Assessments .....	34
Conclusion .....	34
CHAPTER THREE .....	36
Methodology .....	36
Introduction .....	36
Researcher Perspective and Positionality .....	37
Theoretical Framework .....	39

Research Design and Rationale.....	41
Simplified PARI.....	43
CDM .....	44
Results Integration .....	45
Site Selection and Participant Sampling .....	46
Site Selection.....	46
Analyst Sampling .....	47
Data Collection Procedures .....	49
Simplified PARI for Procedural Skills.....	49
CDM for Decision-Making Skills.....	51
General Data Collection Prerequisites and Privacy Considerations .....	53
Data Analysis Procedures .....	55
Preliminary Analysis Steps .....	55
Simplified PARI.....	55
CDM .....	56
Relating Procedural and Decision-Making Skills .....	58
Reliability and Validity .....	59
Ethical Considerations .....	59
Limitations and Delimitations.....	60
Conclusion.....	61
CHAPTER FOUR.....	62
Results and Implications.....	62
Introduction .....	62
Analyst Demographics .....	62
Procedural Skills .....	63
Inquiry Skills.....	65
Evidentiary Skills .....	79
Anomaly Detection .....	85
Network Mapping and Attack Visualization.....	92
Decision-Making Skills.....	95
Decision Cues.....	96
Decision Goals .....	98
A Model of Diagnostic Inquiry .....	105
Implications.....	108

Summary and Conclusion .....	114
CHAPTER FIVE .....	116
Distribution of Findings .....	116
Executive Summary .....	116
Overview of the Data Collection and Analysis Procedures .....	117
Summary of Key Findings .....	119
Informed Recommendations .....	120
Findings Distribution Proposal .....	122
Target Audience .....	122
Proposed Distribution Method and Venue .....	123
Distribution Materials .....	123
Conclusion.....	124
APPENDIX A.....	126
Analyst Screening Survey.....	126
APPENDIX B.....	129
PARI Investigation Scenario .....	129
APPENDIX C.....	130
Consent Form.....	130
APPENDIX D.....	134
Directed Analysis Techniques .....	134
BIBLIOGRAPHY.....	145

## LIST OF FIGURES

<i>Figure 1.1.</i> CTA Analysis distinguishes between knowledge and skills.....	6
<i>Figure 2.1.</i> The six common analyst roles include three primary investigation roles and three supporting roles.....	19
<i>Figure 2.2.</i> The relationship between digital forensic knowledge areas and other computing disciplines .....	24
<i>Figure 2.3.</i> The NICE framework brings logical structure to cyber security roles .....	26
<i>Figure 3.1.</i> Research design using the Simplified PARI and CDM methods for CTA .....	42
<i>Figure 3.2.</i> The relationship between procedures and decisions provides a framework for analyst-driven investigations.....	58
<i>Figure 4.1.</i> Analysts ask preceding and succeeding questions to establish causal relationships and identify other generally suspicious activity relative to another event.....	71
<i>Figure 4.2.</i> Analysts often ask proximate questions when seeking events that correlate with the timing of already known suspicious events .....	72
<i>Figure 4.3.</i> Analysts based capability matching questions on characteristics drawn from internal sources, external sources, and analogs from prior investigations.....	74
<i>Figure 4.4.</i> Analysts formed utility questions to collect information they needed to answer other investigative questions .....	76
<i>Figure 4.5.</i> Analysts excluded known benign to reduce their data set so that they can scrutinize the remaining data for anomalies .....	89
<i>Figure 4.6.</i> Analysts use a process of diagnostic inquiry to reduce uncertainty and make decisions during investigations .....	106

## LIST OF TABLES

Table 3.1. <i>PARI Elements and Associated Interview Questions</i> .....	51
Table 3.2. <i>CDM Deepening Phase Probe Questions</i> .....	54

## ACKNOWLEDGMENTS

I would like to thank the people who helped make this document possible and contributed to the positive step forward it represents. First and foremost, thank you to my wife Ellen, who I kept awake countless nights by storming into the bedroom rambling on and on about the ideas running through my head following late classes.

Nobody becomes a scholar alone, and I was fortunate to have several amazing people on this journey with me. I want to extend my gratitude to my doctoral colleagues who made this whole experience more enjoyable. I also want to thank my instructors at Baylor who shepherded me along this scholarly experience, with special thanks to my advisor, Dr. Sandi Cooper. I don't fit the mold of a typical education student, and I appreciate all of you opening up your mind to learn from me as I did from you.

I want to pay special tribute to my students, whose success helps motivate me, including anyone who has ever taken one of my classes, read one of my books, or sat in on one of my conference presentations. Additionally, I want to thank my colleagues that served as sounding boards and provided feedback on my ideas.

This whole project started over a decade ago when I was a struggling young analyst trying to learn the craft. Someone told me that you are either born with the skills needed to do this job, or you are not. I thought that was nonsense, and I have spent the rest of my career gathering the knowledge and data to prove it. The document you are about to read is a step along that path. I don't remember the name of the person who told me that, but I want to thank them too.

## CHAPTER ONE

### Introduction to the Problem of Practice

#### *Introduction*

A piece of sensitive or identifying information may traverse a dozen or more globally distributed computer systems every time someone orders a chicken sandwich, buys lumber, clicks the thumbs up button on a friend's social media page, or pays a bill. Each system represents an opportunity for an attacker to steal information or disrupt the availability of service delivery. Despite the efforts of the cyber security industry, attacks against the systems people rely on every day outpace the ability to defend them.

Digital forensic analysts are at the front lines of these battles, entrusted with the immense responsibility to detect adversarial attempts to compromise the computer networks that make the world work. Despite the importance of the analyst role, much of the knowledge about how to perform investigative duties remains tacit. Even skilled practitioners are not effective at describing how they perform their duties. Meanwhile, academia struggles to identify the core competencies of the digital investigation disciplines. Difficulties identifying and developing digital forensic analysts have led cyber security to a state of cognitive crisis that is at the heart of its inability to stem the tide of attacks.

The cyber security industry and digital forensic analyst role must move beyond tacit knowledge into more explicit knowledge about the job function. By leveraging skilled practitioners and academic methods, this study constitutes a research-based appraisal of investigative skills allowing for more robust and thorough curricula that meet

the needs of the evolving profession. Through a cognitive task analysis of skilled practitioners, experiences from the digital trenches encapsulate the procedural and decision-making skills that make up the analyst's role.

### *Statement of the Problem*

With dependence on technology comes vulnerability. Each system an individual interacts with has the potential to disrupt their lives to varying degrees should they experience a compromise of confidentiality, integrity, or availability. Breaches at Capital One and Equifax demonstrate a risk to personal financial systems through the theft of credit card and social security numbers (Krebs, 2017; McLean, 2019). Ransomware attacks on local hospitals in Henderson, KY, Tuscaloosa, AL, and elsewhere demonstrate how the interruption in the availability of networks halts critical healthcare services (Andone, 2019; Smith, 2016). At an even larger scale, critical infrastructure attacks against the Ukrainian power grid and an Iranian nuclear plant demonstrate potentially devastating kinetic impacts of cyber-based attacks (Krebs, 2014; Zetter, 2016).

Attacks like the ones mentioned here occur at an increased pace with more dramatic consequences as technology dependence grows. The number of security breaches has gone up 67% in the past five years, with the annual cost of cybercrime up 72% over the same period (Ponemon, 2019). Despite spending an anticipated \$124 billion on cyber security products in 2019 (Aitken, 2018) and increases to the analyst workforce (Bureau of Labor Statistics, 2018), the rate and effectiveness of attacks appear to outpace the industry's ability to defend the information system we depend on.

The inability to stem the tide of attacks is, in part, an education problem. Society depends on digital forensic analysts to detect and respond to adversarial actions against

computer networks before significant breaches occur. When intrusion detection sensors generate alarms, analysts investigate them using multiple sources of evidence, hoping to fend off intruders before they achieve their goals.

Unfortunately, there is little consistency among practicing analysts in how they perform their craft. Analysts often lack metacognitive awareness and cannot describe their process for connecting the dots (Sanders & Rand, 2019). While all analysts rely on data acquisition and parsing tools, there is no standard loadout or classification system for evaluating tools. Every analyst relies on evidence, but the types of evidence available to analysts vary dramatically from organization to organization. Despite efforts, the analyst role rests on a foundation of unwritten tacit information (Sundaramurthy, 2014).

The lack of explicit knowledge surrounding the analyst roles makes it difficult to identify and train new analysts. Most organizations rely on on-the-job training by pairing aspiring analysts with experienced veterans. However, this form of teaching alone is not expedient or entirely effective. When experienced analysts cannot describe how they perform a task themselves, they have little hope of effectively teaching that skill to someone new to the field.

At the same time, academia suffers from an inability to produce job-ready analysts. Cyber security programs are relatively new on campuses, and many industry experts report that universities do not produce job-ready graduates across various security disciplines. According to Downs (2017), 84% of employers believe half or fewer cyber security applicants are qualified for the position. Part of the issue stems from difficulties that universities have determining which skills and technologies to teach in their programs. No universally accepted cyber security or digital forensics curriculum

standards exist that provide a thorough grounding in security theory while instilling practical hands-on investigative skills.

Solving the cognitive crisis affecting cyber security begins with attacking the education problem and better defining universal curriculum standards for digital forensic analysts. This requires a thorough accounting of processes used by analysts to complete their jobs, the skills used during completion of those processes, and the tools that enable and simplify the processes and skills.

### *Purpose of the Study*

The purpose of this cognitive task analysis study is to elicit cognitive skills from digital forensic analysts to provide an accounting of procedural and decision-making skills used during investigations. By identifying these items, the cyber security industry moves towards a better understanding of archetypical professional analyst roles. This research meets these goals by asking two primary research questions:

1. What procedural skills do experienced digital forensic analysts use during investigations?
2. What decision-making skills do experienced digital forensic analysts use during investigations?

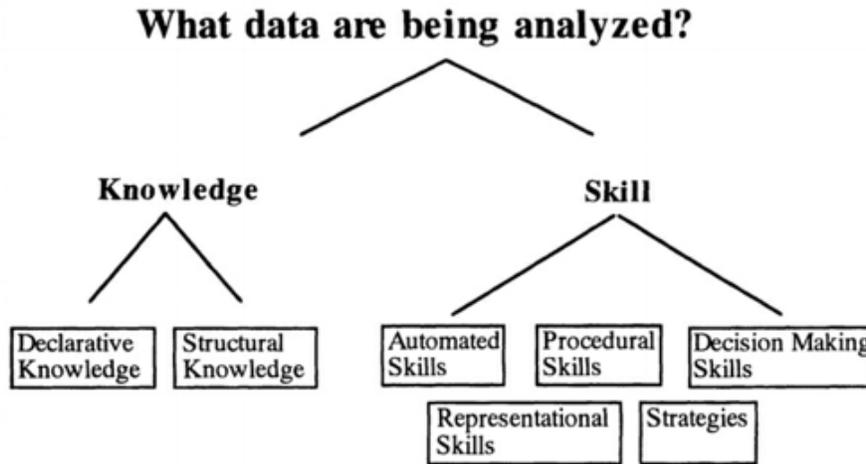
These real-world findings help build a better framework for identifying and developing digital forensic analysts. In academic settings, this information provides educators with critical insight around which to construct curriculum. In practice, organizations should use these findings to develop and enhance their internal training programs and set road markers for analyst career development.

### *Philosophical Assumptions*

Professional mastery requires practitioners to acquire and leverage multiple mental models through which they view a portion of the world. These models define how practitioners classify entities, define relationships, and use tools to interact with data to produce meaningful results. Understanding how to teach or learn the digital investigative disciplines requires knowledge of both direct observable job tasks and underlying cognitive processes. While traditional task analysis meets the needs of the former, a more thorough approach is needed to encapsulate both observable and underlying aspects of expertise. This study leverages cognitive task analysis (CTA) as a theoretical framework for discovering both the observable and cognitive tasks tied to digital forensic analyst roles.

Cognitive task analysis (CTA) is “the extension of traditional task analysis techniques to yield information about the knowledge, thought processes, and goal structures that underlie observable task performance” (Schraagen et al., 2000, p. 3). The methods for completing CTA are wide-ranging. Because of the new and rapidly evolving state of cyber security, this research focused on observational, interview-based, and conceptual CTA with particular attention paid to consistencies among experts.

CTA involves five steps: collection of preliminary knowledge, identification of knowledge representations, knowledge elicitation, analysis of data, and presentation of results (Clark et al., 2007). This process seeks to define declarative and structural knowledge (referred to in this study as “knowledge”) and procedural skills (referred to in this study as “skills”). The CTA knowledge elicitation and analysis methods and practical role of the knowledge define this distinction, as represented in Figure 1.1 (Seamster & Redding, 2017).



*Figure 1.1.* CTA Analysis distinguishes between knowledge and skills. Reprinted with permission from “Applied cognitive task analysis in aviation,” by Seamster and Redding, 2017, p. 91. Reprinted with permission.

This problem of practice focuses on the identification of procedural and decision-making skills scaffolding digital forensic practice. While other knowledge and skills are important, these areas have an essential role in educating future practitioners and feasibility of research mechanisms for knowledge elicitation.

### *Research Design*

Cognitive task analysis (CTA) is “a variety of interview and observation strategies that capture a description of the explicit and implicit knowledge that experts use to perform complex tasks” (Schraagen et al., 2000, p. 3). I adapted the Simplified Precursor, Action, Interpretation, Result (PARI) and Critical Decision Method (CDM) CTA methods to identify the procedural and decision-making skills, respectively. For the CTA interviews, I recruited expert analysts with significant experience in the field and spent the majority of their time conducting investigations. These 9 analysts represent

multiple investigative disciplines, industries, and organizations. I collected data through remote, online interviews.

The Simplified PARI and CDM methods provided a framework for qualitative interviews and data analysis. Using the PARI technique, I conducted interviews with expert analysts framed through an investigation scenario lens. They worked through the scenario aloud, while I asked probing questions to identify facets of the procedures they used. With the CDM technique, the expert analysts recounted a particularly challenging investigation they conducted. I asked probing questions designed to highlight and expand on decisions made during the investigation and posed hypothetical scenarios to elicit unique dimensions of their decision-making skills. Using the collected interview data, I analyzed and consolidated the identified skills and their cognitive context within and across cases. The output of this effort was a unified list of investigative procedural and decision-making skills.

#### *Definition of Key Terms*

*Alert:* A notification that an event or series of events has occurred. This study employs “alert” interchangeably with “alarm.”

*Analyst:* “A person who analyzes digital evidence and investigates computer security incidents to derive useful information in support of system/network vulnerability mitigation” (Newhouse et al., 2017, p. 32). This study employs “analyst” interchangeably with “investigator.”

*Anomaly:* Any evidence that deviates from what an analyst perceives as normal or expected.

*Attacker:* A person who attempts to gain unauthorized access to a computer network.

*Attack Surface*: The attack vectors available to an attacker at given points of presence in the network.

*Availability*: “Ensuring timely and reliable access to and use of information” (Nieles et al., 2017, p. 2).

*Breach*: An event that negatively affects the confidentiality, integrity, or availability of a computer network or data stored on it.

*Cognitive Task*: “Cognitive structures and processes that undergird expertise” (Seamster & Redding, 2017, p. 5).

*Cognitive Task Analysis*: “A variety of interview and observation strategies that capture a description of the explicit and implicit knowledge that experts use to perform complex tasks” (Schraagen et al., 2000, p. 3).

*Computer Network*: Two or more computers connected to share resources.

*Computer Science*: “The study of computers and computing, including their theoretical and algorithmic foundations, hardware and software, and their uses for processing information” (Tucker & Belford, 2019).

*Compromise*: An event where an attacker successfully gains unauthorized access to a computer network.

*Confidentiality*: “Preserving authorized restrictions on information access and disclosure including means for protecting personal privacy and proprietary information” (Nieles et al., 2017, p. 2–3).

*Cue*: Any stimulus with implications for action (Wong, 2004).

*Cyber Security*: “The protection of information and information systems from unauthorized access, use, disclosure, disruption, modification, or destruction in order to ensure confidentiality, integrity, and availability” (Nieles et al., 2017, p. 2).

*Decision-Making Skill*: The algorithms, heuristics and rules of thumb used for deciding among alternative choices. (Seamster & Redding, 2017).

*Digital Forensics*: “The preservation, collection, validation, identification, analysis, interpretation, documentation, and presentation of digital evidence derived from digital sources to facilitate or further the reconstruction of events found to be criminal, or helping to anticipate unauthorized actions shown to be disruptive to planned operations” (Palmer, 2001, p. 16).

*Event*: Any observable occurrence on a computer network. Events usually indicate the formation, change, or dissolution of a relationship between entities.

*Evidence*: Something that serves as proof that an event occurred. In digital forensics, this is usually a log entry, file, packet, software output, or other electronic data.

*Host*: A computer system connected to a network.

*Industrial Control Systems (ICS)*: A catch-all term used to describe the collection of systems responsible for the automation and control of industrial processes (Hayden et al., 2014)

*Integrity*: “Guarding against improper information modification or destruction and ensuring information non-repudiation and authenticity” (Nieles et al., 2017, p. 4).

*Internet of Things (IoT)*: A term referring to the networked interconnection of everyday objects like coffee makers, smart watches, doorbells, and thermostats (Xia et al., 2012).

*Intrusion Detection System (IDS):* A type of security software designed to monitor systems or networks for anomalies and generate alerts.

*Investigation:* The systematic inquiry and examination of evidence to gain an accurate perception of whether a compromise has occurred, and to what extent.

*Malware:* Computer software used for malicious purposes.

*Pattern Matching:* The process by which analysts compare characteristics of an event to those of a familiar model to identify similarities.

*Procedural Skill:* A series of steps performed to accomplish a goal (Konoske & Ellis, 1986).

*Ransomware:* Malware designed to deny access to a computer system, data, or other resource until a ransom is paid.

*Security Operation Center (SOC):* A physical facility or logical division of an organization that houses the information security team tasked with protecting computer network and information assets.

*Sensor:* A hardware or software device that logs data about the environment.

*Threat:* “Any circumstance or event with the potential to adversely impact organizational operations through an information system via unauthorized access, destruction, disclosure, or modification of information, and/or denial of service” (Joint Task Force Transformation Initiative, 2012).

*Threat Actor:* An entity responsible for actions that negatively impact or threaten to impact the confidentiality, integrity, or availability of a computer network.

### *Conclusion*

The future of society depends on information networks and their ability to process data. Digital forensic analysts are the front-line defenders protecting those systems against breaches of confidentiality, integrity, and availability. Despite the importance of the analyst role, existing research fails to thoroughly define the practical and cognitive skills necessary to perform the job well. The research presented in this problem of practice changes that narrative through a cognitive task analysis of digital forensic analysts. This effort pushes the industry toward a standard analyst curriculum that strengthens the state of the profession. The next section highlights a gap in the research literature demonstrating the need for this study and a closer look at digital forensic analyst competencies.

## CHAPTER TWO

### Literature Review

#### *Introduction*

Cyber security is still a new and evolving field that requires a diverse and complex skillset. An array of technical reasons can explain why most computer security breaches occur, but humans lie at the heart of discovering and investigating those breaches. While the field is developing, practitioners and educators still lack a comprehensive understanding of the skills required to perform the duties of digital forensic analysts. This literature review highlights our dependence and vulnerability on technology, describe the role of the digital forensic analyst, and highlight the lack of thorough research focused on identifying practical and cognitive analyst skills. The following review of current scholarship argues that there is a gap in the field's knowledge of digital forensic analyst skills, exhibiting a need for a cognitive skills assessment of the professional analyst role.

#### *Dependence and Vulnerability*

Computer technology is interwoven into the fabric of everything humans do. In 2016, 89% of Americans reported having at least one computer in their household and 82% reported having an internet subscription (Ryan, 2018). At the same time, the number of internet-connected devices grew to 16 billion in 2016; more than twice the human population of Earth. Of these devices, traditional personal computers and tablets only make up about 10% of the total, with mobile and internet protocol-based phones making

up 54%. The remainder comprises the various Internet of Things (IoT) devices like smart watches, internet-enabled appliances, home security cameras, and digital thermostats (Ericsson IoT Forecast, 2018).

Every time someone interacts with a computer to do something as simple as checking their bank balance or purchasing lunch, various bits of information about them traverse dozens of networks and systems distributed globally. Each system their data passes through achieve a specific goal, but they also represent some form of vulnerability that is susceptible to attack. Cyber security attacks impact confidentiality, integrity, and availability of systems and data. Confidentiality refers to protecting restrictions on information access and disclosure, such as safeguarding user credit card information or passwords. Integrity protects data from unauthorized alteration, like preventing modification to figures in an accounting database. Availability refers to protecting reliable and timely access to data, which typically equates to keeping systems running and protecting data against malicious restriction (Nieles et al., 2017).

Attacks against computer networks are growing in number, scope, and impact. Ponemon (2019) found that the number of security breaches has gone up 67% in the past five years, with the annual cost of cybercrime up 72% over the same period. Numerous examples illustrate the growing threat of information systems breaches and how they affect individuals. In December 2013, Target announced that attackers infiltrated their systems, installed malware on their point-of-sale network, and stole guest payment and credit card data from millions of customers (Krebs, 2013). This breach of confidentiality would eventually grow to encompass 40 million stolen credit card numbers and 70 million stolen names, addresses, e-mail addresses, and phone numbers of Target shoppers

(Krebs, 2014). Many Target shoppers became victims of identity theft upon the discovery of the stolen information from the breach online (Krebs, 2013).

Starting in 2016, the rise of cryptocurrencies like Bitcoin facilitated the resurgence of ransomware. This type of malware infects systems and encrypts user data until a victim pays the ransom fee to the malware operator (O’Gorman & McDonald, 2012). Ransomware attacks are particularly damaging when they target businesses that are critical to public safety and health. For example, in October 2019, a strain of ransomware hit three hospitals in Alabama. The malware encrypted patient records, preventing medical providers from accessing them digitally. As a result, the facilities had to divert non-critical patients elsewhere (Andone, 2019). Other hospitals have experienced similar attacks, representing a tangible threat to individuals seeking or undergoing medical care (Smith, 2016).

The potential impact of security breaches grows stronger when considering attacks against operational technology (OT) and industrial control systems (ICS) that facilitate utility services. The attacks against the Ukrainian power grid in December 2015 represent the potential for disruption of utilities resulting from cyber-based attacks (Zetter, 2016). Another attack against energy infrastructure occurred in 2010 when analysts discovered malware eventually named Stuxnet on systems responsible for controlling nuclear centrifuges in Iran. The malware was ultimately responsible for physically destroying almost one-fifth of the countries centrifuges by causing them to spin in a disruptive manner (Kushner, 2013). This attack and others serve as examples of how computer-based breaches could negatively impact lives on a large scale.

Commensurate with rising cyber attacks like the ones described here, defense spending has also increased. Global spending on cyber security products and services reached \$101.54 billion in 2017 and some estimates expect it to grow by over 22% to \$124 billion in 2019 (Aitken, 2018). While the cyber security industry spends more on defense, that investment does not appear to be slowing down the number, scope, or impact of attacks against computer networks.

### *The Role of the Digital Forensic Analyst*

While information system owners strive to protect their systems from attacker intrusion completely, cybercrime and espionage are not entirely preventable. The complexity of modern systems exposes plenty of avenues of attack that can often yield success for opportunistic and highly structured adversaries alike. Opportunistic attackers often take advantage of complex networks by looking for simple oversights such as accidental misconfigurations. Structured adversaries will do the same, but also have the capability to overwhelm simple systems with complexity in the form of custom-written malware or zero-day exploits that evades some protective defenses. In both cases, attackers often leverage human vulnerability to trick system users into opening e-mail attachments or clicking links that can lead to system compromise (Sanders & Smith, 2013).

A prevention-only approach to computer network defense is incomplete because prevention eventually fails. Bejtlich (2004) reasoned this concept as, “If at least some intruders are smarter than you and their ways are unpredictable, they will find a way to penetrate your defenses. This means that at some point your preventative measures will fail” (p. 44). Network security strategy must include detective controls in the same way

that many businesses leverage security cameras along with physical locks to protect their facilities. While preventive controls deter some attackers, detective controls can help spot those who are initially successful with enough time to stop their progress before severe damage is done (Cavusoglu et al., 2004).

When detective controls such as intrusion detection systems (IDS) locate anomalies, human analysts investigate to confirm whether the alert indicates adversarial activity. Palmer (2001) describes the process of computer forensics as:

The preservation, collection, validation, identification, analysis, interpretation, documentation, and presentation of digital evidence derived from digital sources for the purpose of facilitating or furthering the reconstruction of events found to be criminal, or helping to anticipate unauthorized actions shown to be disruptive to planned operations. (p. 1)

For the digital forensic analyst, the investigation process centers on uncovering a timeline of events. By systematically examining evidence sources, the analyst moves toward an accurate perception of whether a compromise has occurred, and to what extent (Sanders, 2016a).

At least six investigative roles characterize the digital forensic analyst profession. While each role leverages differing data sets and tools (with some overlap), they all serve a similar purpose to build or support the building of a timeline of events that transpired in order to uncover malicious activity on computer networks. Organizations often define these roles by evidence type, industry sector, or function in the incident response process. The computer incident response process involves four functional (not preparatory or review) steps: identification (sometimes called detection and analysis), containment, eradication, and recovery (Cichonski et al., 2012). Some analysts focus exclusively within a single role, while others gravitate towards two or more of them. Analysts taking

on multiple roles may do so based on personal interest, organization structure, or budget limitations.

In the first role, triage analysts perform event handling by monitoring alerting queues for IDS output that indicates the presence of anomalous behaviors or relationships on the network. Triage analysts serve as the first line of detective controls for the network and will often triage anomalous events for further investigation. After confirming malicious activity, triage analysts typically hand the investigation off to a forensic examiner or incident responder.

Second, forensic examiner analysts collect and analyze evidence from computers based on some suspicion of wrongdoing. They may begin these investigations from a triage analyst's referral, a crime resulting in warrant for seizure of digital assets, an internal human resources case, or other mechanisms. Forensic examiners build timelines from malicious events that transpired on computers to assess impact so that organizations may seek justice through the legal system, prevent further loss from ongoing compromises, and protect sensitive data and operations. While a triage analyst often hands cases off immediately following identification of malicious activities, forensic examiners dig further into the malicious events in effort to build a more complete timeline.

Third, incident response analysts work with confirmed malicious activity to begin the process of identifying the scope of the incident. From there, responders work to contain the intrusion and eradicate the attacker. The responder achieves their goal after restoring the network to operational status and using intelligence gathered during the breach to strengthen defense posture in hopes of preventing similar breaches in the future.

The incident response role often fully encompasses the forensic examiner role in practice, but often uniquely includes the containment, eradication, and recovery functions of incident response.

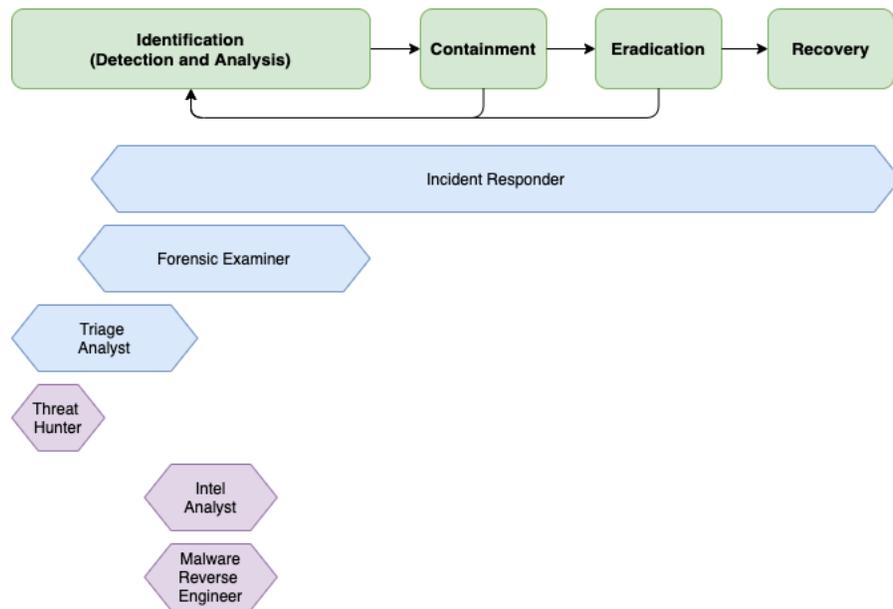
Fourth, threat hunter analysts fill in gaps left by automated signature-based detection tools, which are fallible. Threat hunters perform the human-centric process of proactively searching through networks for evidence of attacks that evade existing security monitoring tools. They are focused primarily on attack discovery and may hand-off suspicious findings to other analyst roles.

Fifth, malware reverse engineer analysts focus on breaking down and understanding the function of malware, which so many breaches rely on. The output of the analysis process allows event handlers and incident responders the ability to leverage gathered knowledge of malware capabilities in their investigations.

Lastly, threat intelligence analysts focus on the characteristics of attackers involved in a breach. Attackers rely on their network infrastructure to facilitate adversarial actions against a target. A threat intelligence analyst leverages community attack and reputation databases, along with other forms of intelligence data, to establish relationships between various aspects of attacker infrastructure. This information informs event handlers and incident responders during their investigations

While all six analyst roles seek to find or characterize intruders in some way, the triage analyst, forensic examiner, and incident responder roles are most directly concerned with building investigation timelines directly. Therefore, I refer to these as primary investigator roles in this study. Even though these roles include some unique skills, the core investigative skills are a consistent thread running through their fabric.

The other three roles are investigation supporting roles. For example, malware reverse engineers and intelligence analysts focus on narrow facets of the investigation; files and tradecraft respectively. Their work enhances the broader investigation by providing context that may aid the uncovering of malicious relationships. At the same time, threat hunters wade through data to find anomalies that launch an investigation, but their primary workflow does not begin with evidence of potential compromise. Figure 2.1 depicts the typical focus of the six analyst roles with respect to the incident identification and response process.



*Figure 2.1.* The six common analyst roles include three primary investigation roles and three supporting roles.

Examining the impactful breach scenarios described earlier highlights the role of analysts in safeguarding society. Analysts discovered and investigated the Target point-of-sale compromise, hospital ransomware cases, and Stuxnet attacks. In each of these cases, the analyst prevented the attack from having broader impacts, but the attackers

accomplished at least some of their goals. Discovery and remediation arriving far too late are a familiar story in security breaches.

Complexity plagues the analyst role. The cyber security domain is broad and rapidly changing with frequent emerging problems that are ill-defined. Analysts must also contend with an overwhelming number of input signals, numerous evidence sources, and incomplete evidence coverage among protected systems. Most analysts continually wrestle with incomplete information that is acquired incrementally over time (Sanders, 2017).

Analyst morale suffers because of the demanding nature of the day-to-day work. Many analysts experience some form of burnout in their job leading to poor work performance or transition to another cyber security career track. Burnout often stems from lack of organizational support, lack of automation for mundane tasks, poor collective operational efficiency, and other factors relating to the broad and complex nature of the role (Sundaramurthy et al., 2015).

In conclusion, the complexity and stress factors described here all work together to present a professional landscape that is cognitively challenging. The importance of the analyst role warrants that researchers and practitioners seek a better understanding of the causes and effects of cognitive demands on analyst practitioners. That understanding must include a thorough inventory of the skills and processes comprising the analyst job role.

### *The Analyst Skills Gap*

Despite the importance of the analyst function, the mechanics of the role remain poorly defined even amongst those who do it every day. In an ethnographic study of

Security Operation Centers (SOCs) where analysts generally work, Sundaramurthy et al. (2014) found that:

SOC analysts often perform sophisticated investigations, and the process required to connect the dots is unclear even to analysts. Incident response isn't just a technical problem; it involves people with various skills interacting in a closed culture, using specific workflows for each incident type. Current solutions aren't informed by these workflows and are only partially helpful. An analyst's job is highly dynamic and requires dealing with constantly evolving threats. Doing the job is more art than science. (p. 55)

High levels of metacognitive awareness among practitioners characterize well-defined fields (Hargrove & Nietfeld, 2015). Individual contributors understand how they think about problems, allowing them to regulate their cognition towards consciously applying specific thinking styles and strategies at appropriate times. Sanders and Rand (2019) found that a lack of metacognitive awareness was prevalent amongst analysts:

While most analysts were able to respond to specific investigative scenarios reasonably, they could not extrapolate on a structured or deliberate investigation process without referencing real-world scenarios. Analysts were able to apply heuristics they had previously developed to the scenarios indicating inductive reasoning had been at work to create "rules of thumb," while deductive reasoning was used to apply them to the current scenario. There were no signs that analysts recognized these processes were occurring. (p. 17)

Expert analysts cannot describe the knowledge that makes them successful at conducting investigations, so it should come as no surprise that organizations struggle to identify and train analysts. Because so much of the required knowledge is tacit, new analysts must learn primarily through observation-based on-the-job training (Sundaramurthy, 2014).

Inconsistent learning experiences amongst practitioners are responsible for a portion of the reported skills shortage in the cyber security industry. "Fifty-three percent of organizations report a problematic shortage of cyber security skills" (Oltsik, 2019, para. 2), and Cyber Security Ventures (Morgan, 2017) estimates that "there will be 3.5

million unfilled cyber security job openings by 2021” (para. 2). Similarly, the ISC2 Cybersecurity Workforce Study (2020) found that 64% of cyber security managers’ report some staffing shortage and 56% of managers’ report that this shortage puts their organization at risk. Worse, Oltsik (2020) reports that 45% of cyber security professionals believe this skills shortage has gotten worse over the past few years. At first glance, the skills shortage may appear as a lack of people to fill roles. However, job board data suggests the problem might lie elsewhere. In a search of open cyber security job roles on Monster.com, there were approximately 55,000 open jobs designated for experienced or senior-level analysts using the searches “senior forensic analyst,” “senior cyber security analyst,” “senior SOC analyst,” and “senior threat analyst.” At the same time, there were less than 4,000 open jobs designated for inexperienced or junior-level practitioners using the same search strings other than replacing the word senior with junior (Monster Jobs, 2019). These results may contain some overlap. Organizations cannot train people for the skills they need for positions they have, limiting their ability to fund slots for inexperienced staff and invest in their career development. However, these organizations cannot overcome a skills gap when their skilled practitioners fail to identify and articulate the necessary skills.

One might expect academia to meet the needs of organizations by producing more job-ready cyber security graduates. However, contention characterizes the relationship between academia and the cyber security industry. Most academic research gains little traction in the security community where practitioners view academic research as dated or irrelevant. Because many of the most experienced practitioners joined the industry when it was brand new, no college degree programs in the field existed. As a result, many

of these individuals did not go to college at all or completed degrees in unrelated fields. At the same time, academic researchers often lack the access to analysts or data needed to perform rigorous research (Lindner, 2010). Outsiders have difficulty gaining trust in SOC environments, partly due to the sensitive nature of the job. Analysts generally have access to corporate secrets and personally identifiable information, requiring a degree of caution when interacting with outsiders. The stressful and time-sensitive nature of the analyst role often leaves analysts little time to work with researchers who are seen as distant outsiders to the analyst's goals (Sundaramurthy et al., 2014).

University cyber security programs are seeing tremendous enrollment interest due to the job demand, broad focus on STEM careers, and general appeal of the industry (Bell & Oudshoorn, 2018; Foresman, 2019). However, information security is still new for many universities and they struggle to identify qualified faculty and curriculum standards. Universities attempting to base curriculum standards on research-based digital forensic skill assessments find the availability of this research lacking (Bell & Oudshoorn, 2018).

### *Prior Work Towards Investigative Skills Assessment*

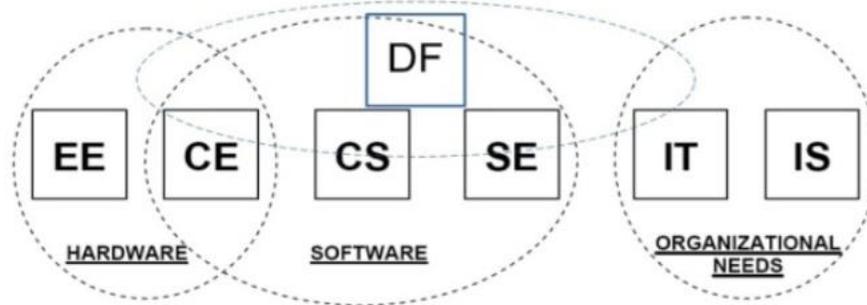
The cognitive demands on analysts and the gap between practitioners and academic researchers have yielded little progress towards understanding the components that comprise the digital forensic analyst discipline. Attempts to characterize digital forensic analyst skills have not trodden deeply enough to produce tangible output useful for professional trainers or academic educators and curriculum planners. This section describes why digital forensics constitutes a distinct knowledge domain and highlights why prior attempts to inventory analyst skills from government, academic, and

professional sources fall short. This section concludes by highlighting common themes among incomplete attempts to assess analyst skills.

### *Digital Forensics as a Unique Domain*

Educators frequently view digital forensics as a subordinate to computer science. Students took the earliest forensic courses as computer science electives. As the scope of forensic skills expanded, so too did the elective courses into dedicated tracks (Liu, 2006). Computer science professors with primarily software or systems development backgrounds often developed curriculum for security courses without any formal education or experience in forensic practice (Lang, 2014).

Digital forensics is a unique discipline. Cooper et al. (2010) mapped knowledge areas related to electrical engineering, computer engineering, computer science, software engineering, information technology, and information systems. Their study (represented in Figure 2) found that digital forensics shared some common knowledge areas with most other computing disciplines, but also brought a substantial set of unique knowledge areas.



*Figure 2.2.* The relationship between digital forensic knowledge areas and other computing disciplines. Reprinted with permission from “Towards standards in digital forensics education,” by Cooper et al., 2010, Proceedings of the 2010 ITiCSE working group reports, p. 89.

Digital forensics is not solely bound to traditional technology and engineering. Forensics relies on human-centric practices such as law, criminology, ethics, psychology, and management (Cooper et al., 2010). A digital forensic curriculum bound only to computing knowledge areas is bound to be incomplete.

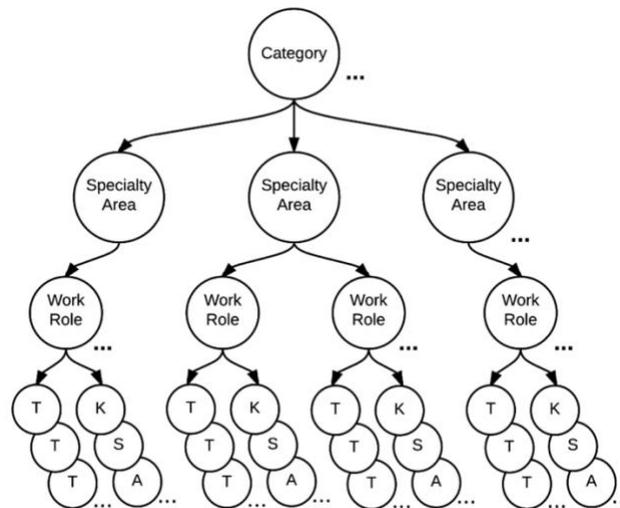
### *Federally Developed Frameworks*

The United States government identifies cyber security as a national defense and workforce priority. In 2010, the Congressional Research Service designated cyberspace as a priority domain for defense spending (Daggett, 2010). Around the same time, the United States Department of Defense established US Cyber Command as a sub-unified command tasked with conducting “full spectrum military cyberspace operations” (“U.S. Cyber Command History,” n.d., para. 7). President Obama elevated Cyber Command to a unified combatant command by signing the National Defense Authorization Act (NDAA) in 2016. Shortly after that, President Trump formally directed the Secretary of Education to increase access to high-quality Science, Technology, Engineering, and Math (STEM) education (Obama, 2017). STEM fields include computer science and cyber security. In 2019, President Trump signed an executive order designed to strengthen America’s cyber security workforce (Executive Order No. 13870, 2019). Through these executive, defense, and education actions along with others, contributes made a few attempts to perform a skills analysis of digital forensic practitioners to pave the way for curriculum standards.

The most impressive of these federal initiatives, conducted by the National Institute for Standards and Technology (NIST), produced the National Initiative for Cybersecurity Education (NICE) workforce framework. The NICE framework “provides

a common lexicon that academic institutions can use to develop cybersecurity curricula that better prepares students for current and anticipated cybersecurity workforce needs” (Newhouse et al., 2017, p. 2). Executive Order 13870 encouraged the adoption of NICE to strengthen the cyber security workforce (Executive Order No. 13870, 2019).

NICE defines individual cyber security work roles and maps them to tasks, knowledge areas, skills, and abilities (represented in Figure 3). For example, NICE describes the Cyber Defense Forensics Analyst role as someone who “Analyzes digital evidence and investigates computer security incidents to derive useful information in support of system/network vulnerability mitigation” (Newhouse et al., 2017, p. 122). This role maps to 39 tasks, 46 knowledge areas, 22 skills, and 2 abilities.



*Figure 2.3.* The NICE framework brings logical structure to cyber security roles. Reprinted with permission from “National initiative for cybersecurity education (NICE) cybersecurity workforce framework,” by Newhouse et al., 2017, NIST SP.800-181, p. 6.

This mapping constitutes the most thorough analysis of the forensic investigator role, but the lack of specificity leaves many unanswered questions for educators striving to teach these skills. For example, NICE maps Task ID T0027 to the digital forensic

analyst role, stating that an analyst should “Conduct analysis of log files, evidence, and other information to determine best methods for identifying the perpetrator(s) of a network intrusion” (Newhouse et al., 2017, p. 25). While accurate, the framework does not define what evidence sources are applicable, the tools used to access and parse them, standards for interpreting them, analytic techniques used to ask and answer investigative questions, the scenarios where the evidence should be consulted, and procedures encapsulating all these things. While NICE is an excellent starting point, the ambiguity surrounding unanswered questions like the ones tied to T0027 highlight where the framework falls short for educators.

The Department of Labor (DoL) developed another framework called the Cybersecurity Competency Model (CCM), which “defines the latest skill and knowledge requirements needed by individuals whose activities impact the security of their organization’s cyberspace” (Cybersecurity Industry Competency Model, 2019, para. 2). The model incorporates competencies defined by NICE in a tiered format that includes (from top to bottom, more general to more specialized) personal effectiveness competencies, academic competencies, workplace competencies, industry-wide technical competencies, and industry-sector functional areas (Cybersecurity Industry Competency Model, 2019). CCM achieves the goal of further mapping NICE competencies to job sector roles along with more general employee competencies. However, this model does not make up for the lack of procedural depth tied to digital forensic analyst workflows.

Another federally developed framework experiencing wide adoption is the National IA Education and Training Programs (NIETP) Center of Academic Excellence in Cyber Defense (CAE-CD) accreditation standard jointly sponsored by the National

Security Agency (NSA) and Department of Homeland Security (DHS). The goal of the CAE-CD program is to “reduce vulnerability in our national information infrastructure by promoting higher education and research in cyber defense and producing professionals with cyber defense expertise” (National Centers of Academic Excellence, 2019, para. 2). These goals aim to address cyber security skills shortages by producing more job ready professionals.

The CAE-CD provides a set of academic standards for associate, bachelor, masters, and doctoral programs. To achieve this designation, institutions must provide courses that encompass specific knowledge units designated by the NSA and DHS. These knowledge units separate into four categories: foundational, technical core, non-technical core, and optional. Each knowledge unit specifies outcomes and topics that institutions should teach, but these areas are left broad and open to the interpretation of the institution (National Centers of Academic Excellence, 2019). For example, within the digital forensics knowledge unit, the CAE-CD designation requires that students of institutions be able to “Discuss the rules, laws, policies, and procedures that affect digital forensics” (Information Assurance Directorate at the National Security Agency, 2019, p. 50). However, the specific rules, laws, policies, and procedures are not listed. Furthermore, the standard provides no guidance describing what processes for which these rules and laws are applicable. While the CAE-CD guidelines provide a useful set of signposts to consider when designing course offerings and high-level requirements within a cyber security degree program, they do not provide enough specific guidance to identify learning objectives within these classes adequately. The CAE-CD standards also fail to recognize the primary investigative roles as a distinct cyber security learning domain,

instead choosing to separate them based on evidence specialty (Information Assurance Directorate at the National Security Agency, 2019).

The federal initiatives defining cyber security skills mentioned here provide a useful step along the path to a comprehensive skills inventory and meaningful curriculum. However, they do not provide a complete picture detailing effective digital forensic analyst skills on their own. In some cases, they also fail to uniquely define the digital forensic analyst role as a unique knowledge area.

#### *Academic Curriculum Development and Research*

As demand for cyber security degree programs rises, universities have made attempts to study digital forensic practitioner skills to aid curriculum development. In one such case, Pennsylvania State University sought to introduce more formal security courses into their computer science and information technology degree programs. University curriculum designers first identified major themes by examining five other university cyber security curriculums. These themes informed surveys of recent graduates and IT professionals to produce a list of concepts deemed critical for the security curriculum. While involving practitioners was a step in the right direction, faculty conducted e-mail-based surveys at a very high level. The tabulated findings merely provide a broad set of key topics and technologies for cyber security practitioners without any focus on individual job roles, including the digital forensic analyst role (Bogolea & Wijekumar, 2004).

In another study, faculty at Marymount University sought to introduce graduate cyber security specialties by assessing practitioner skills. However, rather than performing the assessment themselves, they chose to take queues from several existing

sources. These sources include NICE and NIETP whose limitations we discussed earlier in this literature review, along with input from industry professional certification bodies. While these certification bodies often provide a greater degree of specificity than many of the frameworks already mentioned here, the marketability of the accompanying certification and training classes often drives curriculum rather than skills critical to the everyday practice of the profession (Bicak et al., 2015). The Marymount University process for augmenting cyber security curriculum incorporates useful sources to define curriculum focus areas but falls short of defining specific teaching objectives for those classes.

In one of the more thorough university-based approaches, faculty at the University of Illinois at Urbana-Champaign performed a skills analysis of digital forensic practitioners to create an introductory and advanced digital forensic course as part of a certificate program. This approach centered on a workshop including digital forensic researchers, educators, and professionals with the explicit goal to facilitate dialog about curriculum standards. Eventually, the university established a practitioner-led curriculum committee to define the knowledge taught in the courses (Lang, 2014).

Lang intentionally defines portability as a goal of this curriculum, but this represents a weakness in the overall design. Because of the notable difficulty finding qualified digital forensic faculty, this curriculum targets existing computer science professors for instruction (Lang, 2014). As discussed earlier in this literature review, digital forensics has some overlap with computer science but is not a subset of it or wholly contained by it. The curriculum, as presented, leans too far towards a theoretical

approach rather than a practical one. A thorough skills assessment and the curriculum based on it must treat the digital forensic profession as a standalone discipline.

While this curriculum defines an extensive breadth and depth of forensic analyst skills, Lang (2014) acknowledges:

This program will not be a job-track training program intended to prepare students to directly enter the job market as digital forensic examiners and analysts. Instead, it will provide a broadly applicable education in the field of digital forensics that will be valuable for students going into many disciplines related to digital forensics, such as law, in addition to forensic examiners and analysts. It is expected that these students will receive additional education specific to their career paths and some on-the-job training specific to their eventual professional roles. (p. 6)

The goal of a digital forensic skills assessment and resulting courses should include producing job-ready graduates to meet market demands. Programs failing this charge further pollute the field of job applicants without necessary skills and force organizations to invest in less-than-ideal on-the-job training.

The Forensic Science Education Programs Accreditation Commission (FEPAC) represents a non-university specific collaborative academic effort to establish digital forensic education standards. The purpose of this commission is to “develop and to maintain standards and to administer an accreditation program that recognizes and distinguishes high-quality undergraduate and graduate forensic science programs” (“FEPAC.ORG,” n.d., para. 1). Similar to the CAE-CD accreditation, FEPAC defines a series of topics that course offerings must cover to achieve accreditation. These topics cover the entirety of forensic science, not just digital forensics. The required topics include items such as “Files systems and operating systems” or “Computer networking.” These knowledge areas are incredibly broad and the standard identifies no further components or skills tied to them.

While universities have collectively sought methods for formalizing digital forensic skill inventories, there have been few attempts made by individual academic researchers, and none thoroughly focused on detailed task analysis for curriculum development. For example, Ellis (2009) used CTA methods to assess analyst skills. However, the narrow focus on user interface design resulted in only a simplified set of procedure diagrams lacking meaningful depth.

The academic attempts to define analyst skills and derive curriculum from them represent a sample of many similar attempts. In every case the author reviewed, institutions suffered from similar issues defining the digital forensic role uniquely and to a degree of specificity meaningful to the practice. In each of these programs, a loose set of guidelines dictate high-level topics while the universities relied entirely on individual instructors to define specific learning objectives encompassing common analyst procedures and decisions. Institutions cannot consistently produce job-ready digital forensic analysts without detailed practitioner-informed learning objectives.

### *Industry-Led Attempts*

The cyber security industry also makes attempts to fill the education gap in the digital forensic analyst workforce. However, financial goals that sit in contrast to educational goals define these attempts. Professional training companies consider skill assessment data and curriculum proprietary intellectual property. Therefore, skill assessment and curriculum development processes are rarely transparent or published. It becomes impossible to assess the efficacy of skill assessment processes under these circumstances.

Additionally, private training companies most often target existing practitioners rather than individuals aspiring to enter the digital forensic field (Cybrary Catalog, 2019; Pluralsight Cyber Security Courses, 2019; SANS Cyber Security Courses, 2019). Existing practitioners more frequently receive training reimbursement from their employers (Filkins, 2016), allowing training vendors to charge higher prices. As a result, training companies are often more concerned with skill augmentation rather than holistic analyst education.

Professional cyber security training companies often direct most of their marketing and sales activities toward the sale of standalone courses rather than complete programs. This approach lends itself to focus on individual technologies and specific emerging trends that artificially segment knowledge (Cybrary Catalog, 2019; Pluralsight Cyber Security Courses, 2019; SANS Cyber Security Courses, 2019). The segmentation of knowledge is apparent in the current training landscape's division of analyst training courses by evidence specialty. Well-rounded analysts must interact with a wide array of evidence in any given investigation. However, the focused promotion of ad-hoc courses or evidence specialization tracks leave analysts lacking broad investigative knowledge that prepares them for any primary investigative role. Forgoing a holistic approach fills the immediate needs of employers wishing to augment the skills of existing practitioners but lacks appropriate theoretical grounding for developing new analysts.

While private industry training stepped up to fill the gap provided by traditional academic institutions, the profit-driven nature of these organizations creates an environment of secrecy and knowledge hoarding rather than one of sharing and openness.

Thus, academic institutions struggle to benefit from the curricular innovation of the private sector.

*Conclusion: Common Themes Among Existing Skill Assessments*

Thus far, the federal government, academia, and private industry have failed to produce a rigorous, thorough, open accounting of digital forensic analyst skills applicable to curriculum development. The limitations of existing approaches highlighted in this literature review represent three primary themes.

First, these skills assessments and resulting curriculums are not based on proven theoretical frameworks. Individuals and committees take an ad-hoc approach that varies tremendously. Second, these efforts' output lacks the breadth and depth necessary to encapsulate the critical procedures and decisions associated with analyst duties. Educators cannot produce a detailed or consistent curriculum without a high degree of specificity in learning objectives derived from analyst tasks. Finally, many attempts to inventory analyst knowledge fail to recognize the investigative skill set as a primary role within the cyber security discipline. Instead, educators artificially divide digital forensics using evidence specializations for top-level knowledge domains. Digital forensic analysts must possess skills in every evidence domain, placing general investigative skills as a higher-level knowledge domain with evidence areas falling underneath it. While each of these studies provides useful context for the analyst role, additional work is needed to identify cognitive and functional components of the analyst's skillset.

*Conclusion*

Cyber security exists in a state of cognitive crisis. The inability to identify and educate skilled practitioners partially explains the industry's inability to stem the tide of

attacks despite significant investment in these areas. While there have been many attempts to inventory the skills, processes, and knowledge required to serve in the digital forensic analyst role, no current efforts provide a thorough, research-backed accounting of the procedural and decision-based facets of the profession. This literature review highlighted several of these attempts and identified common themes describing their deficiencies. The identified research gap demonstrates the need for a skills assessment of the digital forensic analyst profession. The next chapter outlines an approach toward solving this problem through a cognitive skills analysis of the digital forensic analyst role. This research is necessary to more adequately train digital forensic practitioners, advance the cyber security profession, and protect society from digital threats.

## CHAPTER THREE

### Methodology

#### *Introduction*

The previous chapter highlighted the importance of the digital forensic role and the lack of a thorough accounting of the procedural and decision-making skills that comprise the practice. This deficiency is partially responsible for a skills shortage, which has led to the industry's inability to defend against crippling attacks despite increased spending. By identifying these skills, academia and industry position themselves to better educate aspiring analysts toward a job-ready state. With that need established, this chapter describes the research methodology I used to elicit procedural and decision-making digital forensic skills. This study's primary research questions are:

1. What procedural skills do experienced digital forensic analysts use during investigations?
2. What decision-making skills do experienced digital forensic analysts use during investigations?

My research employs a qualitative design, so I begin this chapter by describing my perspective as the researcher and a statement on my subjectivity. There, I establish my role as a practitioner in the field and how it has influenced my perception of the analyst profession. Next, I describe cognitive task analysis (CTA) as a theoretical framework underpinning my research and define a CTA-based research design leveraging the Simplified Precursor, Action, Interpretation, Results (PARI) and Critical Decision Method (CDM) knowledge elicitation techniques.

After providing the background and overview of my research design, I describe my decisions regarding site selection, participant sampling, data collection, and data analysis pursuant to answering my research questions empirically. Finally, I describe the ethical considerations and limitations of the chosen methodology.

### *Researcher Perspective and Positionality*

This research employs a qualitative approach to identify and analyze cognitive skills of digital forensic analysts. Denzin and Lincoln (2011) describe qualitative research as “a situated activity that locates the observer in the world” (p. 3). For this cognitive task analysis, I was that observer within the world of cyber security. This world is familiar to me, having spent nearly twenty years working in information technology and cyber security. During this time, I watched cyber security evolve from a small set of tasks handled by systems administrators to a unique discipline encompassing several specialties.

During one of my first dedicated security analyst roles, I asked an experienced analyst, “What makes someone good at this job?” They pompously told me that it was something you must be born with—either you have it or you do not. Did I have it? I was not yet sure, but I did know that this analyst’s views clashed with my own epistemic beliefs about the nature of knowledge. Namely, this idea contradicted my belief that people can learn most things given the right environment, circumstances, and motivation absent externally limiting factors. Much of the rest of my career would focus on understanding the practice and nature of security analysis and how aspiring practitioners learn the craft.

While I am a technologist, I am also an educator, and these perspectives influence my worldview uniquely. I have written several books and articles about cyber security and actively teach aspiring analysts how to leverage their intellect to build investigation timelines and expose network intruders. These diverse experiences shaped my philosophical assumptions about this research toward an eclectic worldview encompassing a diverse range of motivators and outcome goals. My perspective is primarily pragmatic, with a focus on solving real-world problems. Creswell and Poth (2016) describe pragmatism as a perspective focused “on the outcomes of the research—the action, situations, and consequences of inquiry—rather than the antecedent conditions” (p. 26). This cognitive task analysis sought to inductively appraise expert analysts’ reality so that educators may help new analysts better bridge their novice perspective to that reality.

At the same time, this research also represents a transformative worldview serving two distinct causes. First, analysts conduct forensic analysis to reveal criminal acts pursuant to justice. Second, education is a means of transcending the circumstances one is born into. This upward social mobility is particularly accessible with technical education that has the power to end generational poverty through high paying jobs that people may work in locations far away from traditional centers of industry. A greater ability to perform, teach, and learn computer forensic analysis provides transformative societal benefits in both instances. This idea aligns with Mertens (2003) description of the transformative worldview that knowledge should help people improve society. Cyber security professionals’ work inherently safeguards society, and more accessible cyber security education helps extend that transformative profession to more people.

While certain biological dispositions like fluid intelligence and sizeable working memory capacity may make analysis tasks come easier to some than others, an absence of these gifts does not predispose everyone else from making positive impacts within the field. If you can become an effective analyst without inborn cognitive talents, as I believe, then an early step towards enabling this potential for more people must include an accounting of the cognitive skills they must learn and be taught. The combination of these eclectic worldviews, my epistemic beliefs about knowledge, and a hunger for justice for those who may not obtain it for themselves compelled my research.

### *Theoretical Framework*

If someone were to observe a digital forensic analyst performing their job, they would not see much more than a person sitting at a computer while the screen transitions through innumerable forms of evidence. Investigative analysis is a skilled performance whose theater exists primarily in the mind. To ascertain the factors that make the performance impressive is not merely a product of direct observation, but one of careful knowledge elicitation through intentional measures. Cognitive task analysis enables this undertaking by providing a series of mechanisms that “capture what people are thinking about, what they are paying attention to, the strategies they are using to make decisions or detect problems, what they are trying to accomplish, and what they know about the way a process works” (Crandall et al., 2016, p. 9–10). CTA enables researchers to peer behind the curtain of skilled performance to its underlying components.

While domain experts may design university courses and private training, these individuals do not usually capture much of the expertise their work represents completely. Researchers have found significant differences between expert descriptions

of their actions and direct observation of expert performance, with some analyses finding omission rates of up to 70% (Chao & Salvendy, 1994; Clark, 2009; Feldon, 2004). In an attempt to remedy these shortcomings, researchers have employed CTA-based approaches for instructional design since the 1980s (Reigeluth, 1983). Tofel-Grehl and Feldon (2013) performed a meta-analysis of several CTA-based training programs and found that “the gains measured here do robustly support the claim that CTA-based training is more effective than training not based on CTA” (p. 299). This analysis included representation from academia, the military, and multiple private industries. The use of CTA methods allows me to identify and articulate the unseen cognitive procedural and decision-making skills used by digital forensic analysts performing investigations.

Among many forms of CTA, Seamster and Redding (2016) describe a skills-based CTA framework where “a cognitive skill includes the content, organization, and mental manipulation essential for good or superior performance” (p. 136). This approach narrows CTA’s emphasis on reaching operationally relevant results by using practical methods geared toward job readiness training, consistent with a pragmatic approach. These methods focus primarily on knowledge elicitation through a variety of techniques designed to pair with the type of skills the researcher seeks to identify. For this research, I use the pairing of the Simplified PARI technique for collecting and analyzing data about procedural skills and the CDM technique for collecting and analyzing data about decision-making skills.

Simplified PARI and CDM provide frameworks for qualitative, interview-based data collection through scenario and critical incident-based interviews. The semi-structured inquiry prescribed by these techniques compels subjects to describe facets of

their thinking that may not be readily apparent otherwise (Seamster & Redding, 2016). This exercise in making thinking visible allowed me to ask additional questions to probe the characteristics of the task the subject described and yielded rich descriptive narratives.

CTA techniques provide a mechanism for initial and post-collection analysis of collected data. With both Simplified PARI and CDM, the interviews provide an initial degree of analysis throughout the use of deepening strategies and probing questions. The subject is not merely answering questions; they are active participants in identifying and describing the skills they have used to achieve success. After these interviews, many of the cognitive skills are already evident, allowing me to focus on refinement and cross-case analysis (Crandall et al., 2006). CTA provides a rigorous and efficient set of tools to collect and analyze information about cognitive skills and synthesize those findings into practical structures useful for educators.

### *Research Design and Rationale*

I conducted this study by using cognitive task analysis techniques based on a qualitative research paradigm. Schraagen et al. (2000) described cognitive task analysis as “a variety of interview and observation strategies that capture a description of the explicit and implicit knowledge that experts use to perform complex tasks” (p. 3). CTA aligns with a qualitative research paradigm due to its focus on the exploration of an experience that should “make sense of, or interpret, phenomena in terms of the meanings people bring to them” (Denzin & Lincoln, 2011, p. 3). Creswell (2013) describes qualitative research as a process that involves both deductive and inductive data analysis geared towards establishing themes and patterns. For this study, these themes and



### *Simplified PARI*

I used the PARI knowledge elicitation methodology to identify procedural skills. Researchers developed the PARI methodology as a framework to elicit procedural skills from complex real-world work environments (Hall et al., 1995). Seamster and Redding (2017) slimmed PARI down into the Simplified PARI technique to be more time-efficient with an eye toward curriculum development and task analysis conducted for operational environments. In Simplified PARI analysis, a domain expert designs a problem and poses it to another expert who works through and solves the problem while providing a concurrent verbal report of their actions. Throughout this process, the interviewer poses questions to elicit information surrounding the subject's precursors, actions, results, and interpretations. The researcher examines interview transcripts to further identify these elements through a process of decomposition. The action statements identify tangible procedures, while precursor and interpretation statements identify cognitive procedure skill elements (Seamster & Redding, 2017). Simplified PARI allows the researcher to serve the role of domain expert and interviewer, provided their experience justifies both roles. I assumed both roles in this study based on my lengthy experience in digital forensic analysis.

Simplified PARI was an appropriate research technique for this study because it provided a scenario-based framework for qualitative data collection. This framework provided the basis for identifying and analyzing seemingly complex practitioner skills. These methods facilitated the identification of procedural tasks for curriculum development and domain-specific education.

Researchers have used PARI collection and analysis to identify procedural skills in multiple fields, including computer game development (Pillay et al., 1999), medicine

(Riggle et al., 2013; Garcia, 2015), and instruction (Depeda-McZeal, 2014). While these fields are unique from cyber security, they share similar traits in complexity, cognitive load, and reliance on numerous mental models for concept simplification. In a meta-analysis of CTA techniques' efficacy, Tofel-Grehl and Feldon (2013) found that instruction based on CTA methods yielded more effective outcomes than non-CTA approaches with a large effect size ( $g = 0.871$ ). Within those results, instruction based on skills derived from the PARI method yielded better outcomes than other CTA methods, with a large effect size ( $g = 1.598$ ). The existing research body demonstrates that PARI-based skills analyses are an effective technique when used to elicit skills with instructional design in mind.

### *CDM*

I used the CDM knowledge elicitation methodology to identify decision-making skills. Researchers developed the CDM methodology as a framework to elicit decision-making skills from environments “characterized by high time pressure, high information content, and changing conditions” (Klein et al., 1989, p. 462). These characteristics describe security operations well.

In CDM analysis, an interviewer conducts a semi-structured interview with a domain expert and asks them to recall a particularly challenging incident they encountered and provide a retrospective verbal report of it. The interviewer makes multiple sweeps through the experience, digging deeper each time by asking the subject probing questions. These questions focus on decision cues, patterns, goals, errors, and other statements that may identify decision points and their context. After that, the

interviewer may ask a series of hypothetical “What if?” questions to explore additional opportunities for decisions (Klein et al., 1989; Seamster & Redding, 2017).

CDM was an appropriate research technique for this study because it provided a structured framework for retrospective verbal reporting of events likely to trigger decision-making skills. By identifying information that leads to, characterizes, or results from decisions, CDM provided the interview tools necessary to identify decisions framed through real-world analyst investigations.

Like PARI analysis, researchers have also used CDM extensively for task analysis in multiple fields. Researchers have used CDM to identify decision-making skills in military intelligence analysis (Hutchins et al., 2004), medicine (Gazarian et al., 2010; Harenčárová, 2015; Pauley et al., 2011), civil engineering (Cattermole et al., 2015), and other cognitively demanding fields. Tofel-Grehl and Feldon’s (2013) meta-analysis of CTA-derived training found that instruction based on skills derived from CDM yielded more positive outcomes than those that did not, with a small effect size ( $g = 0.329$ ). CDM-based skills analysis has been proven as a reliable way to elicit decision-making skills from expert practitioners, making it suitable for this study focused on digital forensic analysts.

### *Results Integration*

While the procedural and decision-making skills provide significant value to instructional designers on their own, these skills are related in the performance of computer forensic tasks. Therefore, after I completed the Simplified PARI and CDM analysis, I conducted an additional qualitative analysis of these data sets together to

identify potential relationships between decisions and procedures that may provide additional value to cyber security educators.

### *Site Selection and Participant Sampling*

The methods of cognitive task analysis selected for this study require lengthy interviews with expert digital forensic analysts. The interviews' depth and required expertise warranted specific considerations about participant sampling and site selection, which I describe in this section.

#### *Site Selection*

Digital forensic analysis typically takes place in security operation centers (SOCs) within offices where analysts are employed (Sundaramurthy, 2014) or at the physical location of the affected organization when conducted in a consulting role. However, many analysts also work from home offices since this work is amenable to remote work where online tools facilitate data sharing and collaboration. Ultimately, the analyst experience defines the common link between the population rather than the physical location where they conduct the task.

With this in mind, I collected data for Simplified PARI and CDM analysis through interviews using web conference technology. Remote data collection offered several benefits. First, remote collection allowed a larger pool of potential applicants by permitting geographically diverse subject participation that might have been prevented by travel costs or site access restrictions. This strategy also enabled the selection of analysts working for multiple organizations, ensuring a more diverse representation of the analyst population. Second, it decreased human-to-human exposure during the COVID-19 pandemic, which may have limited participation from analysts with vulnerable pre-

existing health conditions. Third, it allowed me to break the interviews into multiple sessions, which was critical considering the potential fatigue resulting from lengthy combined PARI and CDM interviews. I did not anticipate that remote data collection would negatively affect the quality of the gathered data or the conclusions of the task analysis.

### *Analyst Sampling*

This study relies on purposeful sampling, which Creswell and Poth (2016) describe as a process that will “intentionally sample a group of people that can best inform the research about the research problem under examination” (p. 14). The nature of the digital forensic profession and the chosen CTA knowledge elicitation methods dictate the use of criterion-based purposeful sampling. This technique allowed me to obtain a sample of people who can best inform this research.

The Simplified PARI and CDM methods both require knowledge elicitation from domain experts regularly engaged with the skills I am assessing so that the derived skills accurately reflect mastery of task performance (Klein et al., 1989; Seamster & Redding, 2017). According to Ericsson and Charness (1994), experts have generally spent at least a decade in their fields. However, the computer security field is still in its infancy, limiting the number of practitioners with lengthy dedicated experience. Also, because security jobs sprang from systems and network administration roles, digital forensic practitioners may have gained significant and relevant evidentiary experience before transitioning to an investigation-centric career. Therefore, I selected participants with at least eight years of dedicated computer security experience or individuals with at least five years of dedicated investigation experience combined with three years of system or network

administration experience. Because analyst job roles are often split among multiple tasks, I also selected individuals who currently spent at least 50% of their time actively conducting security investigations at the time of data collection.

Digital forensic analysts have six differing investigative roles, as described in the literature review. I selected subjects from the three primary investigative roles for this study: triage analysts, forensic examiners, and incident responders. This selection provided specialty variation in the sample so that I could identify procedural and decision-making skills universal to the primary investigative roles. I did not select any malware reverse engineers, threat hunters, or intelligence analysts due to the primarily supportive characteristic of those roles' workflow.

I invited applicants to participate in the study through postings on social media, semi-private industry chat rooms, and mailing lists. As an incentive, I offered selected analysts a free seat in one of my online training courses of their choosing once they completed their participation. I asked prospective subjects to complete a short questionnaire describing their experience in terms of current role, years in practice, and their self-reported investigative ability along four dimensions: investigative heuristics, analysis tools, confidence, and evidentiary knowledge (see Appendix A). From this initial applicant pool, I excluded analysts from the study who did not self-identify as experts, did not meet the minimum time-experience requirements, did not spend at least 50% of their time performing investigations, or represented some other conflict of interest. I randomly selected nine initial participants from the remaining applicant pool, choosing three from each investigation role. I reviewed the selected participants' self-assessment scores to ensure alignment with field expertise; namely, that they felt confident in the

majority of the listed abilities. Since the CDM portion of the interview is contingent on analysts discussing a challenging incident they investigated, I sent follow-up emails to selected participants to ensure they could recall one such experience for discussion.

For cognitive task analysis to be complete, research using these techniques should seek to reach a saturation point. Charmaz (2006) described saturation as the point at which gathering additional data no longer reveals new insights. The breadth and complexity of investigation tasks makes complete saturation challenging. However, I was able to achieve saturation within the initial sample for the findings presented in Chapter Four.

#### *Data Collection Procedures*

Simplified PARI and CDM data collection rely on qualitative interviews, but each requires a unique method for conducting those interviews to elicit their targeted skill types. In this section, I describe the data collection processes used in this study.

##### *Simplified PARI for Procedural Skills*

A procedural skill is a series of steps performed to achieve a goal (Konoske & Ellis, 1986). Examples for analysts might include deciphering an obfuscated text string or retrieving sensor data. The fundamental nature and repeatability of consistent procedures within a profession makes procedural skills ideal for early task analysis efforts, particularly when the task analysis is geared towards education and job proficiency. Researchers can easily represent step-by-step procedures on paper once elicited (Seamster & Redding, 2017). They are also transferrable in an educational setting through various instructional designs and teaching methods (Hall et al., 1995).

In some cases, analysts repeat procedures so many times by analysts that they become trivial and require little cognitive effort. These may warrant classification as automated skills, another CTA skill type that serves as a “foundation of skilled performance by contributing to the efficient execution of multiple tasks” (Seamster & Redding, 2017, p. 137). Identifying and explicitly teaching procedural skills gives novices the foundation necessary to become adept at the procedures and work towards automating them in their journey to expertise.

I used the Simplified PARI qualitative research method to gather data that helped me identify procedural skills used by digital forensic analysts. To perform this collection, I generated a realistic investigation scenario with multiple paths (see Appendix B). This scenario inserts the subject into the middle of an attack timeline, giving them the option to move backward or forward in time relative to the initial event.

In a single interview session, I presented the initial evidence for the scenario to each subject along with a list of available evidence types. Using the PARI elements in Table 3.1, I asked strategic questions to collect information about procedures as the analysts worked through the problem. After the initial problem and evidence presentation, I asked subjects about their interpretation of the evidence. Then, I asked for their next action based on the available evidence and why they chose it (the precursor). Based on their choice, I provided the results of their action. At the point, the process repeated through the steps of interpretation, action, precursor, and result until the analyst completed the investigation.

Table 3.1

*PARI Elements and Associated Interview Questions.*

PARI Element	Questions
Precursor	Why are you taking this action?
Action	What would your first/next action be in solving this problem?
Result	None
Interpretation	What does the result tell you? Can you eliminate any causes from suspicion at this point? What causes do you suspect? Why do you suspect them?

*Note.* Adapted with permission from “Applied Cognitive Task Analysis in Aviation,” by Seamster and Redding, 2017, p. 145–146.

The analysts indicated when an investigation was complete. This occurred when the analyst assessed a disposition about the investigation (whether it represent malicious activity) and expressed satisfaction with the timeline they uncovered (the attacker actions that occurred). These steps are consistent with the Simplified PARI method (Seamster & Redding, 2017) framed through digital forensic investigations.

*CDM for Decision-Making Skills*

Decision-making skills are algorithms, heuristics, and rules of thumb used for deciding among choices (Seamster & Redding, 2017). Digital forensic analysts contend with a litany of decisions every day as they choose which investigative leads to follow, what evidence to collect, and whether a case warrants an immediate response. These decisions dictate the flow and eventual success or failure of an investigation. Poor decisions during the investigation can lead to mistaken conclusions, but they can also lead to significant decreases in time efficiency, allowing the attacker to further their

foothold while going undiscovered (Sanders, 2016b). Understanding the types of decisions analysts must make and their surrounding context provides a mechanism to ensure accuracy and efficiency in those crucial determinations.

I used the CDM qualitative research method to gather data that helped me identify digital forensic analysts' decision-making skills. The CDM interview consisted of four phases: incident identification, timeline verification, deepening, and optional "What if?" queries (Klein et al., 1989). For the incident identification phase, I instructed analysts to recall a challenging investigation they conducted prior to the interview. This incident formed the structural grounding for my examination in the interview.

At the start of the interview, I asked the analyst to describe the investigation at a high level to continue the initial incident identification. After they painted broad strokes, I proceeded to the timeline verification phase by asking them to walk through the investigation again while I asked specific questions designed to identify the particular sequence of the investigation and potential areas where they made decisions. With a clear timeline established, I advanced to the deepening phase.

The purpose of the deepening phase is to induce the subject to reveal cognitive facets of their decision making through a series of probing questions. Crandall et al. (2006) described this as,

I know what happened, who did what, and I know a bit about their role in the event. But what did they know, when did they know it, how did they know, and what did they do with what they knew? That's what Sweep 3 is designed to figure out. (p. 77–78)

In this third pass through the investigation, I asked relevant probing questions based on the list of CDM interview questions in Table 3.2.

Choosing which question to ask was an exercise in carefully evaluating the analyst's description of the event in real-time and asking questions that helped identify decisions or identify important context surrounding them. The probing questions in Table 3.2 provided some structure and tools to elicit information about decision-making skills. However, those questions did not constrain me, and I asked other relevant probing questions as needed.

After completing the third phase, I moved to the "What if?" query phase where appropriate. In this part of the interview, I used information from the analyst's investigation to form questions that allowed the analyst to speculate regarding different incident variables. I included questions surrounding expert-novice contrasts, hypotheticals, and investigative aids. Questions included, but were not limited to:

1. Expert-novice contrasts: "What would have happened differently if a novice were in this scenario?"
2. Hypotheticals: "If [evidence] were different, what would you have done differently?"
3. Aids: "What evidence or tools might have helped at this point?"

These questions provided more opportunities to identify and examine decisions that the analyst-provided scenarios did not allow on their own. I asked these questions when analyst scenarios lacked the necessary complexity to gather enough information about decisions or context surrounding them (Crandall et al., 2006).

#### *General Data Collection Prerequisites and Privacy Considerations*

I required all analysts to sign research release forms before enrollment in this study. Prior to the interviews, I reminded analysts that the information I collected would be kept confidential and nonattributable to them or their employer. Since the CDM

interview involves describing actual investigations, I instructed analysts to refrain from using company or victim names and to anonymize identifying details about the technical aspects of the breach to the best of their ability. The data collection procedures described in this document were pursuant to and approved by Baylor University’s Institutional Review Board.

Table 3.2

*CDM Deepening Phase Probe Questions*

Category	Questions
Cues	What were you seeing, hearing, smelling, and noticing, etc.?
Information	What information did you use in making this decision or judgment? How and where did you get this information, and from whom? What did you do with the information?
Analogs	Were you reminded of any previous experience? What about that previous experience seemed relevant for this case?
Standard operating procedures	Does this case fit a standard or typical scenario? Is it a type of event you were trained to deal with?
Goals and priorities	What were your specific goals and objectives at the time? What was most important to accomplish at this point in this incident?
Options	What other courses of action were considered or were available to you? How was this option chosen or others rejected? Was there a rule that you were following in choosing this option?
Experience	What specific training or experience was necessary or helpful in making this decision?
Assessment	Suppose you were asked to describe the situation to someone else at this point. How would you summarize the situation?
Decision making	What let you know that this was the right thing to do at this point in the incident? How much time pressure was involved in making this decision?
Guidance	Did you seek any guidance at this point in the incident? How did you know to trust the guidance you got?

*Note.* Reprinted with permission from “Working Minds: A Practitioner’s Guide to Cognitive Task Analysis,” by Crandall et al., 2006, p. 78–79.

I recorded interviews using web conferencing software and sent audio-only transcripts to a third-party service for transcription to text. I cataloged interview and transcript data using best practices for securing data-at-rest and labeled it with unique numbers, rather than identifying information.

### *Data Analysis Procedures*

Although Simplified PARI and CDM data both take the form of interview transcripts, each requires its own analysis technique. In this section, I describe those analysis techniques and how I derived skills from each dataset.

#### *Preliminary Analysis Steps*

After data collection for both Simplified PARI and CDM interviews, I submitted the anonymized audio recordings to a third-party transcription service to generate text transcripts. I reviewed the transcripts manually to identify industry jargon, acronyms, and other misinterpretations anticipated due to leveraging transcriptionists without domain experience. The final transcripts represent what analysts said, verbatim. I loaded the transcripts into qualitative data analysis software (ATLAS.ti) for further analysis. In some cases, I annotated these transcripts with time-stamped notes taken manually during the interviews.

#### *Simplified PARI*

The Simplified PARI method has the benefit of including preliminary data analysis steps during the data collection process (Seamster & Redding, 2017). Using the previously defined semi-structured interview questions framed through an incident, I had already begun the process of decomposing analyst investigations into precursors, actions,

and interpretations that define procedural skills. The actions represent procedure steps, while the precursor and interpretation provide the cognitive context surround the procedure: where it might occur, what it might lead to, what thought processes are involved, and what other skills might relate to it (Hall et al., 1995).

From here, I continued decomposition analysis with careful review and refinement of the collected interview transcript data. This process included the following steps:

1. I reviewed case transcripts and coded analyst responses based on the four PARI steps: precursor, action, result, and interpretation.
2. I grouped and coded responses based on similar actions. These formed the top-level identity of individual procedural skills.
3. For each procedural skill, I collected the precursors and interpretations to derive the context surrounding the procedure. I applied sub-codes (second level), linking the precursors and interpretations to their top-level procedural skill.
4. I deduplicated and combined similar actions, precursors, and interpretations as necessary while analyzing additional cases.
5. Finally, I built descriptive models for each procedural skill using the final set of precursors, actions, and interpretations.

I performed data analysis for each interview shortly after conducting it and before proceeding with the next. This strategy allowed me to use insights from data analysis to refine my interview techniques. Interleaving collection and analysis also helped me develop interview questions that elicited richer descriptions of certain procedures and clarified ambiguity and overlap between analysts.

### *CDM*

Similar to Simplified PARI, the CDM semi-structured interview defines the initial context for data analysis. By making several verbal sweeps through a prior investigation

with analysts, strategic questions yielded an initial set of decision points for which additional questions reveal their surrounding context. Therefore, the post-interview analysis process focuses on refining that data by comparing and contrasting similar decisions and their context within and across cases (Klein et al., 1989; Seamster & Redding, 2017). This analysis yielded a "model of the decisions in a way that describes the flow of the cognitive events" (Seamster & Redding, 2017, p. 187).

The CDM analysis involved multiple steps for interpreting decisions and their surrounding context from analysts:

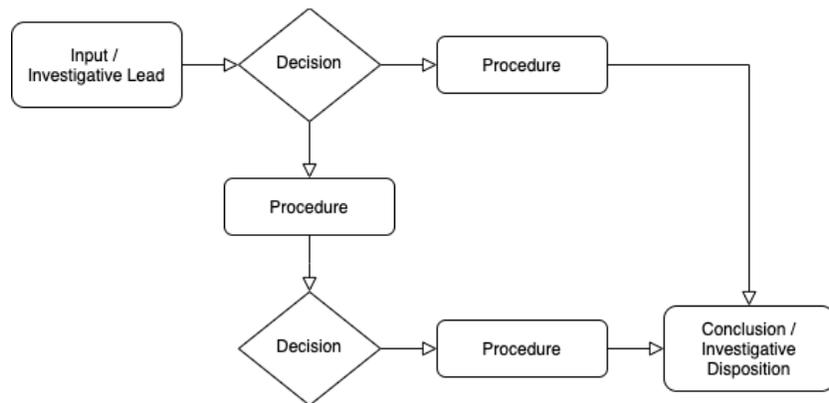
1. I reviewed case transcripts and coded analyst responses based on the decisions identified during the interviews. I grouped similar decisions to form the top-level identity of individual decision-making skills.
2. I identified defining features of the decisions based on a review of the analyst's responses to probing questions. These features included critical cues helpful for identifying the situation, goal alternatives for characterizing the decision, and potential errors that might arise from making the decision. I applied sub-codes (second level), linking the decision features to their top-level decision-making skill.
3. I reviewed decision-making skills and their features within and across cases to deduplicate and combine elements as necessary and normalize descriptive vocabulary.
4. With a final set of decisions and features, I built a descriptive model representing each decision-making skill. I also constructed a critical cue inventory, which contained a collection of informational and perceptual clues identified during the analysis (Klein et al., 1989).

Like the Simplified PARI analysis, I performed data analysis for each interview shortly after conducting it and before proceeding with the next. This strategy allowed me to refine my interview techniques based on insights gained from data analysis. For example, the analysis of a decision for one subject left unanswered questions about the decision's goal. With that in mind, I was able to ask more specific probing questions

about decision goals when the same decision point came up in a later interview. Similar analytic insights allowed me to pose useful expert-novice contrasts and hypothetical situations to garner useful facets of decisions, as described by Crandall et al., 2006.

### *Relating Procedural and Decision-Making Skills*

While procedural and decision-making skills warrant examination on their own, their relationship in the context of domain achievement reveals an additional opportunity for scholarly attention. Practitioners follow procedures that lead to decisions, and those decisions yield more procedures, as shown in Figure 3.2. The collective depiction of these skills represents more than the total of their parts by providing context regarding when and how analysts might encounter the need to leverage each one.



*Figure 3.2.* The relationship between procedures and decisions provides a framework for analyst-driven investigations.

As part of my analysis, I recognized links between procedures and decisions. Where encountered, I recorded these observations and reviewed within and across cases to highlight additional evidence supporting potential links between the manifestation of these skills. By applying added focus on relationships between skill types, my results provide additional opportunity for educators to incorporate these relationships into

analyst curriculum. These relationships help analysts better understand when they might encounter the need to employ specific skills and provide a better understanding of the consequences of their performance and decisions.

### *Reliability and Validity*

As a final step in analyzing both skill types, I shared the completed set of procedural and decision-making skills with analyst subjects for their feedback. Lincoln and Guba (1985, p. 314) identified member checking as “the most critical technique for establishing credibility.” This process provided additional validity to the study by ensuring that I captured the essence of analyst's experience performing investigations (Creswell & Poth, 2016; Lincoln & Guba, 1985; Merriam & Tisdale, 2015).

By showing analysts the collective results of the analysis, this step also provided them with the opportunity to comment on and confirm the use of skills that they had omitted in their interviews, further validating them and providing some sense of reliable occurrence. The use of CTA techniques focused on expert-interviews was a significant advantage since those experts could directly comment on the validity and reliability of the collective analysis findings. Combined with the collection of rich, thick descriptions of analyst experience from several expert subjects, this additional layer of confirmation for the CTA results speaks to the study's enhanced trustworthiness.

### *Ethical Considerations*

Cyber security investigations are typically shrouded in secrecy because of the potential negative impact of the public disclosure of breaches or safety risks for analysts associated with them. Because the CDM interview involved describing real

investigations, protecting participant identities was of the utmost importance for this study.

For maintaining anonymity, I assigned each participant a unique number to identify their associated data during collection and analysis. I stripped any personally identifying information from transcripts, notes, and other documents and relied exclusively on identification numbers. I stored research data on encrypted disks and followed industry best practices for the protection of data-at-rest.

I was the only person with access to the research data, except for the third-party service provider contracted to perform interview transcription. I did not provide participant names or other identifying information to the contractor. Data collection for this study was pursuant to Baylor's Institutional Review Board (IRB). All methods and procedures met ethical and legal standards defined by the university and relevant governing laws.

#### *Limitations and Delimitations*

The Simplified PARI interview relies on verbal reporting of thought processes while working through an investigation. Concurrent verbal reporting can have limited effectiveness when paired with tasks that strenuously tax working memory limitations (Seamster & Redding, 2017). To decrease the investigation tasks' cognitive load, I encouraged the analysts to take notes and take all the time they needed. I also provided an on-screen listing of evidence sources and simplified complex numbers required by the investigation (IP addresses, file hashes, etc.). While these measures helped, the concurrent verbal reporting method may still have limitations on the subject's ability to describe their procedures.

The CDM interview relies on analysts reporting actual investigations that they worked through (Klein et al., 1989). In some cases, analysts conducted these investigations several months before the interview. Therefore, these analysts may not have a perfect memory of the events, resulting in the omission of decisions and other key facts. I partially mitigated this limitation by priming the analyst to think through the investigation and review their notes before the interview. The CDM interview technique itself, which makes several sweeps through the events, may also inherently aid event recall (Crandall et al., 2006).

Both Simplified PARI and CDM analysis require expert subjects. No broadly accepted standard exists for certifying security investigation expertise. While I undertook significant efforts to identify experts during participant selection, these methods are not based on psychometrically tested and validated quantitative measures.

### *Conclusion*

This cognitive task analysis sought to identify digital forensic analysts' procedural and decision-making skills used during the investigation process. By performing qualitative interviews based on the Simplified PARI and CDM methods, I elicited cognitive skills that would not have been accessible through simple observation or verbal reports. My analysis results have a direct and immediate impact on the education of security practitioners and serve as cornerstones for the development of analyst curriculum. The following chapter describes the cognitive task analysis results and examines the implications of how educators might leverage the identified skills for analyst education.

## CHAPTER FOUR

### Results and Implications

#### *Introduction*

This Problem of Practice identifies the procedural and decision-making skills that experienced digital forensic analysts use during investigations. By using the Simplified PARI and CDM cognitive task analysis techniques, I elicited an accounting of these skills in a way that meaningfully informs forensic practitioners and educators within the domain. These analyses yielded a significant understanding of analyst procedural and decision-making skills that converge around a model of diagnostic inquiry.

This chapter presents the results and implications of the cognitive task analysis of digital forensic analysts in five parts. First, I briefly describe the experience-related demographics of the research participants. Second, I present the results of the Simplified PARI analysis to highlight analyst procedural skills. Third, I discuss analyst decision-making skills derived from the analysis of CDM data. Fourth, I provide a diagnostic inquiry model of analyst investigations based on integrating procedural and decision-making findings. Finally, I describe the implications of these findings for practitioners and educators.

#### *Analyst Demographics*

I collected demographic information for participants to ensure that they met the sampling criteria outlined in Chapter Three. The nine participants consisted of three individuals each from the triage, incident response, and forensic examiner analyst roles.

Of these analysts, the average age was 38.1 years ( $SD = 6.5$ ), the average number of years of IT experience was 10.8 ( $SD = 4.2$ ), and the average number of years of information security experience was 13.2 ( $SD = 5.0$ ). The analysts reported spending an average of 86.2% ( $SD = 13.2$ ) of the time in their current role performing investigations.

While all analysts met the minimum guidelines for participation in this study, there was a notable amount of variance in their years of experience in information security. This variance likely is a product of individual analyst expertise relying more on the quality of experience rather than its duration, consistent with broader research on expert performance (Ericsson, 2004; Ericsson, 2008; Ericsson et al., 2018). Several things may impact the quality of experience for analysts, including the available technology and evidence sources in an individual role, access to the expertise and mentorship of others in that role, access to formal training, or the nature of the job itself. For example, analysts in consulting positions are more likely to receive exposure to more diverse attack investigations than analysts focused on a single company that experiences a limited number of attacks and consistently uses the same technology components. At the same time, there also is no standardized path into cyber security like with other, more well-established professions such as law, medicine, or financial services. Therefore, practitioners following different career paths may take differing lengths of time to achieve expertise within the digital forensic domain.

### *Procedural Skills*

A procedural skill is a series of steps performed to accomplish a goal (Konoske & Ellis, 1986). While analysts may have a unified purpose of figuring out what malicious actions occurred on a system or network, their skills to achieve this goal are numerous

and complicated. By presenting analysts with an investigation scenario and carefully probing their actions, I accounted for many of these skills.

The Simplified PARI cognitive task analysis technique provided both the data collection and analysis framework for this portion of the study and the construct for describing cognitive skills: precursors, actions, results, and interpretations. The basic unit of analysis for procedural skills centered around investigative actions, where an analyst performs some operation to understand events on the computer network further. A precursor justifies that action. The action typically produces a result that the analyst interprets before taking another action.

This study identified 308 investigative actions ( $M = 34.2$ ,  $SD = 10.3$ ), 210 unique precursors ( $M = 23.3$ ,  $SD = 8.5$ ), and 95 discrete interpretations ( $M = 10.6$ ,  $SD = 4.6$ ) across nine subjects. Although every analyst worked through the same investigation exercise, there were multiple paths through the investigation by design. Some paths yielded more meaningful results earlier than others, consistent with Sanders's (2016b) research on the effects of opening move speed in investigations. For example, analysts who chose to focus on execution analysis first arrived at more definitive conclusions about what happened faster than those who focused on network traffic analysis. This difference likely explains much of the variance between the number of investigation actions, precursors, and interpretations recorded per analyst. The remainder of the variance may result from differences in analyst level of expertise, experience with the available data sources, or other underlying cognitive differences.

Using this data, I identified procedural skills inherent to the investigation process by performing a thematic analysis of the actions, precursors, and interpretations exhibited

by analysts within and across cases. In the remainder of this section, I characterize and describe those skills across five categories: inquiry skills, evidentiary skills, anomaly detection, network mapping and attack visualization, and metacognitive skills.

### *Inquiry Skills*

Analysts began investigations based on an input that represented an anomalous or malicious relationship. However, that initial evidence only described a single event on a timeline. The analysts aimed to uncover the remaining parts of that timeline to determine if someone attacked the network. During investigations, the most prevalent cognitive skill expert analysts displayed was forming investigative questions to resolve uncertainties surrounding the events, disposition, and nature of a potential attack.

An investigative question is an inquiry that, when answered, helped an analyst identify and understand what events have transpired on the network they were protecting. Analysts formed investigative questions in their minds based on existing evidence, prior experience, researched threat capabilities, theories, hunches, or input from other analysts. Analysts answered investigative questions by collecting, manipulating, and interpreting evidence, another skill domain I discuss later in this chapter.

In most occurrences, analysts expressed their inquiries in the explicit form of an interrogative question, indicated by their sentence structure and voice intonation. In other instances, they expressed their questions as statements. However, these statements still contained prominent elements of inquisition with the expectation of an answer. One analyst said, “What was the name of the process?” while another one said, “I want to know the name of the process.” Both statements were functionally equivalent forms of

inquiry. I refer to both interrogative questions and question statements simply as questions or investigative questions throughout the remainder of this study.

The formation of an investigative question served as the fundamental action analysts took in their pursuit of understanding what events transpired. Analysts asked investigative questions at the onset of the investigation based on an initial input and continually asked additional questions throughout the investigation as they interpreted new information collected based on answers to other questions. Further questions arose directly from answers to prior questions in a linear fashion or the holistic interpretation of answers from multiple questions. In either case, a series of questions and answers represented lines of inquiry and characterized the analyst's investigation path.

Examination of expert analysts' investigative questions showed that not only was the formation of questions the core unit of action for the analyst but that expert analysts excelled at asking high-quality questions. A high-quality investigative question was one whose answer had strong potential to move the investigation forward in a meaningful way by identifying or characterizing important relationships. The vast majority of these questions had three characteristics: high relevancy, high specificity, and clear answerability. The following questions from analysts exhibited these characteristics:

- “What is the parent process for KZM.exe on all of these systems?”
- “I looked at the proxy logs to see if I can see any calls to that particular Google Drive path.”
- “That fgrimes user, since we know that 9.9.9.9 was used to communicate with all three of those systems, was that same user account used on all three of those systems?”

Analysts asked relevant questions when they based them on the interpretation of existing evidence within the current investigation or other investigations involving

similar components and techniques. Analysts interpreted evidence and used it to forecast theories about events that may have happened. The questions they pursued reflected a test for artifacts that validated or invalidated their theories. In the above examples, an analyst based the first question on the discovery of a malicious process named KZM.exe, the second based theirs on the presence of a Google Drive link in a phishing email, and the third based their question on the discovery of the FGRIMES user account logging in to a machine they previously observed interacting with an attacker command and control channel. Questions that lacked relevance seemed like shots in the dark, with no apparent ties to the existing investigation or others like it.

A specific question had a narrow range of answers or answer types, such as any possible value that fits within the expected format of a file name, username, event log ID, IP address, and so on. Analysts asked questions whose answers took a specific form knowing those answers may fill gaps in the attack timeline or the context surrounding identified relationships. Said more simply, specific questions ensured analysts only had to look through four-sided pegs when they sought to fill a square hole. In the above examples, the first question can only yield a process name. The second and third questions can only yield yes or no answers based on the presence of specific log results. The limited range of potential answers or answer types makes these questions specific. Questions lacking specificity did not point to any evidence source or yield meaningful information when answered.

Analysts could realistically resolve answerable questions through the examination of evidence. In the above examples, all three questions are answerable with commonly used forms of digital evidence: Windows event logs for the first and third questions and

HTTP proxy logs for the second. Although the PARI scenario was consistent across analysts, the answerability of a question would be environment-specific, as organizations do not all collect the same forms of evidence.

While analysts in this study consistently asked high-quality investigative questions, there were some instances where that was not the case. In events where an analyst formed a question that lacked relevance, specificity, or answerability, they often arrived at unsatisfying answers that did not move the investigation meaningfully forward. In some cases, the analyst abandoned that line of questioning. In others, they used this opportunity to refine their questions and increase their quality. For example, one analyst asked, “Do I see any kind of spam that was sent to JSMITH?” While relevant based on threat intelligence the analyst had about phishing messages delivering the malware, this question is not answerable in its current form based on the available evidence. The analyst eventually formed questions leveraging techniques that looked for messages from senders the recipient had never seen before combined with searches for mixed case and non-business oriented subject lines. These questions were answerable using the available mail transaction logs. Scenarios like these demonstrate that question refinement is also a helpful skill as analysts proceed down a line of inquiry and discover limitations of their evidence sources.

It is notable that just because a question did not yield an answer or confirm an analyst’s theory does not mean their question was not well-formed or valuable. In some cases, quite the contrary. Lack of data for some questions proved to invalidate theories and terminate investigative paths quickly. These fast failures allowed the analyst to draw quick conclusions and move to other lines of inquiry in some cases.

Asking high-quality investigative questions was, to some degree, self-rewarding. These questions were more likely to provide positive feedback by resolving uncertainty or identifying new and vital facets of an attack timeline. Therefore, it may be that analysts who encountered the benefit of specific high-quality questions in previous experiences were likely to rely on them and pursue similar lines of inquiry more frequently going forward.

The entirety of unique questions that analysts may ask across all possible investigations is not realistically identifiable due to the diversity of technology and evidence sources that analysts use. However, my examination of investigation questions in this study yielded three core categories of questions across the domain. I classified investigative questions as event-relative questions, capability matching questions, or utility questions. I also identified clusters of questions that comprise directed analysis techniques.

*Event-relative questions.* Analysts primarily formed investigative questions relative to existing timeline events. When asking these questions, analysts oriented themselves at a specific event in the attack timeline and asked questions that looked backward in time, forward in time, or directly at characteristics of the event itself. These orientations represent the three categories of event-relative questions: preceding, succeeding, and context questions.

Analysts asked preceding questions when they oriented from a specific event in the attack timeline and asked questions about events leading up to it. For example, one analyst discovered the execution of a malicious file named ReportReview84.exe. Orienting themselves from that timeline event, they began asking preceding questions to

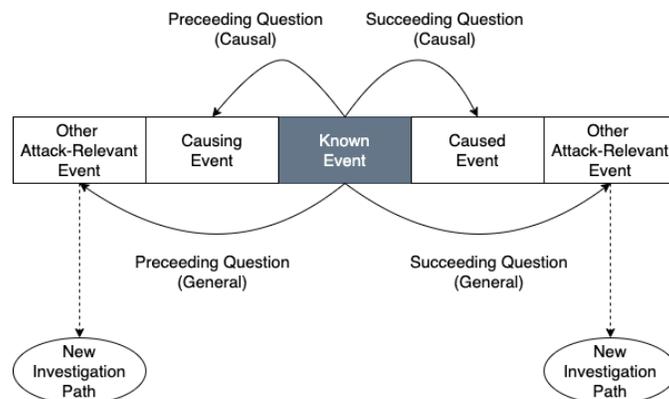
work their way backward. One of those questions was, “How did the user get the link to be able to download ReportReview84[.exe]? I have proxy data. Can I check and see what URLs the user visited?” This question led the analyst to a Google Drive connection that occurred from the infected system just before the malware executed. Knowing that attackers used phishing messages to deliver this specific malware, they then asked another preceding question, “Can I do an e-discovery search and see what emails that might’ve included a Google Drive link that were sent to JSMITH in the past week?” This line of inquiry led them to discover several other events occurring before the program execution where they initially oriented themselves, including the phishing message that began the attack.

Analysts asked succeeding questions when they oriented from a specific event in the attack timeline and asked questions about events following it. For example, one analyst discovered a compromised system where an attacker executed the PsExec tool. While PsExec is a legitimate tool administrators use to execute commands on other systems remotely, attackers can utilize this same functionality for malicious purposes. To determine if the tool successfully ran commands against remote hosts, the analyst asked a succeeding question to search for Windows event log entries that indicated successful remote PsExec service execution on target hosts.

We’re going to search for PsExecSvc on systems to see. It should be a 7045. It’s a process that starts, and what we’re looking for is to see how many systems have that process running. What we have is we have an infected system running PsExec a few times, so there’s that and then there’s going to be a corresponding event log entry. There should be a system event log entry for that service process running on the additional systems that that Sales2 is talking to.

This line of inquiry led the analyst to discover which target hosts the attacker ran commands on remotely.

Analysts typically asked preceding or succeeding questions to determine the cause or effect of a known event. However, some event-relevant questions sought to find relevant events beyond a single causal link. For example, analysts who discovered malicious process executions commonly asked a preceding question to determine what parent process launched the known malicious process and a succeeding question to determine what child processes the malicious process launched. These questions both sought to establish causal relationships that identify relevant attack timeline events and further current lines of inquiry. However, in multiple instances like this one, analysts also asked additional preceding and succeeding questions to look for other unique or interesting process executions around (before or after) this time frame. These analysts knew it was possible that an attacker was in partial control of a system at this point and may have executed more than one process to achieve multiple goals. By asking additional preceding and succeeding questions related to process execution, they could find evidence of other malicious activity independent of a causal link to a known event. These questions yielded additional lines of inquiry for further exploration. I provide a model of this process of causal and generalized preceding/succeeding questions in Figure 4.1.



*Figure 4.1.* Analysts ask preceding and succeeding questions to establish causal relationships and identify other generally suspicious activity relative to another event.

Analysts often combined two separate preceding and succeeding questions into a single proximate question when they knew that relevant events might exist on either side of the timeline. In one example, while referencing another event, one analyst asked, “Was anything downloaded around that time?” This single question achieved the same net result as splitting it into its preceding and succeeding components and asking them separately. Analysts often formed these combined questions with correlation in mind, as depicted conceptually in Figure 4.2. In these instances, they had already identified a suspicious event in one evidence source and wanted to know if any relevant events occurred in another evidence source surrounding the timestamp of the known suspicious event. For example, multiple analysts searched for suspicious execution logs that correlated with the timestamp of a malicious connection. These proximate questions were common, and analysts generally used them along with varying time intervals (within milliseconds, seconds, minutes, hours, days) to set a bounding on their query.

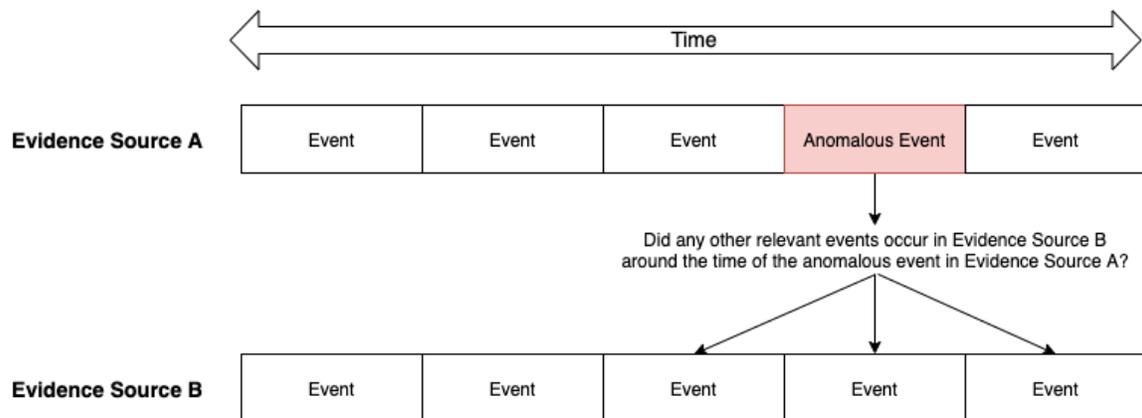
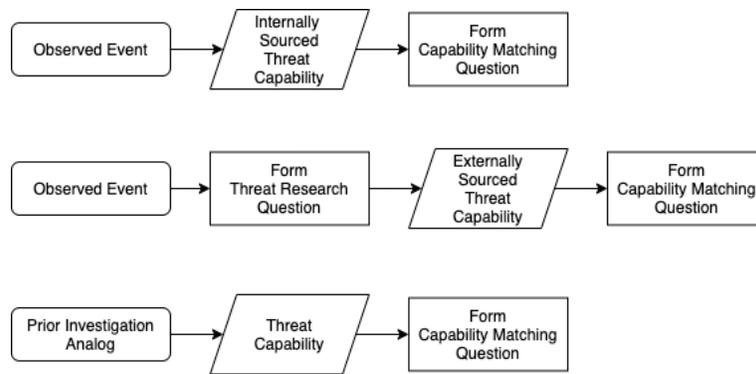


Figure 4.2. Analysts often ask proximate questions when seeking events that correlate with the timing of already known suspicious events.

Analysts asked context questions when they oriented at a specific event and asked questions about the characteristics of the entities and relationships associated with it.

These questions sought to reveal details that helped the analysts to understand the event's relevance to a potential attack or characterize its disposition as malicious or benign. For example, after discovering a suspicious process was running on a host, multiple analysts asked some version of, "Do I have a file path?" Analysts reported seeking this information for several reasons, such as to match known malware activity observed in prior cases or threat intelligence reporting, to associate the file with a specific user, to indicate the source of the malware, or to provide some other helpful context about the disposition of the file. In another example, while examining the same suspicious process, analysts asked several context questions about the role of the user account that ran the process, including questions such as "is this a normal privileged user?" or "Is this an elevated user, like a systems administrator or a domain administrator?" or "Was this running as [the] system [user]?" These questions helped analysts establish the potential impact of the attack, the threat surface available to the attacker, and the extent of their capabilities. The answers also provide some clues as to the source of the infection based on the user's role.

*Capability matching questions.* Analysts asked capability matching questions when they sought to confirm or refute the presence of a known threat capability. They identified these capabilities from internal and external sources tied to the current investigation, and analogs from prior investigations (Figure 4.3).



*Figure 4.3.* Analysts based capability matching questions on characteristics drawn from internal sources, external sources, and analogs from prior investigations.

In its simplest incarnation, analysts observed an attacker performing some action on a system and formed a question to determine if they took the same action on other systems. For example, one analyst discovered a system downloading malware from a Google Drive URL. Based on this behavior, they asked, “What I might want to do, though, is look for that Google Drive URL anywhere else across the network.” At this point, they knew the attacker was capable of hosting and delivering malware using Google Drive, so they used that finding to form an investigative question that helped them determine the prevalence of this facet of the compromise. In this example, the analysts derived the capabilities they based the question on from internal events within the current investigation that they observed first-hand.

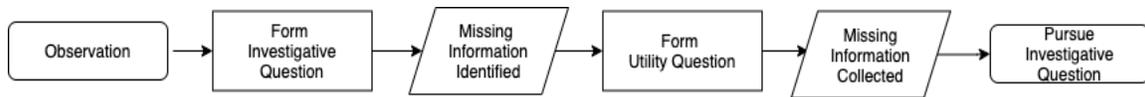
Analysts derived some threat capabilities from external sources, such as public threat intelligence databases. In a recurring example, many analysts suspected a strain of malware was present on a host based on an IDS alert. The alert provided a domain name associated with the malware, so the analysts began searching public evidence repositories for associations with this domain and malware strain. Analysts were asking versions of the question, “What have other analysts found out about this domain name and

malware?” Through this research, the analysts eventually found that other analysts investigating similar attacks on other networks had observed the malware communicating with the IP address 9.9.9.9 for command-and-control purposes. Understanding that capability, the analysts asked some version of the capability matching question, “Is their [network] traffic to that 9.9.9.9 address?” By asking this question, they attempted to match the known capability of the malware to network traffic in or out of their network. The answer to this question helped them to determine if the malware was present on the network and which hosts it infected. While they still found the initial observation that compelled them to research the threat capabilities within the current investigation, the analysts relied on external sources to perform the necessary research and find capabilities that drove capability matching questions. This additional step required the formation of intermediary threat research questions focused on investigating the history and capability of the threat itself, independent of the currently defended network.

In addition to drawing capabilities from internal and external sources related to the current investigation, analysts also formed capability matching questions through analogs from prior investigations. In one such case, an analyst asked, “Does the VPN log show anything for Sales2?” This question was not particularly relevant to their current line of questioning or a recent event. Instead, the analyst’s precursor for this event was a prior investigation they had been a part of years earlier. In that investigation, the analyst observed an attacker leveraging stolen credentials to maliciously access a host using the VPN appliances remote access capability. That attacker in that investigation had also used similar persistence techniques as what the analyst found in this one. In the current investigation, the analyst did not yet know how the attacker had gotten initial access to

the system named Sales2. Knowing that VPN logs were available, they formed this capability matching question that sought to fill the gap in the attack timeline based on previous experience.

*Utility questions.* Analysts asked utility questions when they sought to collect data that was necessary to answer another question. Answers to the utility question did not provide immediate value toward an understanding of the attack timeline on its own. Instead, analysts described these questions as those whose answers would fill in the blanks of other future (and usually immediate) investigative questions. In these cases, analysts appeared to form their target investigative question first before identifying the piece of missing information needed to finalize it. At that point, they asked utility questions to identify the missing information (illustrated in Figure 4.4).



*Figure 4.4.* Analysts formed utility questions to collect information they needed to answer other investigative questions.

As an example, while pursuing a series of preceding questions focused on malicious process execution, an analyst asked, “Which users are on this computer?” In this case, the analyst did not ask any context questions about the users themselves, and the answer provided no additional information about the attack timeline. Instead, the analyst then asked, “I’d search any mail coming from not his company. So, no internal email and I’d start with looking for any outliers, anybody who’s never mailed him before.” In this case, “his” referred to the user of the system the analyst uncovered with the previous question. Another analyst approached the same portion of the investigation

similarly, asking, “For the FIN7, Sales2, and Sales3 systems, do we have a list of known users that use those systems? I’m going there because it’s going to pivot over into the exchange logs.” Knowing the usernames associated with these systems was only relevant to the analysts because it allowed them to answer the questions they were planning to ask next, focused on identifying phishing attempts. The usernames served as a pivot point between the process execution evidence they were examining and the mail transaction logs they hoped to examine next.

*Directed analysis techniques.* While most investigative questions analysts asked stood alone as independent queries, sometimes repeated patterns of questions emerged. These questions often spawned from similar inputs and always targeted specific goals. Analysts also repeated them frequently within and across cases. By labeling these question clusters, I identified a series of directed analysis techniques (DATs). A DAT is a grouping of related questions whose line of inquiry works toward a specific analysis goal.

An example of a commonly observed DAT used by analysts was lateral movement analysis. Once analysts recognized that an attacker compromised a system, they asked questions like these:

- “Are there any other systems that Acronis has logged into from sales2?”
- “Are there any log ins to those systems that aren’t their typical user?”
- “What I’m looking for is that PSEXEC 4624 or 7045 where it’s actually starting the service and connecting to a remote machine.”
- “I’d check for log ins because is jsmith’s user logged in to Sales3 or Fin7? That’d be odd because those aren’t his computers, right?”

While these questions are all different, analysts described similar precursors for all of them: to help determine if the attacker leveraged their access on one host to move

to other systems within the network and to identify those systems and the mechanism(s) used for lateral movement. The examination of these questions allowed for the development of archetype questions commonly associated with each DAT that are abstracted from specific evidence types or characteristics of an individual investigation.

Example archetype questions for the lateral movement DAT included:

- Was there any connection between [Compromised Host] and other internal hosts?
- What internal systems did [Compromised Host] communicate with after [Process] executed?
- What systems did [Compromised User Account] authenticate to?
- Are there any authentication attempts to [Suspected Target Systems] that were not their typical user?

A DAT typically contained some event-relative, capability matching, or utility questions. However, analysts often started down DAT lines of inquiry independently of the three core question types. For example, after establishing that an attacker had a foothold on the system, some analysts began asking questions associated with the data exfiltration DAT. The analysts had not discovered anything within the current investigation to suggest that exfiltration occurred. However, analysts described knowing that data exfiltration was a common goal of many attackers and that its consequences were potentially dire. Therefore, they chose to follow this line of inquiry to attempt to determine if the attacker exfiltrated data.

Analysts frequently chose to pursue DATs after exhausting investigation paths spawned from event-relative or capability matching questions. This finding suggests that these techniques may occasionally serve as fallback investigation paths drawn from analogs to prior investigation experience or investigative playbooks. The use of various

DATs explained investigative questions analysts asked that were not classified as independent event-relative, capability matching, or utility questions.

I identified strong evidence for nine DATs used by analysts throughout the investigation scenario I provided, listed in Appendix D. This list does not constitute a complete accounting of all possible techniques and is limited by the scope of the investigation scenario I created for this study. However, it may serve as a framework for clustering additional analyst inquiry through future research.

### *Evidentiary Skills*

Analysts' ability to conceptualize and interpret evidence was critical to their skill in asking and answering investigative questions that advanced their knowledge of digital events. By analyzing the function of evidence relative to analysts' actions, precursors, and interpretations, I characterized the evidentiary skills that analysts leveraged during investigations. These skills manifested across four dimensions: interpretation, capability comprehension, collection, and manipulation.

*Interpretation.* Analysts formed investigative questions with the intent to answer them using evidence. Interpreting evidence to answer questions required significant declarative knowledge about individual evidence sources, which included the availability of the source, the fields contained within it, each field's data structure, and typical stimuli and responses manifested within the fields. Additionally, analysts demonstrated an understanding of how individual fields and combinations of fields represented the formation, change, or dissolution of relationships.

The small slices of data analysts took from evidence did not provide significant meaning on their own, absent contextual interpretation. For evidence to give meaning to

the analyst and relevance to the investigation and attack timeline, analysts needed to interpret and understand what the evidence represented in multiple contexts. They considered the significance of evidence within the context of the question they asked, other events in this specific investigation, analogs from other investigations, and their perception of normal versus abnormal activity in the evidence source. Given shifts in any of those contexts, analysts were likely to arrive at different conclusions based on their interpretation of the same evidence. For example, multiple analysts discovered the execution of a file named ReportReview84.exe on a host. While most analysts correctly suspected this file was first-stage malware that executed the second stage of an infection chain, one analyst thought this application was more likely to be legitimate and a potential victim of software exploitation. Two context changes were responsible for this analyst's differing interpretation. First, they had observed a similarly named file in a prior investigation that was a legitimate business application. Second, they arrived at this finding from a different path than the other analysts. The other analysts had uncovered evidence that the host in question was the initial foothold the attacker leveraged to get into the network. This analyst had not established that fact yet, so they thought it was equally likely that the attacker had moved to this host laterally from another compromised host inside the network. While the analyst eventually arrived at the correct conclusion regarding the file's role in the attack, their varying interpretation of this evidence sent them along different investigation paths toward that end.

This example and others like it highlight how analysts interpreted the same evidence multiple times beyond its initial surfacing. As analysts discovered new evidence, they sought to analyze that data alongside other findings and not just as an

isolated answer to a particular question. Analysts consistently and frequently recalled their prior observations and conclusions as they discovered new evidence and relationships. They recognized that their interpretation of evidence could change as their understanding of the attack timeline evolved. This periodic recall and reassessment formed the basis for some level of error-checking and a mechanism to stimulate additional investigative questions.

*Capability comprehension.* Analysts' ability to form effective investigative questions depended on their understanding of individual evidence source attestation capabilities. By understanding the types of relationships and events that an evidence source could describe, analysts gained the ability to form investigative questions relevant to that source. This comprehension provided the knowledge for analysts to ask investigative questions that were answerable, which was one of the characteristics of good investigative questions.

While every analyst participating in this study is an expert, there were differences in each person's comfort level with individual evidence sources. I expected these differences, as the array of evidence sources available for analysts is large and diverse. Throughout this study, it became clear that the analysts who had a broader comprehension of more evidence sources could ask a more varied set of questions. When analysts understood the relationships and context available from the interpretation of an evidence source, they were more likely to leverage that source where appropriate. For example, every analyst that participated in this study performed some form of process execution analysis. However, they relied on different evidence sources to form their questions and seek answers. Some analysts developed their questions based on the

capabilities of Windows Prefetch evidence, others came up with questions targeting Windows Security event logs, and some sought system memory with their questions. While all three sources could provide a list of processes that the system or someone using it executed, each came with nuances and limitations that analysts accounted for in their questions.

In some cases, the investigation scenario did not provide analysts with their preferred evidentiary mechanism for forming and answering a question. This situation caused some analysts to reform their questions based on an evidence source with which they were less comfortable. When forced to leverage unfamiliar evidence sources, analysts showed some signs of confusion or discomfort. They did not always get the answer they sought or failed to recognize additional context the data source provided, even when it was particularly relevant to the investigation.

Analysts who were skilled with an evidence source sometimes demonstrated competence by assessing the configuration of the data source before forming investigative questions. For example, analysts who were well versed in Windows Sysmon logs knew that the Sysmon tool requires configuration and that these configurations vary from network to network and affect what data the tool records. One analyst hoped to leverage Sysmon evidence to learn more about a suspicious DNS query. Before sharing the investigative question, they first asked me to confirm that IT staff configured Sysmon to record DNS queries.

Analysts who were unskilled with an evidence source asked exploratory questions about the capabilities of the source. Other analysts indicated that they would likely spend time researching the source if the time constraints of our interview did not bind them. In a

few cases where data limitations forced the analysts to rely on an evidence source with which they were uncomfortable to form and answer a specific question, they abandoned their current line of inquiry and took a different investigative path altogether. These scenarios demonstrate how an analyst's incomplete understanding of the attestation capabilities of an evidence source may limit their ability to form specific investigative questions and follow worthwhile investigation paths. At the same time, they also demonstrate how analysts can be adaptable by seeking additional knowledge about data sources or approaching timeline gaps in another way.

Analysts' application of their comprehension of evidence capabilities was present across every type of investigative question. However, investigative questions did not always represent an exclusively one-to-one relationship with evidence sources, as shown by the presence of utility questions. When an analyst formed an investigative question, they considered the evidence source's capability to answer it and the data they needed to perform the query and achieve an answer. When they lacked some essential data required for the query, they asked utility questions to obtain that data. Other evidence sources typically answered these utility questions. Therefore, analysts sometimes required comprehension of the capabilities of multiple evidence sources to achieve a single investigative goal.

*Collection.* Typically, analysts use software tools to collect the data they want from evidence sources to answer their questions. Since this study relied on a verbal interview mechanism, analysts were not responsible for manual data collection. However, I expected them to request data in a manner that was compatible with standard tools used for retrieving whatever data type they sought. I probed this skill by asking questions such

as, “How would you answer this?” and “What would you search for to find the answer?” Responses to these questions provided demonstrations of analysts’ skills translating investigative questions into formats compatible with the tools used to access whatever evidence source they referenced. For example, while performing prevalence analysis related to suspicious file hashes, one analyst asked, “So since we have those hashes, I’m going to grab those hashes and potentially use whatever they have in their network, even if just PowerShell, to see if those files exist anywhere else across their network.” When asked how they would answer the question, they elaborated with this response:

I guess preferably I would use, if they have anything in their environment, any type of EDR in place. You could use PowerShell to push out a query to see if that file is anywhere else. We actually have an EDR solution . . . we could push that out and basically get that on the systems to look and see. That’s a little bit extreme with only one system showing signs of compromise. So, I’d almost rather use a local method to see if those files are there by hash.

Pointing to these tools required that the analyst understand which tools could accept a file hash and perform the search across the environment.

This example also indicated a preference for one tool over the other, highlighting the analyst’s skill in evaluating multiple collection options and choosing a preferred choice based on the context of the given situation. Analysts consistently asked questions related to the capabilities of their forensic tools for the collection of data from evidence sources before forming investigative questions, while refining questions, or when clarifying how they would seek answers to questions.

*Manipulation.* Although I provided evidence to analysts verbally, they still frequently indicated a preference for how they would manipulate the data for their eventual interpretation. Evidence exists in its most basic form as raw text, but analysts

stated a preference for that data in other formats in different scenarios. Those formats included:

- Tables: When viewing data tables, analysts could line up similar fields and look for anomalies across multiple entries. Analysts could also sort tabular data by field.
- Statistics: Analysts performed certain types of anomaly detection by using various statistics like averages or sums.
- Aggregations: Analysts grouped and counted unique values to find outliers.
- Graphs: Analysts used graphs to identify patterns that indicated specific types of activity.

These data formats aligned with the purpose of their investigative question. For example, while examining authentications between systems, one analyst asked, “when I’m looking at all of the logs for Acronis authentications, I would be looking at the least frequent logon types.” Manipulating the data into an aggregation like this allowed the analyst to spot outliers by examining the least frequently occurring logon types. The analysts used this strategy to identify potential lateral movement facilitated by the Acronis account.

### *Anomaly Detection*

Much of the analysts’ job involves disposing events as either benign or malicious. Not only did these actions require that the analyst figured out what happened, they also had to determine if the event represented legitimate network activity or the actions of an attacker. These decisions were crucial because once an analyst decided that an action was malicious, events and relationships associated with that event were typically deemed malicious by association. An error in disposing an event could lead the analyst to miss an entire sequence of malicious activity or waste time pursuing a series of activity that is unrelated to an attack.

Analysts primarily relied on pattern matching for anomaly detection. They routinely asked investigative questions with the intent to review evidence and compare it against patterns they observed elsewhere. These patterns took many forms and varied based on the evidence sources considered at the time. They leveraged known patterns of both benign and malicious activity to compare observations within the current investigation to model examples.

In some cases, analysts focused pattern matching on specific behaviors. For example, analysts performing exfiltration analysis asked questions like, “But do I see any unusual data spikes, any large data transfers off of SALES3 or FIN7 out to that CDK12.kazam domain?” and “What rough volume of traffic am I looking at here? Is it gigabytes, or is it megabytes?” They explained that this pattern of large outbound data spikes was typical in cases where data exfiltration occurred. By seeking evidence of this pattern, they attempted to determine if an attacker had exfiltrated data from the network.

In other cases, analysts examined data while attempting to match a broad set of patterns, rather than a single specific pattern. In one such example, an analyst expressed their intent to examine the list of running processes on a system to find evidence that a malicious process was executed. They expanded to the following investigative questions when pressed for how they would identify such a malicious process:

Is there anything that’s running out of the downloads folder? Or is there anything running out of System 32 that looks out of place? Or any other common Windows directories that would suggest that maybe someone had downloaded a piece of malware from their email and ran it really quickly? Do I have anything that’s running on the desktop? Do I have anything that’s running in common user directories that kind of smells like something being delivered via phishing?

Each one of these questions represented some attempt to match a pattern of malicious activity the analyst had observed or researched at some point in their career. By analyzing

analysts' actions, precursors, and interpretations during this scenario, I was able to identify several subset pattern detection methods that analysts used to arrive at dispositions regarding the benign or malicious nature of events.

*Baseline comparisons.* Analysts frequently sought to establish baselines of activity over extended periods for certain types of events. They used these baselines to develop a standard of normalcy. Then, they took a smaller subset of similar activity (perhaps even just a single event) and compared it to the baseline to determine if it was an outlier or something that occurred with some frequency. In one such example, an analyst discovered suspicious activity on a system tied to the FGRIMES user account. So, they asked, "Is that FGRIMES user, is that a normal user for the FIN7 workstation?" They established a baseline by reviewing the FIN7 authentication logs:

I would just look at, again, looking at the authentication logs just to see if it was maybe a console login on a daily basis, something to that extent. Something where I could paint a frequency pattern. Especially if like, okay, we'll say business hours are 8:00 to 5:00 just to make it easy. At 8:05 I see an authentication from that user Monday through Friday, that would establish that baseline for me.

By understanding the baseline authentication activity for the user, they could then compare that against the actions the user account was currently taking to determine if the account was under the control of the legitimate user or an attacker.

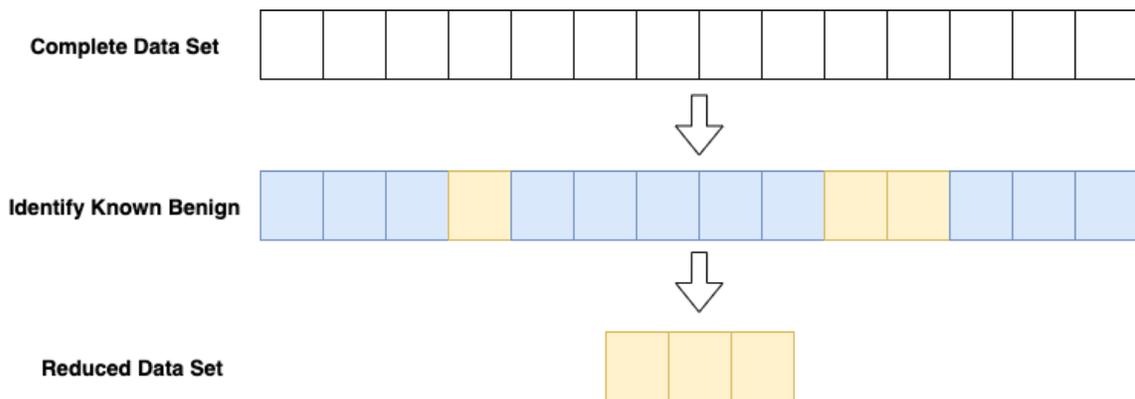
Baseline comparisons are a unique subset of pattern matching both because they involve establishing a known benign pattern over a more extended period and because the analyst derived the model pattern through evidence gathered within the scope of the active investigation and currently defended network. In these scenarios, analysts first encountered an anomalous event. To better understand the event, they explicitly sought

out additional data containing the same characteristics as the newly encountered event and built a baseline data pattern. Then, they compared the original event they discovered with the established baseline to aid in making a disposition.

*Frequency assessment.* In some cases, analysts sought to calculate the frequency of occurrence of specific events to provide clues about the disposition of the events. In one instance, an analyst encountered network communication from a unique HTTP user agent associated with a month's out-of-date version of the Firefox web browser. To determine if the application responsible for the traffic was malicious, they asked, "I'd search for that user agent to see if there's any more additional IP's that that user agent is talking to." The analysts counted the number of hosts communicating with that user agent. Since they only found a few hosts communicating with it on a large network, they concluded that the network traffic and application responsible for it were likely malicious. They expressed that if many hosts communicated with that user agent, they would likely interpret the traffic and application as benign absent any other findings.

Frequency assessment also involved analysts examining the dichotomy between one-to-one, one-to-many, and many-to-many relationships. For example, an analyst who suspected three hosts were compromised asked, "Was there any one-to-many connections taking place from these three systems internally?" Based on prior experience, they knew that a one-to-many internal network communication pattern sourced from a user workstation could indicate that an attacker was performing network reconnaissance. Asking this question and collecting evidence in this form allowed them to leverage frequency-based pattern matching to aid in deciding about the disposition of these hosts.

*Excluding known benign.* While analysts primarily focused on matching patterns associated with malicious activity, they also sought to compare patterns of known benign activity, as shown in Figure 4.5. For example, several analysts described filtering out built-in Windows processes from a list of executed processes because they knew these applications were benign. Analysts hoped that excluding known benign patterns from a data set would reduce the data set to a small size. By reducing the data set significantly, they could apply more scrutiny to the remaining entries through additional pattern matching or external research. Alternatively, it was possible that excluding known benign entities could leave no remaining data for examination and no candidates for further analysis. This result would indicate that no malicious processes executed.



*Figure 4.5.* Analysts excluded known benign to reduce their data set so that they can scrutinize the remaining data for anomalies.

*Timing anomalies.* Analysts identified some anomalies based on their temporal characteristics. For example, one analyst suspected a group of hosts was compromised and sought to assess the network communication activity to and from those systems. They asked, “Is there any other traffic that matches beaconing, just focusing on those three systems to other IP addresses? Do I see any regular traffic, every five minutes or 15

minutes or 30 minutes?” These beaconing patterns of traffic occurring at similar intervals are often associated with malware command and control channels. By seeking evidence of this time pattern, the analysts hoped to find further evidence that the systems were compromised and that malware was actively running.

While analysts sometimes focused on specific pattern matching based on event timing, they also compared the timing of separate events to find anomalies based on time correlation. They drew timestamps from evidence sources where known malicious activity had occurred and took narrow slices of data that occurred proximate to that timestamp from other evidence sources. They scrutinized anything happening around this timestamp further. This tactic was a pervasive entry point into the undirected executed analysis DAT. Whenever analysts found malicious events occurring with a specific host, they frequently pursued a line of inquiry focused on finding anomalous process executions correlating to the same time frame. They worked under the belief that process executions occurring around the same time as the previously discovered malicious event were more likely to be abnormal than events occurring at other times.

*Syntax anomalies.* In some instances, analysts focused on content characteristics within evidence fields to detect anomalies that matched or violated specific syntax conventions. Analysts accumulated a library of syntax heuristics throughout their careers, expecting that individual evidence source fields usually follow specific syntax conventions. For example, an analyst performing undirected execution analysis examined process names a system executed, saying, “Well, the first obvious the low-hanging fruit is some funky file name ... long miscellaneous characters and numbers.” Process names are human-defined values, so analysts expected them to follow human-language syntax

conventions. Things like random sequences of characters, mixed or unexpected character casing, or unclear use of numbers violate these conventions. Additional syntax anomalies included things like the presence of alternate character sets, alternate languages, or unexpected obfuscation or encryption.

This list of anomaly detection techniques does not constitute a complete accounting of all possible methods. However, it may serve as a framework for clustering additional analyst investigative actions through future research. Analysts used many of these techniques in conjunction with one another, and some investigative lines of inquiry relied on several of them. For example, when examining processes executed on a system, analysts frequently employed techniques related to baselining, excluding known benign, and looking for syntax anomalies simultaneously or in succession. Each method provided an opportunity to uncover irregularities that they would not find with other approaches. Furthermore, sometimes differing techniques highlighted the same anomaly, which had an additive effect on the analyst's perception that the finding was malicious.

In most cases, analysts were familiar with the specific fields provided by evidence sources. This familiarity enabled anomaly detection methods because analysts understood the expected type of content for individual fields. However, some evidence sources provided fields that could contain a myriad of data in multiple formats. One example was network packet data, where some protocols may include payload data formatted in countless ways. When encountering data whose formatting was unknown, analysts engaged in lines of inquiry focused on feature detection to make better sense of the data. For example, when discovering that packet data to a suspicious host existed, one analyst asked, "What kind of traffic? Can you be more specific? Is it again encrypted versus

unencrypted?” Analysts performed feature detection like this to characterize unknown evidence data and organize it into individual parts. This exercise helped the analysts understand if anomaly detection was possible using this data and what anomaly detection methods they would use.

### *Network Mapping and Attack Visualization*

As analysts worked their way through investigations, they provided descriptions of how they mapped the network they were defending and visualized its parts as a physical dimension. In these explanations, borders and connections existed at multiple layers of abstraction. For example, sometimes analysts viewed entities at the network level, where individual hosts defined boundaries and transferred data between each other through network links. At a different layer of abstraction, analysts viewed entities as the process level, where individual files defined boundaries and transferred data between each other through mechanisms provided by the operating system. While performing investigations, analysts described observations using movement-based terms. A hard drive did not merely write data onto a disk; a process downloaded a file from an upstream host. Data did not just traverse network cables; the attacker moved laterally from one host to another. These movement-based terms were ubiquitous: up, down, in, out, forward, backward, to, from, around, through, and so on.

The prevalence of movement-based terms and deliberate effort to map digital realm activity to physical realm counterparts indicated that analysts relied on internal visualizations to consider how information flowed through their network and make sense of attack sequences. When asked, analysts confirmed that they frequently visualized their

network as they attempted to make sense of what they were seeing and what steps they might take next. One analyst provided this overview of their visualization process:

The way I visualize it in my head, the only way I can like describe it is like a network graph. Just connecting the points together and just thinking logically, like to get from point A to point B, like what path does it have to travel?

Multiple analysts described similar network/link style visualizations for conceptualizing the defended network and attacker movement.

In most cases, analysts sought to understand the layout and organization of the defended network as a baseline for which to map attacker actions. They asked non-investigative questions to paint a picture of the network. For example, one analyst offered the following action and precursor after discovering a successful compromise:

So they've got some beachheads here from what it looks like. So now what I'd like to try and do is establish what possibly is taking place or has taken place. But real quick to understand because again, we've kind of smoke jumped into this environment, don't have a lot of visibility as far as greater understanding of the network infrastructure or domain structure. Are we dealing with a single domain here? Are we dealing with a flat network or is it segmented? And the reason why I'm asking those, it helps me really just visually understand the domain. So single domain, flat network. If I'm a threat actor, okay, that actually makes my life very easy because now all I have to do is enumerate a single domain. And as far as on the network side I'm working with an easily obtainable network that I don't have to worry about crossing security zones or anything to that extent.

This analyst asked these questions once they suspected lateral movement might be possible after the compromise of an individual host. Likewise, other analysts asked network mapping questions when they required that specific understanding. However, some analysts asked network mapping questions earlier in the investigation, perhaps anticipating this need later.

As part of network mapping, analysts also assessed attack vectors available to an intruder at given points of presence in the network. These combined vectors provided that

analyst with an idea of the attack surface of a given process, host, or network. Analysts described the attack surface as a physical space, akin to a gemstone. A larger attack surface has multiple facets, with each representing some exposed mechanism that an attacker could attempt to compromise. They considered these facets when forming investigative questions. For example, one analyst provided the following statement when trying to determine how an attacker may have compromised a host:

I want to check the level of patches that those machines have. Because I'm thinking in another theory, that maybe could be some kind of, I don't know, vulnerability-related kind of worm, that it's exploiting something on the machines that it's not patched.

Had the system been significantly out of date on software patches, the analyst may have pursued a line of inquiry focused on the likelihood of a software exploitation-based initial compromise. When faced with a similar situation, another analyst asked, "And there's no other way for someone else to remote into the network?" Here, they hoped to determine what external access pathways existed that an attacker could leverage to reach the compromised host from the internet. Each of those remote access pathways represented a potential attack vector that the analyst could explore through inquiry and examination of related evidence sources. In examples like these, analysts visualize hypothetical scenarios that may have occurred given currently known events. This technique allowed them to consider what actions attackers may reasonably take next and where evidence of those actions may appear in evidence. Therefore, attack visualization seems to represent a vital skill for investigative question formation and not just evidence interpretation.

While most analysts indicated that they visualized network movement in their heads, some said they had physically drawn diagrams to understand network events

better. They expressed that this was common, particularly in complex networks or when the attack involved multiple attackers, hosts, or networks. Some analysts also shared that they were likely to physically draw diagrams when sharing their findings with other analysts or investigation stakeholders.

### *Decision-Making Skills*

Making a decision involves leveraging algorithms, heuristics, and rules for choosing among alternatives (Seamster & Redding, 2017). Analysts must make numerous decisions while performing investigations. Reasons for these alternatives include analysts beginning investigations with a minimal understanding of all the events that transpired, meaning they have to consider many potential attacks. Analysts also must contend with numerous evidence sources and the multitude of data points they provide during interpretation. By walking through prior novel investigations that expert analysts experienced and carefully probing their actions, I accounted for facets of their decision-making skills.

The CDM cognitive task analysis technique provided both the data collection and analysis framework for this portion of the study, as well as the construct for describing cognitive skills. The basic unit of analysis for decision-making skills centered around decision points, where analysts chose one action over alternatives. I identified 100 decision points ( $M = 11.1$ ,  $SD = 2.5$ ) across investigations from nine subjects. I expected some degree of variance between subjects since each analyst identified and described a unique investigation with varying events and complexity.

Using this data, I identified features of decision-making skills inherent to the investigation process by performing a CTA-based thematic analysis of the decision points

described by analysts within and across cases. This analysis revealed types of cues analysts discovered during their investigations and how characteristics of those cues led analysts to specific choices that formed lines of inquiry and investigative paths. In this section, I describe two facets of analyst decisions. First, I describe the cues that compelled analysts to make decisions and their characteristics. Second, I explain the goals pursued by analysts at each decision point and their relation to decision cues.

### *Decision Cues*

A cue is any stimulus with implications for action (Wong, 2004). In security investigations, cues primarily come in the form of data obtained from evidence sources, although they may come from human input as well. While cues may be present in an environment, analysts must first find and interpret them to develop courses of action that involve making decisions. For each decision point analysts arrived at in their verbally reported investigation, I asked probing questions to elicit and characterize the cues that led them there. I identified four cue types that led to analyst decisions: relational cues, dispositional cues, novelty cues, and operational cues. It is important to note that these cue types were not mutually exclusive because investigative cues sometimes had more than one meaning when interpreted by analysts. For example, cues often had relational and dispositional characteristics, resulting in classification as both types.

*Relational cues.* These cues indicated the presence of other, yet to be discovered, relationships that were relevant to the attack timeline. For example, one analyst identified that an attacker used the PsExec tool on a system. Since PsExec is a remote command execution tool, its invocation on one system generally implies the presence of a target

system where an attacker attempted to issue a command. Hence, the company of other relationships relevant to what the attacker did.

*Dispositional cues.* These cues indicated whether some event or relationship was suspicious or benign. For example, after discovering an attacker was in control of a service account, an analyst searched for other systems that the service account had accessed after the attacker was in control of it. Their search revealed that the account had authenticated to the network domain controller and other systems hosting sensitive data, providing a disposition as to their status as a part of the attack timeline. Dispositional cues also revealed some relative impact to the organization that warranted a particular response or heightened priority. In this example, the nature of the sensitive data hosted on some of these compromised systems indicated significant organizational impact that influenced analyst decision-making.

*Novelty cues.* These cues indicated the presence of some unknown threat, capability, or technology that the analyst did not understand well. They most frequently arose when analysts discovered indicators whose association, history, or abilities were unclear, like IP addresses, domain names, file hashes, malware family names, or strings. In other cases, analysts encountered novelty when they found legitimate services they did not understand well. For example, one analyst discovered a potentially compromised Microsoft Internet Information Services (IIS) web server during an investigation. They were not familiar with how this service functioned normally, prompting further research into it.

*Operational cues.* These cues came from some external input (not found by the analyst in evidence) and indicated the potential to affect the analyst's ability to conduct the investigation. They were usually related to evidence availability, shifts in organization priorities, and other organizational and social factors affecting the investigation. In one example, an analyst requested a disk image from a system they believed the source of an intrusion. However, the organization would not provide it because they had data privacy concerns relating to the analyst's status as a third-party consultant. The organization ultimately only offered a subset of the host data, which limited the analyst's ability to pursue this line of inquiry and affected the subsequent decisions they made.

#### *Decision Goals*

As part of the CDM knowledge elicitation interviews, analysts described the goals they pursued at each decision point. While the analysts' overarching goal was to identify the disposition of network events and a timeline of potential attacks, each decision represented some more specific goal in pursuit of those broader objectives. My analysis of investigation-related decisions revealed six unique goal types. Analysts chose each investigative action with one of these goals in mind.

*Identify relationships.* The most predominant goal of any analyst decision was the choice to seek out evidence that would identify relationships relevant to the investigation. These decisions constituted over half of all decisions made by analysts, drove most investigative questioning, and defined the majority of each analyst's investigative path. By identifying the start or end of a relationship between two entities or characteristics of those relationships, analysts identified events relevant to an attack

timeline. For example, analysts frequently chose to pursue lines of inquiry that helped them determine which applications executed on a system. A known malicious process running on a system constitutes a relationship between that process and the system with many implications for attack identification and response.

Analysts were most likely to choose to pursue a line of inquiry based on relational cues. They cited goals focused on figuring out what led to a cueing event, followed after it, or happened in conjunction with it. While analysts sometimes pursued the first relationships that came to mind, a single relational cue often presented multiple investigative paths to consider. In these cases, dispositional cues frequently influenced path selection, particularly when the analyst perceived that the potential impact of an attacker action warranted priority discovery so that they could initiate response actions. In other cases, operational cues swayed the analyst's decisions based on the tooling available to them and organizational priorities.

Multiple paths presented difficulties to the analysts' decision-making process when they lacked additional cues to help prioritize one approach over another. In these cases, analysts typically described choosing to pursue relationship identification focused on the investigative path that could either yield the quickest disposition of events or uncover the events with the most significant business impact. Analysts also reported making relationship identification choices based on convenience surrounding their available tooling, the most straightforward data access, and the work they could do themselves without involving other people.

*Assess threat capability.* Analysts were frequently concerned with the potential actions taken by human threat actors or malware on their defended networks. When

encountering an indication of these threats, analysts chose to perform research surrounding their capabilities. These capabilities included threat actors' common tactics, techniques, and procedures, and actions taken by malware in previously discovered samples. Findings from this research often represented relational cues that led analysts to form specific investigative questions that identified relationships between their network and the threat.

Analysts most frequently chose to assess threat capabilities based on novelty cues. In these situations, analysts discovered a link to a threat, such as an IDS alert that referenced it directly or a domain name tied to the threat in internet searches and public threat reporting. Without prior or complete knowledge of the threat, its novelty compelled the analyst to assess its capabilities, which cued additional decisions. Notably, every novelty cue that compelled analysts to assess threat capability was also a relational cue. For an analyst to research threat capabilities, they needed to encounter novelty and some implication that the novel entity had interacted with something on their network. This finding was consistent with expert analysts' desire to ask relevant questions.

*Assess forensic capability.* Some analysts who participated in this study described investigations where they were not employed by the affected organization and took part as a third-party consultant or law enforcement. As a result, they were not entirely familiar with the forensic analysis capabilities available to them. In some cases, analysts chose to deliberately assess these capabilities before taking additional investigative actions. Their steps included inquiring about the availability of specific evidence sources, the prevalence of endpoint forensic agents, the visibility of network collection sensors, and

the availability of other forensic tooling. Aside from technical capabilities, assessing forensic capability also included determining the availability of network and systems support staff to aid in collecting forensic artifacts and facilitating response actions.

Analysts primarily chose to assess forensic capability in two situations. First, they made this decision situationally as they interpreted relational cues. Since they did not know if the evidence source they needed to pursue relationship identification was available, they chose to seek that clarity first. If the evidence source was available, they could proceed with their preferred investigative path. Second, some analysts sought to perform a baseline assessment of forensic capability at the outset of an investigation. One analyst described this process as:

I don't jump right into the investigation. I do spend a little bit of time to understand their environment. So, that way, I can understand what does their security stack look like? What administrative tools do they have? What capabilities does the team have? Because I have to leverage them to do things because again, I'm not going to be able to do everything, so I need to know what type of assistance they can provide.

In cases where the analyst assessed forensic capabilities at the beginning of an investigation, they described having to rely less on situational assessments. However, it did not usually eliminate the need for them completely.

*Perform bulk collection.* Analysts occasionally decided to collect significant amounts of evidence data from one or more systems in anticipation of needing it to answer future investigative questions in addition to current ones. They primarily performed bulk collection when discovering relational or dispositional cues indicating one or more systems were compromised. For example, one analyst found malware running on a system they were focused on. They took an indicator from that malware and

used a log aggregation tool to search for all other systems where it executed. After identifying the affected systems, they ran a custom tool developed by the analyst's firm that retrieved commonly used evidence source data from all these hosts and stored it for later analysis. The tool collected much more information than the analyst initially needed to identify relationships they were immediately concerned with. However, they justified their decision by explaining how it gave them the information they needed now, along with other data they might need in the future. Analysts described going back to bulk collected data frequently as they needed to reference it to answer investigative questions they formed. The bulk collection tools also automated labor-intensive data retrieval tasks, which saved analysts time.

Across cases, the availability of appropriate tooling facilitated analysts' ability to perform bulk collection. However, this availability was not ubiquitous since it typically involved some custom-developed scripts or expensive commercial endpoint collection tools. Some analysts reported wishing they had this tooling at specific points in their investigation but instead having to resort to the manual collection of individual evidence sources. One analyst noted using input from an investigation to build out an automated bulk collection tool tied explicitly to a prevalent malware family they encountered. They used that tool in later investigations, demonstrating how analysts may develop bulk collection tools iteratively over time themselves.

The scope of the investigation also influenced analysts' decisions to perform bulk collection. In one example, an analyst began an investigation with a notification that an attacker had compromised dozens of systems on the network. The analyst explained that it would not be feasible to perform a bulk collection on all these hosts due to the storage

requirements of the collected data and the unlikely scenario that they would need all of it.

They described their course of action, saying:

So, knowing that I am not going to triage 40 to 50 systems, what I did was to identify or limit the scope of analysis around the systems, and ask them for the critical systems to look at. So, with that, we trim the data set down to I believe it was about five systems.

With a more reasonable number of hosts to work with, they initiated bulk data collection for these systems.

*Take response action.* Although the CDM interviews focused on analyst identification of a network compromise, several of them described their decisions to take some containment or eradication action based on what they learned. These actions included things like resetting passwords, blocking connections, purging malicious emails from inboxes, and removing systems or entire networks from the internet. In some cases, these steps required that analysts took action themselves, mainly when the victim organization employed them. In the other cases, analysts reported their recommendation to organization stakeholders who initiated the response actions with their network and systems engineering staff.

Response actions primarily occurred after analysts observed relational or dispositional cues. More specifically, analysts chose to take some immediate response action when they discovered the potential for significant negative impact to network or data confidentiality, availability, or integrity if they did not act quickly. For example, one analyst described blocking malicious file execution using malware file hashes they discovered an attacker using, “So you get the hash, we did have a solution at the time that would block hashes, like software based off of a hash. So we would block list

those.” Their discovery of a phishing attack delivering this malware served as a relational cue, because every user that received that email was a potential victim of the attack if they were to click on the link in the message. By blocking these hashes, they prevented automated delivery of this malware and slowed the attacker down.

Analysts described being careful to initiate or recommend response actions due to the potential impact to an organization or push back from the organization’s leadership. They also anticipated that response actions might compel the attacker to shift tactics or somehow slow the investigation process. Because of these concerns, analysts noted reviewing findings multiple times to make sure their story was straight and that they had arrived at the correct decision.

*Seek advice or help.* Even expert analysts sometimes sought advice or help from colleagues as they worked their way through investigations. They primarily made these decisions when encountering novelty cues related to attacker behavior or unfamiliar technologies. For example, one analyst suspected that an attacker had manipulated an organization’s web server to host a web shell. However, they were unfamiliar with how that might manifest in the web server configuration. They sought advice from another colleague who had experience with that web server software and could point the analyst to the appropriate configuration section to find what they were looking for.

Analysts also chose to seek advice or help when they discovered dispositional cues that represented a significant increase in perceived impact to the organization. These actions were usually tied to the initial incident declaration or the discovery of a potentially damaging event like data exfiltration. When analysts observed these events,

they often engaged colleagues as “another set of eyes” to validate their findings before moving forward.

In the remaining occurrences, analysts sought help as a matter of convenience or resource management. They deferred some portion of the investigation to specialists like malware or threat intelligence analysts. For example, one analyst relied on a colleague to handle forensic collection and analysis of a mobile phone:

I sent the phone back to the lab for someone else. Because the phone takes hours to download because you can't just go, “I want that text message,” in a forensically sound manner, you've got to do a full download of the phone. Because even if you just say when you download, “I only want pictures,” it still does the whole download it just then only gives you the pictures.

This decision freed the analyst up to continue working through the investigation and focusing on the readily available data.

#### *A Model of Diagnostic Inquiry*

While the Simplified PARI and CDM findings provide insight into analyst procedures and decision-making on their own, together, they form an investigation model greater than the sum of its parts. Across expert analysts, forming and answering investigative questions was where their experience and abilities culminated into a process of diagnostic inquiry, initially conceptualized by Sanders (2016a) and depicted with more detail based on these results in Figure 4.6. This process provided analysts with a mechanism to reduce uncertainty and chart an investigative path. Along this path, analysts identified relevant events, discovered their causes and effects, and characterized their dispositions. Ultimately, they arrived at a complete or more thorough understanding of what events transpired so that they could make response recommendations allowing for the containment and eradication of attackers.

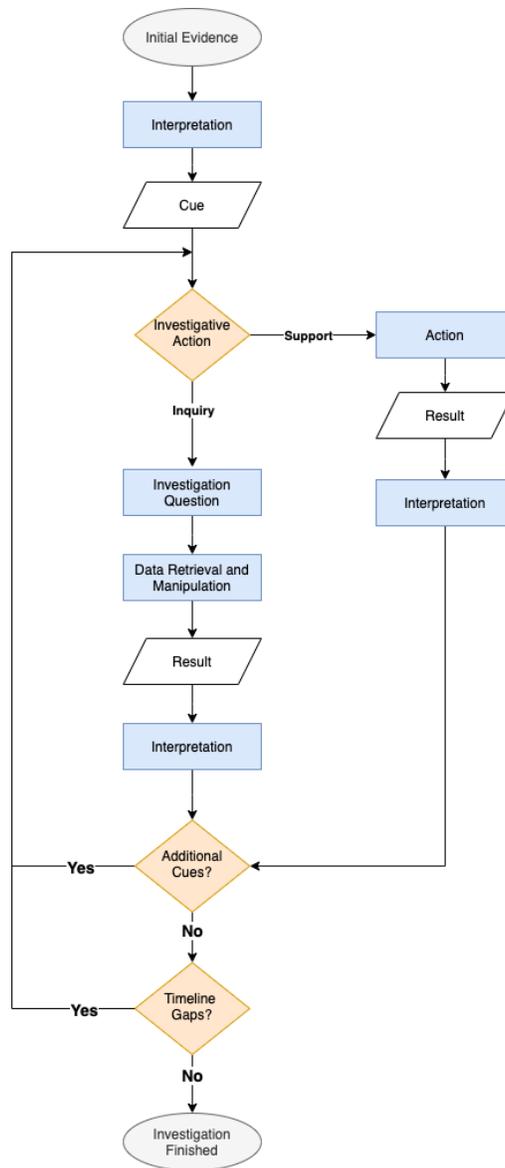


Figure 4.6. Analysts use a process of diagnostic inquiry to reduce uncertainty and make decisions during investigations.

The diagnostic inquiry investigation model represents the relationships between how analysts interpreted evidence, formed investigative questions, and made investigative decisions. Every analyst investigation in this study aligned with this model, and it encapsulates the procedural and decision-making skills observed and reported in this study. Hence, I may refer to the collective skills revealed in this study and their

interactions as a process of diagnostic inquiry where analysts seek to discover events that occurred and diagnose their cause, effect, and disposition. The model uses standard flowchart notation with ovals representing the start and end of the investigation, rectangles representing processes, parallelograms representing input and output, and diamonds as decisions.

Analysts began their investigations with the interpretation of some initial evidence. This initial interpretation is where they brought meaning to the data based on their existing knowledge and experience. Based on this interpretation, they identified one or more decision cues that led to choosing their first investigative action. With that input, they could decide to pursue inquiry or support actions. If they chose inquiry, they formed investigative questions designed to help them identify characteristics of relationships or events relevant to the attack timeline. They retrieved and manipulated data necessary to answer the question, then interpreted that result for meaning and additional cues.

Analysts chose support actions when they did something that indirectly aided or affected the investigation but was not a direct function of inquiry, like assessing forensic capability or taking response actions. After performing either type of investigative action, they determined if additional cues or unresolved timeline gaps were present. If so, the analyst reached another decision point regarding their next investigative action, and the process started again. This process repeated until the analyst had an adequate understanding of the events that transpired and their disposition, until they ran out of cues, until they ran out of answerable investigative questions, or when an operational cue dictated the end of the investigation.

Notably, the diagram in Figure 4.6 is fairly linear, primarily due to analysts' limited ability to perform only a single action at a time. While analysts often considered multiple actions at any given point, they had to decide which action to prioritize. They often prioritized one action knowing what their following action or several actions would be. They usually stuck to this plan. However, in some cases, they interpreted cues that caused them to shift to another line of inquiry or divert to other supporting actions.

### *Implications*

Organizations and academic institutions struggle to train analysts, leading to a shortage of skilled people needed to fill investigative roles (Morgan, 2017; Oltsik, 2019). Those shortages pose problems for society when so much of it depends on computer networks and the analysts who protect them. The cyber security industry is experiencing a cognitive crisis, characterized by a poor understanding of how analysts skillfully perform their jobs. The cognitive underpinnings of successful investigations have been obscure and elusive, making it challenging to teach and learn the analyst's craft (Sundaramurthy et al., 2014).

This cognitive skills analysis provides one of the first and most detailed insights into expert analyst skills and performance during investigations. These results aid comprehension of the cognitive procedures analysts work through to make sense of network attacks while also shining a light on facets of the decision-making that leads them down different investigative paths. When appropriately considered by educators and practitioners, these findings have far-reaching implications that can move the cyber security industry forward in meaningful steps.

By conducting this cognitive task analysis, I have taken steps to make analyst thinking visible to equip practitioners with the language needed to better articulate and describe the investigation process. Using the vernacular of diagnostic inquiry, interested parties may discuss active and past investigations in a way that allows for a clearer shared understanding of the cognitive processes that went on during the investigation. This understanding allows for scrutiny of that process, identifies biases inherent to it, and provides a framework for developing and applying mental models for teaching and learning.

The development of mental models is critical for learning complex fields like cyber security because it allows practitioners to share common experiences and educators to bridge knowledge gaps. Markman and Gentner (2001) describe mental models as “a representation of some domain or situation that supports understanding, reasoning, and prediction” (p. 228). These representations help simplify complex systems so that analysts can learn concepts and make more informed decisions. The findings in this study reveal several new mental models and enhance understanding of existing mental models, including:

- Characteristics of good investigative questions
- Types of investigative questions
- The attack timeline
- Directed analysis techniques
- Analysts’ interaction with evidence
- Common anomaly detection techniques
- Network mapping

- Attack visualization
- Investigative paths
- Common decision cues
- Common decision goals
- The diagnostic inquiry investigation model

These mental models and those derived from them represent immense value to expert analysts, inexperienced analysts, and educators.

For expert analysts, diagnostic inquiry and related models provide a real-world framework that helps them explain what they are already doing. When experts describe their investigations in terms of inquiry, investigative paths, decision cues, and other concepts identified in this research, they map their expertise to clearly defined constructs that allow them to communicate with other analysts more efficiently. This improved ability to think aloud may help experienced analysts refine their craft and bridge the gap between themselves and their less experienced counterparts. For example, an analyst could outline their process working through an investigation by identifying the investigative questions they asked, evidence sources used to answer those questions, and anomaly detection techniques they applied to evidence. They could review these artifacts independently to validate their findings or as part of a group lessons learned session. The latter may have similar benefits to physician's morbidity and mortality conferences (Orlander et al., 2002).

For inexperienced analysts, the mental models described here provide a framework for deliberate practice, facilitating clear learning objectives and feedback (Ericsson, 2008). Each cognitive skill serves as a road marker that allows analysts to set

goals and monitor their progress towards them as they advance in their careers.

Leveraging concepts of deliberate practice focused on enhancing the application of these procedural and decision-making skills can help analysts accelerate their ability to accumulate experience. For example, an inexperienced analyst could work through an investigation scenario while documenting their investigative path using the constructs identified in this study. When they finish, they could compare their investigative path and discovered attack timeline with similar artifacts from an expert analyst. Discrepancies between the path and timeline then become points of discussion and opportunities for meaningful feedback and learning.

For educators, these mental models provide an empirical foundation for the development of cyber security courses. A review of post-secondary and professional cyber security education programs found extreme inconsistency in how institutions designed courses and curriculum (Bicak et al., 2015; Bogolea & Wijekumar, 2004; Lang, 2014). Instructional designers can now reference these identified procedural and decision-making skills when developing investigation-centered education.

The findings of this study may manifest in several ways when incorporated into curricula. Primarily, identified skills that are relatively self-contained may directly represent learning objectives. For example, instructors should teach analysts anomaly detection techniques and measure their retention and application of that knowledge. In other cases, identified skills provide reference models for teaching broader domain knowledge. For example, educators can leverage findings on how analysts interact with evidence to design lessons for individual evidence types that encompass the concepts critical to their investigative uses: interpretation, capability comprehension, collection,

and manipulation. Additionally, some identified skills may allow instructors to connect more general information technology skills to the investigative domain. For example, instructors may design activities that build upon students' knowledge of network architecture diagrams to work through the steps an attacker might take to steal sensitive data. Then, the instructor can describe the evidence sources that would contain artifacts of the attacker's actions. Activities like this start from a place where the student is comfortable (network architecture diagramming) before moving into new territory (evidence sources). The whole time, the student is engaging in a process that enhances their network mapping and attack visualization skills. Bridging related domains like IT and digital forensic analysis is helpful since many analysts may begin their careers or education with IT experience.

In all these examples, curriculum designers and instructors stand to benefit from using these cognitive task analysis results as a base for their work. The identified constructs provide insight into processes that analysts actively use but have poorly understood before now. Additionally, Tofel-Grehl and Feldon's (2013) meta-analysis showed that courses leveraging CTA-based input were more effective than those that did not, particularly when leveraging PARI and CDM techniques for CTA. While others have done work to identify facets of analyst professional roles (Cybersecurity Industry Competency Model, 2019; Newhouse et al., 2017), this work tends to treat these roles as though they are more different than they are the same. While incident response, triage, and forensic analyst roles have unique constraints, these findings identify the common investigative skills inherent to all of them. This work provides an empirical baseline for

analysts and educators to converge on a common investigative framework for digital forensic analysts.

While this cognitive task analysis revealed significant components of digital forensic analyst's skills, like much research, it also raised additional questions for future consideration. I identified some of these questions throughout the findings in this chapter. For example, while I identified a framework for directed analysis techniques and several examples, my conclusions were limited to the Simplified PARI scenario chosen to assess procedural skills. More work is needed to evaluate analysts across a broader array of scenarios to identify additional DATs and enhance the DATs I identified based on clusters of investigative actions.

Some of the skills identified in this study require different research techniques to capture their nuance fully. For example, analyst's anomaly detection processes indicated the presence of some elements of feature detection. A more focused study could allow the researcher to fully characterize that relationship and explain how analysts use feature detection with specific evidence sources. Similarly, the CDM mechanism for elicitation of decision-making skills provided a useful initial step for characterizing analyst decisions, but my findings were limited in their depth. For instance, while I could identify analysts' use of relational cues to pursue relationship identification goals, there may be additional sub-cues and sub-goals associated with these cognitive constructs whose identification would prove useful. In both of these examples, further research would require the use of research techniques beyond cognitive task analysis, including observational methods.

This research was focused exclusively on procedural and decision-making skills. While these were the most crucial skill types to understand for this initial assessment of the field, more work is needed to analyze other skill types, including representational and automated skills. Additionally, this research focused exclusively on experts. Conducting similar research on novice or journey person analysts would provide an opportunity for novice-expert contrasts that may provide insight into the gaps that exist between how experienced and inexperienced practitioners approach job tasks differently.

The skills identified in this research highlighted the broad amount of knowledge required to conduct investigations, but this study did not serve as a complete knowledge inventory for the field. Additional work towards completing a knowledge inventory would provide immense value to the cyber security industry and educators supporting it. Mental models identified in this research may serve a role in designing that research and organizing its findings.

The academic and practitioner research communities focused on identifying the cognitive skills of digital forensic analysts are small. However, the findings in this cognitive task analysis represent an opportunity for increased engagement in this area. With the use of additional research methods, each study helps make analyst thinking more visible. These collective works have the potential to improve digital forensic analyst education and practice dramatically.

### *Summary and Conclusion*

In this chapter, I presented the findings of the Simplified PARI analysis that revealed analyst procedural skills and the CDM analysis that revealed facets of analyst decision-making skills. Then, I described diagnostic inquiry as a foundational model for

how analysts conduct investigations. After that, I discussed the implications of these findings for cyber security practitioners and educators. Finally, I identified opportunities for additional research that builds off the findings in this study.

The presence of overwhelming uncertainty characterizes the analyst experience. It is not merely something they contend with; it defines the entire landscape they operate within. The findings in this cognitive task analysis revealed how analysts navigate this landscape through a process of inquiry. By asking investigative questions, analysts navigate a path through mountains of evidence to build a timeline of events. Along the way, they visualize attacks, identify anomalies, manipulate data, interpret cues, and perform other complex cognitive tasks that leverage their experience and expertise. By identifying these processes and decisions, this research provides concepts and mental models that aid practitioners and educators in moving the digital forensic practice forward in tangible, meaningful leaps.

## CHAPTER FIVE

### Distribution of Findings

#### *Executive Summary*

As the world relies more on computer networks to support even the most basic human needs, so too does it depend on the digital forensic analysts who protect these networks. Organizations task these analysts with detecting and investigating attacks that negatively affect the confidentiality, integrity, and availability of computer networks. By investigating and characterizing these attacks, victim organizations can begin to contain and eradicate attackers, ideally before experiencing harmful impacts.

While the analyst's role is increasing in importance, the cyber security industry is currently in the midst of a cognitive crisis. Despite significant investment in the growing field, the number and severity of attacks against computer networks are outpacing its ability to defend against them. As part of this crisis, 64% of cyber security managers report a lack of skilled practitioners to fill security roles, with most of them stating that the shortage puts their organizations at greater risk (ISC2 Cybersecurity Workforce Study, 2020). Universities and colleges have failed to meet this need and are unable to produce job-ready graduates consistently, and many organizations lack the ability to train new hires themselves.

Many of these workforce development issues exist because much of the analyst's job is poorly understood. Even highly skilled analysts cannot describe how they perform investigations (Sundaramurthy, 2014). This reliance on tacit knowledge in the profession makes it challenging to teach aspiring and novice analysts how to conduct investigations

and interpret evidence. There has been little work that examines analyst cognitive skills at a detailed level, making it difficult for educators to design courses and build curriculum meaningful to analyst job functions (Bell & Oudshoorn, 2018).

In this problem of practice, I used cognitive task analysis research methods to identify skills leveraged by analysts in digital forensic investigations. This work focused on identifying procedural and decision-making skills analysts relied on to build attack timelines. The findings from this study provided a model of diagnostic inquiry that makes the analyst cognitive process explicit. This model and the associated findings provide the foundations for analyst curriculum development by educators and a mechanism by which practicing analysts can better understand their craft. The remainder of this chapter provides an overview of the research methodology, a summary of key findings, informed recommendations stemming from the findings, and a discussion regarding the impactful distribution of these results.

#### *Overview of the Data Collection and Analysis Procedures*

I based this study on qualitative interviews using cognitive task analysis (CTA) methods. Schraagen et al. (2000) described CTA as “a variety of interview and observation strategies that capture a description of the explicit and implicit knowledge that experts use to perform complex tasks” (p. 3). I relied on the Simplified Precursor, Action, Response, Interpretation (PARI) and Critical Decision Method (CDM) knowledge elicitation techniques to collect and analyze data.

The Simplified PARI technique allows researchers to elicit procedural skills from complex, real-world environments with an eye toward curriculum development and operational environments (Hall et al., 1995; Seamster & Redding, 2017). In this research

phase, analysts participated in a tabletop exercise where they worked through an investigation scenario. They thought aloud while answering questions about the processes they used to identify elements of the attack and determine the disposition of network events.

The CDM technique allows researchers to elicit decision-making skills from work environments “characterized by high time pressure, high information content, and changing conditions” (Klein et al., 1989, p. 462). In this phase of the research, analysts described a novel or interesting investigation they took part in at some point during their career. I asked probing questions to identify key decision points during the investigation and characterize unique facets of analyst decisions.

For each technique, I collected data from nine expert analysts who predominantly spend their time performing investigations. This subject pool included three analysts from each typical analyst role: triage analysts, incident response analysts, and forensic examiner analysts. By collecting data across these roles, I identified cognitive skills universal to all of them. The CTA methods chosen for each phase of this research provided the conceptual model necessary to elicit the analyst’s procedural and decision-making knowledge and analyze the results using procedure and decision constructs. I analyzed this data within and across cases to identify an investigative model used by analysts. I also provided analyst subjects with these results and allowed them to provide feedback as a form of member checking and to enhance the credibility and trustworthiness of the results.

### *Summary of Key Findings*

The cognitive task analysis of digital forensic analyst's cognitive skills revealed a diagnostic inquiry-based investigation model. Analysts interpreted evidence to find cues that led them to take investigative actions. Those actions primarily involved forming investigative questions and interpreting evidence to answer those questions. The analysts used those answers to identify additional cues. They repeated this process to discover more about potential attacks and determine if the disposition of network events was malicious or benign. They completed this process when they answered all their investigative questions and no cues remained, until they arrived at an adequate understanding of the attack, or until some external factor caused them to stop.

The Simplified PARI analysis of analyst procedural skills revealed four primary skills: inquiry, evidentiary, anomaly detection, and network mapping and attack visualization skills. Analysts use inquiry skills to form investigative questions, including event-relative, capability matching, and utility questions. Some clusters of questions shared common goals, allowing their classification as directed analysis techniques. Analysts used evidentiary skills to answer questions by interpreting, collecting, and manipulating evidence data. They also used evidentiary skills to assist in question formation by assessing the attestation capability of evidence sources. Analysts performed anomaly detection by employing numerous pattern matching techniques to spot things that could represent malicious activity. They leveraged network mapping and attack visualization skills to conceptualize the network they were defending, understand where attackers may leave evidence after performing specific actions, and characterize the attack surface available to attackers in different locations on the network.

The CDM analysis of analyst decision-making revealed information surrounding the goal of analyst decisions and the cues that led to analyst decisions. I identified four cues that commonly led to analyst decisions: relational, dispositional, novelty, and operational cues. I also characterized five unique goals that defined analyst's decisions. They chose goals based on identifying relationships, assessing threat capabilities, assessing forensic capabilities, performing bulk collection, taking response actions, or seeking advice or help.

The diagnostic inquiry investigation model accounted for the procedural and decision-making skills identified in analysts' investigations. It provided a mechanism for encapsulating and describing how analysts conduct investigations and built attack timelines. Collectively, these results demonstrate new insight into the mind of the analyst.

### *Informed Recommendations*

By identifying procedural and decision-making skills analysts use during investigations, I have taken steps to make components of the analyst's job more explicit. These findings can drastically inform and alter the shape of analyst education in academic and professional settings. Therefore, the most appropriate use of these findings centers on their incorporation into curriculum and course development.

Educators in academic institutions and professional education environments should adopt the diagnostic inquiry investigation model as a fundamental component of the digital forensic analyst curriculum. I recommend that any introductory level digital forensic analysis course include analyst procedures relayed in these findings, emphasizing the formation of investigative questions and network and attack visualization. Where possible, instructors should include learning objectives and

examples related to directed analysis techniques, as these are common among analyst workflows and tie multiple investigative questions into unified goals. This curriculum should also include dedicated instruction focused on developing analyst anomaly detection skills across many scenarios.

These findings highlighted the importance of evidence sources for forming and answering questions. Educators should include learning objectives tied to evidence comprehension throughout analyst training. However, the numerous and diverse nature of evidence sources does not allow for instructors to cover every source or even the most common sources in detail in any single class. Therefore, I recommend that educators also teach conceptual models that analysts can use to comprehend additional evidence sources as they encounter them. These mental models should include elements of evidence interpretation, collection, manipulation, and capability comprehension. The amount of data analysts must contend with may be overwhelming, particularly to new practitioners. It also is essential that educators help analysts understand facets of decision making, including the types of decision cues and goals and the relationship between them.

Finally, organizations employing analysts have a role to play as well. I recommend formalizing diagnostic inquiry into standard operating procedures used to guide analyst's workflow. This step provides analysts with a common vocabulary and framework for describing their investigations. Adopting these concepts places the industry on a trajectory toward a more common analyst experience across organizations, defined by empirical findings from expert analysts rather than the ad-hoc whims of software vendor marketing or biased individual experiences.

These recommendations provide guidance for leveraging the results of this study to better the digital forensic analyst profession. Understanding the cognitive skills that make expert digital forensic investigations possible is an essential step along that path. The next step is applying those findings to how we educate analysts and facilitate their work. The combined efforts of the academic and professional communities can help cyber security emerge from its cognitive crisis.

### *Findings Distribution Proposal*

The findings presented in this study represent a significant step forward in the collective understanding of how digital forensic analysts perform investigations. With this in mind, the following sections describe considerations for the distribution of these results, including the target audience, proposed distribution method and venue, and potential distribution materials.

#### *Target Audience*

I conducted this study primarily with the distribution of findings to analyst instructors and instructional designers in mind. Individuals in these roles will benefit the greatest by understanding how expert analysts think in the context of analyst job performance. With this information, they can design curriculum and instructional activities that will move aspiring or less experienced analysts toward an expert practice of the craft.

I also will distribute these findings directly to analysts. Because of the ad-hoc nature of education in the cyber security industry, many practitioners spend time on self-education outside of formal schools or courses. By providing these findings in a format

digestible to practitioners directly, they have a more significant opportunity to adopt and apply diagnostic inquiry and the related concepts in a directly impactful manner.

### *Proposed Distribution Method and Venue*

As necessity often is the mother of invention, so too was this study. I spent most of my career as a digital forensic analyst before becoming an analyst instructor.

Therefore, I sought to develop a better understanding of expert analyst performance to incorporate these findings into multiple formats that allow them to better reach those seeking to become analysts or increase their expertise in performing investigations.

This distribution plan includes multiple methods and venues for educators and practitioners alike. First, I will publish this Problem of Practice dissertation in its entirety. This step will ensure the research findings are accessible to a variety of stakeholders. Educators and practitioners alike can interpret these findings, incorporate them into their work, and build upon them as necessary. Second, I will incorporate these findings into multiple vehicles for direct analyst instruction, including incorporation into analyst-focused courses that I teach. Lastly, it also includes the eventual publication of a book that seeks to help analysts understand how to perform investigations based on the cognitive task analysis of expert practitioners.

### *Distribution Materials*

The findings of this study provide the most meaningful impact when shared through the development of educational materials. These exemplar materials provide mechanisms that leverage the elements of diagnostic inquiry and include components like lectures, formative assessments, participatory exercises, and lab scenarios. By constructing and sharing these materials, I provide a starting point for educators and

analysts alike to begin learning diagnostic inquiry principles and applying them to real-world investigations. This approach better prepares analysts to consciously leverage procedural and decision-making skills observed in experts.

### *Conclusion*

In this cognitive task analysis, I probed the experience of expert digital forensic analysts to reveal the procedural and decision-making skills they leverage to perform investigations. This research revealed a model of diagnostic inquiry that encompasses and explains several procedural and decision-making skills in a complex domain. These findings provide one of the first and most significant insights into the analyst's mind and provide educators and practitioners with mental models necessary to advance the profession.

## APPENDICES

## APPENDIX A

### Analyst Screening Survey

I used these survey items to screen applicants for enrollment into the study, as described in the Analyst Sampling section of Chapter Three I identified the answer format of the question by using square brackets.

Q1. How would you rate your expertise as an investigator? [Multiple Choice / Select One]

- Junior
- Intermediate
- Expert

Q2. How many years of experience do you have in information security? [Numeric Drop-Down Selection]

Q3. How many years of experience do you have in network or systems administration? [Numeric Drop-Down Selection]

Q4. What types of investigations do you normally conduct? (Select all that apply) [Multiple Choice]

- Alert Triage/Event Analysis/SOC Analysis
- Forensic Examination
- Malware Reverse Engineering
- Threat Intelligence
- Incident Response

- Threat Hunting

Q5. What percent of your time at work is spent actively spent performing security investigations? That includes threat hunting, alert investigation (SOC), incident response, threat actors (intelligence), or malware reverse engineering. [0-100 Selection]

Q6. Please rate how strongly you agree with the following statements regarding your use of investigative tools. [Rated on a 5-point Likert scale]

- I'm comfortable using the tools required for analysis relevant to my work.
- I can explain the function and purpose of the tools required for analysis relevant to my work.
- I often develop or customize tools to solve problems related to my work.

Q7. Please rate how strongly you agree with the following statements regarding your heuristic knowledge. [Rated on a 5-point Likert scale]

- I rarely see anomalous activity I can't explain.
- I can think of at least ten ways to hunt for evil on the network.
- I can think of at least five ways to prove a program executed on a system.
- I know the most common ways attackers exploit systems to gain an initial foothold on to a target network.
- I know the most common ways attackers "live off the land" to extend their control in a compromised network.

Q8. Please rate how strongly you agree with the following statements regarding your confidence. [Rated on a 5-point Likert scale]

- When I encounter something I've never seen before, I know how to find the answers to move forward.
- I'm nervous in unfamiliar investigative scenarios.
- I'm confident investigating attacks that I've never seen before.

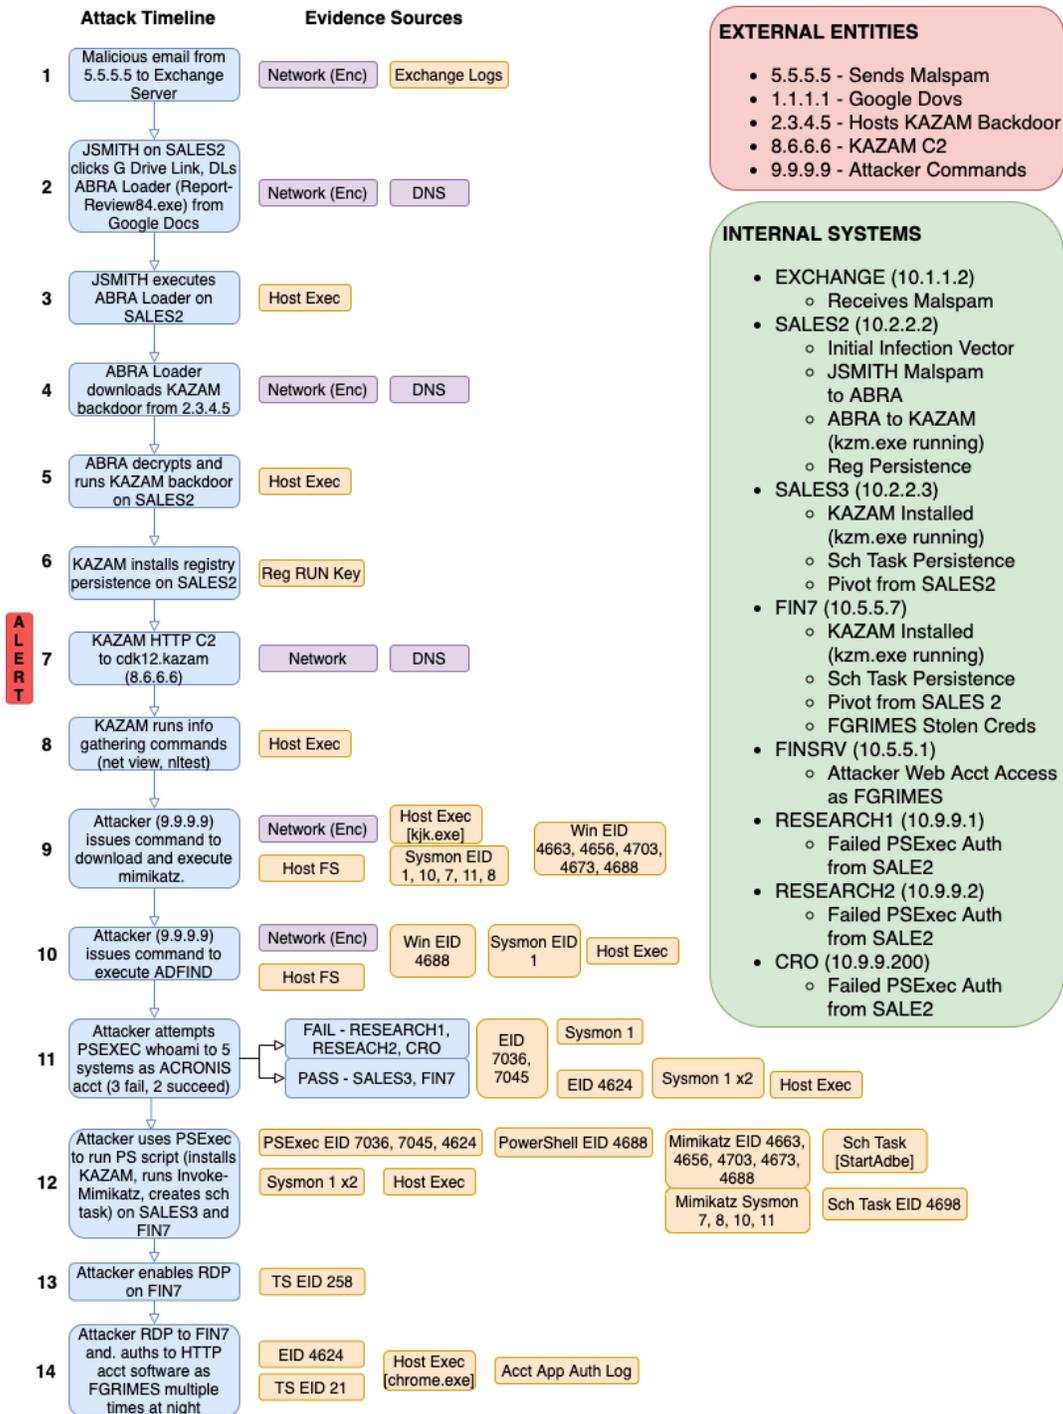
- I never feel intimidated in investigation scenarios.

Q9. Please rate how strongly you agree with the following statements regarding your evidentiary knowledge. [Rated on a 5-point Likert scale]

- I have an in-depth understanding of how operating systems work.
- I'm comfortable examining Windows operating system logs.
- I'm comfortable examining Linux operating system logs.
- I'm comfortable performing file system analysis.
- I have an in-depth understanding of how network communication works.
- I'm comfortable examining packet captures.
- I'm comfortable examining firewall and/or flow logs.
- I have an in-depth understanding of how memory works.
- I'm comfortable performing analysis on memory captures.
- I'm comfortable performing dynamic malware analysis.
- I'm comfortable performing static malware analysis.

# APPENDIX B

## PARI Investigation Scenario



APPENDIX C

Consent Form

Baylor University  
**Department of Education**

Consent Form for Research

PROTOCOL TITLE: **Cognitive Task Analysis of Digital Forensic Analysts**  
PRINCIPAL INVESTIGATOR: **Chris Sanders**

**Invitation to be part of a research study**

You are invited to be part of a research study focusing on digital forensic practitioners' cognitive skills. This consent form will help you choose whether or not to participate in the study. Feel free to ask if anything is not clear in this consent form.

**Why is this study being done?**

The purpose of this study is to gain a better understanding of the cognitive skills used by expert digital forensic analysts. These skills include decision-making and procedural skills that comprise the investigation process.

**What will happen if I take part in this research study?**

If you agree to take part in this study, you will be asked to **participate in three tasks:**

**TASK ONE:** Participate in a web-based interview where you complete a verbal tabletop exercise based around a fictional investigation scenario. (Approximately 1 hour)

**TASK TWO:** Participate in a web-based interview where you describe a complex or interesting investigation you worked through within the past five years. (Approximately 1 hour)

**TASK THREE:** Review my summarized findings from the analysis of your interviews and provide specific feedback based on prompts I will provide by email (Approximately 1 hour)

These tasks will occur at separate scheduled times.

Audio recording is required for this study. Video recording is optional for this study. If you do not want to be recorded at all, you should not be in this study. If you would like to be audio recorded only, you should indicate that at the beginning of your interview.

**How long will I be in this study and how many people will be in the study?**

Participation in this study will last approximately 3 hours spread across several weeks. This study will use nine initial subjects, with more added as necessary.

**What are the risks of taking part in this research study?**

I don't believe there are any risks from participating in this research.

**Are there any benefits from being in this research study?**

You might benefit from being in this study because the published results may help you better understand your thought processes.

**How Will You Protect my Information?**

A risk of taking part in this study is the possibility of a loss of confidentiality. Loss of confidentiality includes having your personal information shared with someone who is not on the study team and was not supposed to see or know about your information. I plan to protect your confidentiality.

I will keep the records of this study confidential by using unique numbers instead of your name to store and reference the recorded interview and interview transcripts. I will leverage best practices for ensuring the confidentiality of data at rest and data in transit. I will make every effort to keep your records confidential. However, there are times when federal or state law requires the disclosure of your records.

Representatives of Baylor University and the BU Institutional Review Board may review your study records for purposes such as quality control or safety.

The anonymized audio-only recording of your interview will be shared with a third-party service that will convert the recording to a text transcript.

The results of this study may also be used for teaching, publications, or presentations at professional meetings. If your individual results are discussed, your identity will be protected by using a code number or pseudonym rather than your name or other identifying information.

**Will information you collect about me be used for future research studies?**

Information collected from you as part of this research may be shared with the research community at large to advance science and health. I will remove or code any personal information that could identify you before the information is shared with other researchers to ensure that, by current scientific standards and known methods, no one will be able to identify you from what is shared.

### **Will I be compensated for being part of the study?**

You will receive a scholarship to an Applied Network Defense online training course of your choosing (limited to courses taught by Chris Sanders). Course access will be granted upon completion of the study in 2021. You must complete every task of the study to receive the scholarship.

### **Your Participation in this Study is Voluntary**

Taking part in this study is your choice. You are free not to take part or to withdraw at any time for any reason. No matter what you decide, there will be no penalty or loss of benefit to which you are entitled. If you decide to withdraw from this study, the information that you have already provided will be kept confidential. You cannot withdraw information collected prior to your withdrawal.

If you are a Baylor student or faculty/staff member, you may choose not to be in the study or to stop being in the study before it is over at any time. This will not affect your grades or job status at Baylor University. You will not be offered or receive any special consideration if you take part in this research study.

### **Contact Information for the Study Team and Questions about the Research**

If you have any questions about this research, you may contact:

Chris Sanders  
Phone: [Redacted]  
Email: [Redacted]

Or

Dr. Sandi Cooper  
Phone: [Redacted]  
Email: [Redacted]

### **Contact Information for Questions about Your Rights as a Research Participant**

If you have questions about your rights as a research participant, or wish to obtain information, ask questions, or discuss any concerns about this study with someone other than the researcher(s), please contact the following:

Baylor University Institutional Review Board  
Office of the Vice Provost for Research  
Phone: [Redacted]  
Email: [Redacted]

## Your Consent

### **SIGNATURE OF SUBJECT:**

By electronically signing this document, you are agreeing to be in this study. I will give you a copy of this document for your records. I will keep a copy with the study records. If you have any questions about the study after you sign this document, you can contact the study team using the information provided above.

*I understand what the study is about and my questions so far have been answered. I agree to take part in this study.*

## APPENDIX D

### Directed Analysis Techniques

#### *Prevalence Analysis*

Analysts performed prevalence analysis when identifying what proportion of a population shared one or more characteristics. The most common goal of prevalence analysis was to determine what other hosts may be affected by malicious activity that the analyst already identified. Analysts did this to identify the scope of the investigation based on which hosts experienced similar events. Archetypical questions included:

- Is [Malicious Process] running on any other system on the network?
- Did any other user receive an email with [Phishing Message Characteristics]?

Analysts also used prevalence analysis to facilitate anomaly detection based on the frequency of occurrence. Analysts used this technique to determine the disposition of entities and events working from the assumption that widespread events were less likely to be malicious and rare events were more likely to be malicious. Archetypical questions included:

- Did any other system on the network talk to [IP Address]?
- What other hosts resolved [Domain Name]?
- How many other hosts on the network use [Web Browser]?

Another goal of prevalence analysis was to characterize events to help analysts properly order the attack timeline. For example, a malicious event only occurring on one

system could indicate it was the initial point of compromise, while that event occurring on multiple systems may mean that the attacker targeted them using a similar lateral movement mechanism after already establishing an initial foothold.

Investigative questions that analysts used with prevalence analysis always included at least one characteristic as the basis for the prevalence search. However, some questions also had additional bounding elements to limit the scope of the investigation. For example, analysts typically limited prevalence inquiries based on the hosts that were of interest to them: a single host, hosts within a department, internal or external hosts, or hosts matching some other criteria. These archetypes provide an example of this scoping:

- Was [User Account] used to authenticate to any of [List of Systems]?
- Are there any login attempts to [Suspect System] using [Attacker-Controlled User Name]?

#### *Directed Execution Analysis*

The goal of directed execution analysis was for analysts to prove whether a specific process executed or determine some characteristics of that execution. This execution included files directly executed by the operating system (like Windows EXE files) or files opened by another process (like DOCX files opened by the Microsoft Word process). Analysts frequently encountered observations that suggested a system executed a process. They used information obtained from these observations or through additional threat capability research to form investigative questions centered on a characteristic of the process and the file linked to it, like its name or hash. Archetypical questions for proving execution included:

- Are there any event log entries indicating [File Name] executed?
- Is there any record of software blocking the execution of [File Hash]?

- Is there a file indicating that [File Name] executed?
- Is [Process Name] actively running on this system?
- Which process made a DNS Query for [Domain Name]?

After identifying suspicious process execution, analysts often looked for things that would help them identify additional relationships and events based on their characteristics. Archetypical questions included:

- When was the first and last time that [File Name] executed?
- How many times has [File Name] executed?
- What is the parent process of [File Name]?
- Were other processes are running at the same time as [Process Name]?
- What command did [Process Name] run?
- Was the [File Hash] executed as an elevated user?

### *Undirected Execution Analysis*

Analysts performed undirected execution analysis when examining all of the process executions during a specific time window to determine if one or more anomalous executions occurred. This technique is unique from directed execution analysis because the analyst did not have a characteristic of a specific process or file to base their investigative questions on. Instead of trying to prove the execution of a particular process, they performed anomaly detection to attempt to find the execution of any suspicious process. In most cases, analysts performed undirected execution analysis based on the time when other suspicious events occurred. Archetypical questions included:

- Did any other suspicious process execute just before or after [Time of Suspicious Event]?
- Are any other suspicious processes currently running right now?
- Were there any newly executed processes within a few hours after [When User Received Phishing Message]?

In other cases, analysts performed undirected execution analysis when they suspected that an attacker had compromised a system but had not yet confirmed it or did not yet comprehend the full scope of the attack. Without event-relative or capability matching questions to drive this inquiry, analysts usually formed investigative questions around common attacker techniques observed in other investigations. Archetypical examples included:

- Does evidence indicate the execution of any processes that have never been seen before around [Time of Suspicious Event]?
- What non-default Windows processes are running right now?
- Did PowerShell.exe run within a few minutes before or after [Time of Suspicious Event]?
- Have any processes executed from a user download or desktop directory?

Once analysts identified the execution of suspicious processes, they frequently sought additional information about that execution. Since they now had characteristics of the process to help form investigative questions, these searches now met the criteria for directed execution analysis.

### *Reputation Analysis*

Analysts performed reputation analysis when assessing the disposition of an entity based on reporting of its past behavior or relationships. The goal of reputation analysis

was to help form an opinion about whether events related to the entity were likely to be malicious or benign. Archetypical questions included:

- What is the reputation of [IP Address]?
- Has anyone reported [Domain Name] as malicious?
- Has [IP Address] ever been associated with malware command and control?
- Does [Domain Name] appear on any email blocklists?

Analysts performed reputation analysis when encountering unknown external IP addresses and domains in a large variety of situations. In most cases, they performed this analysis on individual entities once they suspected their disposition might be malicious. However, in some cases, analysts performed a bulk analysis of entities to help find evidence of additional malicious activity. For example, one analyst suspected a host was compromised and then performed bulk reputation analysis of every external IP address that host communicated with over a proximate time frame.

### *Malware Capability Analysis*

Analysts performed malware capability analysis when they sought to understand how suspected malware functions. In most cases, they identified characteristics of the malware for the eventual purpose of asking capability matching questions that helped them determine if the malware was active on one or more systems. In other cases, analysts were not yet confident that a suspicious file was malware, so they performed capability analysis to decide the file's disposition. Analysts typically used this DAT after encountering evidence that a system executed a suspicious process, a suspicious file existed on a system, or after observing some other malware characteristic. They may

have based their concern on characteristics of a process or file itself or the context surrounding its execution or existence. In any of those cases, their focus eventually turned to a suspicious file. Analysts usually had at least one of four possible items to base their continued investigation on: the file name, the file hash, the malware/family name, or the file itself.

When analysts observed the name of a suspicious file, they asked the following archetypical questions:

- What does an internet search on [File Name] tell me about its capabilities?
- When was [File Name] first and most recently seen on networks?
- What does [File Name] do?
- Does [File Name] serve a legitimate purpose?

Analysts often identified files uniquely with their cryptographic hash. They substituted that value into the above questions or leveraged the hash and file name to identify the common name of the malware if it was not already known.

Analysts also described the concept of malware families, which are collections of malware variants that function almost identically or were likely developed by the same person or organization. They asked the following archetypical questions when they discovered a malware name or family name:

- What does an internet search on [Malware/Family Name] tell me about its capabilities?
- What does [Malware/Family Name] do?
- What time frame was [Malware/Family Name] active? Is it still active?

In cases where analysts could recover the file itself, they leveraged it to better understand its capabilities by executing it (or using a service to execute it) in a controlled environment while observing its interaction with the system. Question archetypes included:

- What behaviors does [File] exhibit when executed in a sandbox?
- What indicators of compromise does [File] produce when executed that I can search for?

### *Lateral Movement Analysis*

Analysts performed lateral movement analysis to determine if an attacker leveraged their access on one victim host to compromise other systems within the network. Their goal was to identify those other compromised systems and the mechanism(s) used for lateral movement. In most cases, analysts formed broad investigative questions that might locate any relationship between a compromised system and other hosts on the network. This included archetypes like:

- Was there any connection between [Compromised Host] and other internal hosts?
- What internal systems did [Compromised Host] communicate with after [Process] executed?
- What systems did [Compromised User Account] authenticate to?
- Are there any authentication attempts to [Suspected Target Systems] that were not their typical user?

There are dozens of techniques attackers use to move laterally within networks. Therefore, analysts also formed investigative questions whose goal was to identify specific characteristics of these techniques they were aware of or had directly observed in past investigations. These questions included focusing on techniques like lateral tool

transfer, remote services use, and use of alternate authentication material. Some observed archetypes include:

- Is there evidence confirming [Tool] successfully interacted with [Target System]?
- Were there any remote desktop protocol connections to internal systems from [Compromised Host]?
- Were there remote authentication attempts to any systems that do not usually receive these attempts?
- Were any unique authentication mechanisms used on systems they are not generally used on?

### *Staging and Exfiltration Analysis*

Analysts performed exfiltration analysis to determine if an attacker stole sensitive information from compromised network hosts. This technique includes inquiry related to attackers moving sensitive data around the network (staging) and pulling it off the network (exfiltration). In cases where analysts confirmed that attackers had compromised systems, they often first sought to assess the sensitivity of data that attackers could access given their current privilege level and the nature of the hosts they controlled. This strategy helped them determine if they should perform additional exfiltration analysis.

Archetypical questions included:

- What is the role of [Compromised System]?
- Is any sensitive data hosted on [Compromised System]?
- Does [Compromised User] have access to any sensitive data?

When analysts decided to look for signs of exfiltration, they pursued lines of inquiry centered around common behavioral techniques observed in previous attacks.

These questions included the following archetypes:

- What files were accessed on [Compromised System] during the time frame of the compromise?
- Did anyone access a large number of files in a short time frame on [Compromised System]?
- Were there any suspicious archive files created on [Compromised System]?
- Are there any uncommon large outbound transfers coming from any hosts on the network?

More than most techniques, analysts performed exfiltration analysis absent any indication that exfiltration occurred. Instead, they cited performing this analysis because of the potential impact exfiltration could have on their organization and its customers or constituents.

### *Phishing Analysis*

Analysts performed phishing analysis when attempting to locate phishing messages or understand their role in an attack. The most common occurrence in this scenario was when analysts suspected phishing as an initial attack vector. In these cases, they took what information they knew about the attacker's initial foothold into the network or early attack timeline events and probed for phishing messages sent to that user. Archetypes included:

- Did [User Account] receive any suspicious emails within a few hours prior to [Malicious Observation]?
- Do any messages to [User Account] have suspicious subjects?
- Do any messages to [User Account] have suspicious attachments or links?
- Did [User Account] receive any messages from suspicious domains?
- Did [User Account] receive any messages from spoofed domains?

- Did [User Account] receive any messages from accounts they have never received a message from before?
- Did [User Account] receive any messages containing content formatting anomalies?
- Did any messages contain a link to [Malicious Website Contacted by Compromised System]?
- Did any messages contain an attachment of [Type of File Executed by Compromised System]?

In the occasion that one or more of these questions led to the analyst's identification of a relevant phishing message, they reformed the above archetypes to focus them on the identification of malicious artifacts from a single message rather than a large set of them.

#### *Account Role Analysis*

Analysts performed account role analysis when they sought to understand the normal function of a user account and its role within an organization. In these instances, they often encountered an attacker using an account for malicious purposes or identified a user account associated with a compromised system. By understanding the role of the account, they could determine if the behavior they observed was typical for the account, controlled by the attacker, or merely adjacent to the attack. They also used this information to determine the potential impact of the attack. Archetypical questions included:

- What groups is [User Account] a member of?
- Is [User Account] a member of any administrator groups?
- Is [User Account] a service account?
- Is [User Account] a local account or domain account?

- When was [User Account] created?
- What user created [User Account]?
- When was [User Account] last used?

In cases where the analyst associated the user account with a human user (as opposed to a service account), they formed investigative questions to identify that person's role. This strategy helped them better characterize the disposition of observed events. That included these archetypes:

- What functional part of the organization is [Human User] part of?
- What are the typical working hours for [Human User]?
- What time zone does [Human User] work from?

## BIBLIOGRAPHY

- Aitken, R. (2018, August 19). *Global information security spending to exceed \$124B in 2019, privacy concerns driving demand*. Forbes.  
<https://www.forbes.com/sites/rogeraitken/2018/08/19/global-information-security-spending-to-exceed-124b-in-2019-privacy-concerns-driving-demand/>
- Andone, D. (2019, October 11). *3 Alabama hospitals are accepting patients again after a ransomware attack on its computers*. CNN.  
<https://www.cnn.com/2019/10/11/us/alabama-hospital-ransomware-attack/index.html>
- Bejtlich, R. (2004). *The Tao of network security monitoring: Beyond intrusion detection*. Pearson Education.
- Bell, S., & Oudshoorn, M. (2018, October). Meeting the demand: Building a cybersecurity degree program with limited resources. In *2018 IEEE Frontiers in Education Conference* (pp. 1–7). IEEE.
- Bicak, A., Liu, X. M., & Murphy, D. (2015). Cybersecurity curriculum development: Introducing specialties in a graduate program. *Information Systems Education Journal*, *13*(3), 99–110.
- Bogolea, B. & Wijekumar, K. (2004). Information security curriculum creation: A case study. *Proceedings of the 1st Annual Conference on Information Security Curriculum Development* (pp. 59–65). ACM.
- Bureau of Labor Statistics. (2018). *Job outlook - information security analysts data set, 2018 to 2028* [Data set]. <http://data.bls.gov>
- Cattermole, V., Horberry, T., Burgess-Limerick, R., Wallis, G., & Cloete, S. (2015). Using the critical decision method and decision ladders to analyse traffic incident management system issues. *Proceedings of the 2015 Australasian Road Safety Conference* (pp. 1–11). Australasian College of Road Safety.
- Cavusoglu, H., Mishra, B., & Raghunathan, S. (2004). A model for evaluating IT security investments. *Communications of the ACM*, *47*(7), 87–92.  
<https://doi.org/10.1145/1005817.1005828>

- Chao, C. J., & Salvendy, G. (1994). Percentage of procedural knowledge acquired as a function of the number of experts from whom knowledge is acquired for diagnosis, debugging and interpretation tasks. *International Journal of Human-Computer Interaction*, 6(3), 221–233.  
<https://doi.org/10.1080/10447319409526093>
- Charmaz, K. (2006). *Constructing grounded theory: A practical guide through qualitative analysis*. Sage.
- Cichonski, P., Millar, T., Grance, T., Scarfone. (2012). *Computer security incident handling guide*, NIST Special Publication, 800-61. U.S. Department of Commerce, National Institute of Standards and Technology.
- Clark, R. E. (2009). How much and what type of guidance is optimal for learning from instruction? In S. Tobias & T. M. Duffy (Eds.), *Constructivist instruction: Success or failure* (158–183). Routledge.
- Clark, R. E., Feldon, D., van Merriënboer, J., Yates, K, and Early, S. (2007). Cognitive task analysis. In J. M. Spector, M. D. Merrill, J. J. G. van Merriënboer, & M. P. Driscoll (Eds.), *Handbook of research on educational communications and technology* (3rd ed., pp. 577–594). Lawrence Erlbaum Associates.
- Cooper, P., Finley, G. T., & Kaskenpalo, P. (2010). *Towards standards in digital forensics education*. 2010 ITiCSE Working Group Reports, June 26–30, 2010, Bilkent, Ankara, Turkey, 87–95.
- Crandall, B., Klein, G., Klein, G. A., & Hoffman, R. R. (2006). *Working minds: A practitioner's guide to cognitive task analysis*. MIT Press.
- Creswell, J. (2013). *Qualitative research inquiry and design: Choosing among five approaches*. Sage.
- Creswell, J. W., & Poth, C. N. (2016). *Qualitative inquiry and research design: Choosing among five approaches*. Sage.
- Cybersecurity industry competency model. (2019).  
<https://www.careeronestop.org/CompetencyModel/competency-models/cybersecurity.aspx>
- Cybrary Catalog (2019). <https://www.cybrary.it/catalog/>
- Daggett, S. (2010). *Quadrennial defense review 2010: Overview and implications for national security planning*. Congressional Research Service.
- Denzin, N. K., & Lincoln, Y. S. (Eds.). (2011). *The Sage handbook of qualitative research*. Sage.

- Downs, D. (2017). *ISACA state of cybersecurity 2017*. ISACA, Schaumburg, IL. [http://www.isaca.org/Knowledge-Center/Research/Documents/state-of-cybersecurity-2017\\_res\\_eng\\_0217.pdf](http://www.isaca.org/Knowledge-Center/Research/Documents/state-of-cybersecurity-2017_res_eng_0217.pdf)
- Ellis, B. L. (2009). *The human analysis element of intrusion detection: A cognitive task model and interface design and implications* [Doctoral Dissertation, Nova Southeastern University]. ProQuest Dissertations Publishing.
- Ericsson IoT Forecast. (2016, November 9). <https://www.ericsson.com/en/mobility-report/internet-of-things-forecast>
- Ericsson, K., & Charness, N. (1994). Expert performance: Its structure and acquisition. *The American Psychologist*, 49(8), 725–747. <https://doi.org/10.1037/0003-066X.49.8.725>
- Ericsson, K. A. (2004). Deliberate practice and the acquisition and maintenance of expert performance in medicine and related domains. *Academic Medicine*, 79(10), S70–S81.
- Ericsson, K. A. (2008). Deliberate practice and acquisition of expert performance: a general overview. *Academic Emergency Medicine*, 15(11), 988-994. <https://doi.org/10.1111/j.1553-2712.2008.00227.x>
- Ericsson, K. A., Hoffman, R. R., Kozbelt, A., & Williams, A. M. (Eds.). (2018). *The Cambridge handbook of expertise and expert performance*. Cambridge University Press.
- Exec. Order No. 13870, 3 C.F.R. 20523 (2019).
- Feldon, D. (2004). *Inaccuracies in expert self-report: Errors in the description of strategies for designing psychology experiments* [Doctoral Dissertation, University of Southern California]. ProQuest Dissertations Publishing.
- Fepac.org. (n.d.). <https://www.fepac-edu.org/>
- Filkins, B. (2016). *IT security spending trends*. SANS Institute.
- Foresman, B. (2019, July 26). *Universities are expanding cybersecurity education to meet broad demand*. <https://edscoop.com/universities-cybersecurity-education-expanding-workforce-demand/>
- Garcia, C. (2015). *Using cognitive task analysis to capture how expert anesthesia providers conduct an intraoperative patient care handoff* [Doctoral Dissertation, University of Southern California]. ProQuest Dissertations Publishing.
- Gazarian, P. K., Henneman, E. A., & Chandler, G. E. (2010). Nurse decision making in the prearrest period. *Clinical Nursing Research*, 19(1), 21–37. <https://doi.org/10.1177/1054773809353161>

- Hall, E. M., Gott, S. P., & Pokorny, R. A. (1995). *A procedural guide to cognitive task analysis: The PARI methodology* (AL/HR-TR-1995-0108). United States Air Force.
- Harenčárová, H. (2015). Structured analysis of critical decision method data—emergency medicine case study. *Human Affairs*, 25(4), 443. <https://doi.org/10.1515/humaff-2015-0036>
- Hargrove, R., & Nietfeld, J. (2015). The impact of metacognitive instruction on creative problem solving. *The Journal of Experimental Education*, 83(3), 291–318. <https://doi.org/10.1080/00220973.2013.876604>
- Hayden, E., Assante, M., & Conway, T. (2014, August). *An abbreviated history of automation & industrial controls systems and cybersecurity*. SANS Institute.
- Hutchins, S. G., Pirolli, P. L., & Card, S. K. (2004). *A new perspective on use of the critical decision method with intelligence analysts*. Naval Postgraduate School Dept of Informational Sciences.
- ISC2 Cybersecurity Workforce Study. (2020). <https://www.isc2.org/Research/Workforce-Study>
- Information Assurance Directorate at the National Security Agency. (2019). *CAE-CD knowledge units*. [http://www.iad.gov/NIETP/documents/Requirements/CAE-CD\\_2019\\_Knowledge\\_Units.pdf](http://www.iad.gov/NIETP/documents/Requirements/CAE-CD_2019_Knowledge_Units.pdf)
- Joint Task Force Transformation Initiative. (2012). *Guide for conducting risk assessments*, NIST Special Publication, 800-30. U.S. Department of Commerce, National Institute of Standards and Technology.
- Klein, G. A., Calderwood, R., & Macgregor, D. (1989). Critical decision method for eliciting knowledge. *IEEE transactions on systems, man, and cybernetics*, 19(3), 462–472.
- Konoske, P. J., & Ellis, J. A. (1986). *Cognitive factors in learning and retention of procedural tasks*. Navy Personnel Research and Development Center.
- Krebs, B. (2013, December 18). *Sources: Target investigating data breach*. <https://krebsonsecurity.com/2013/12/sources-target-investigating-data-breach/>
- Krebs, B. (2013, December 20). *Cards stolen in target breach flood underground markets*. <https://krebsonsecurity.com/2013/12/cards-stolen-in-target-breach-flood-underground-markets/>
- Krebs, B. (2014, May 6). *The Target breach, by the numbers*. <https://krebsonsecurity.com/2014/05/the-target-breach-by-the-numbers/>

- Krebs, B. (2017, September 17). *Breach at Equifax may impact 143M Americans*. <https://krebsonsecurity.com/2017/09/breach-at-equifax-may-impact-143m-americans/>
- Kushner, D. (2013). The real story of Stuxnet. *IEEE Spectrum*, 3(50), 48–53.
- Lang, A. (2014). *A new portable digital forensics curriculum* [M.S. Dissertation, University of Illinois at Urbana-Champaign]. University of Illinois.
- Lang, A., Bashir, M., Campbell, R., & DeStefano, L. (2014). Developing a new digital forensics curriculum. *Digital Investigation*, 11, S76–S84.
- Lincoln, Y. S., & Guba, E. G. (1985). *Naturalistic inquiry*. Sage.
- Lindner, F., (DATE). “On Hackers and Academia,” keynote, *European Conference on Computer Network Defense*, 2010; <http://2010.ec2nd.org/program/keynote2>.
- Lipshitz, R., & Strauss, O. (1997). Coping with uncertainty: A naturalistic decision-making analysis. *Organizational behavior and human decision processes*, 69(2), 149–163.
- Liu, J. (2006). Developing an innovative baccalaureate program in computer forensics. In *Proceedings of the 36th ASEE/IEEE Frontiers in Education Conference*. October 28–31, 2006, San Diego, CA.
- Markman, A. B., & Gentner, D. (2001). Thinking. *Annual Review of Psychology*, 52(1), 223–247.
- McLean, R. (2019, July 30). *Capital One data breach: A hacker gained access to 100 million credit card applications and accounts*. CNN. <https://www.cnn.com/2019/07/29/business/capital-one-data-breach/index.html>
- Merriam, S. B., & Tisdell, E. J. (2015). *Qualitative research: A guide to design and implementation*. John Wiley & Sons.
- Mertens, D. M. (2003). Mixed methods and the politics of human research: The transformative-emancipatory perspective. In A. Tashakkori, C. Teddlie (Eds.) *Handbook of Mixed Methods in Social and Behavioral Research* (3<sup>rd</sup> e., pp. 135–164).
- Morgan, S. (2017, June 8). *Cybersecurity jobs report 2018–2021*. <https://cybersecurityventures.com/jobs/>
- Monster Jobs. (2019). <https://www.monster.com/>
- National Centers of Academic Excellence. (2019). <https://www.nsa.gov/resources/students-educators/centers-academic-excellence/>

- Newhouse, W., Keith, S., Scribner, B., & Witte, G. (2017). *National initiative for cybersecurity education (NICE) cybersecurity workforce framework*, NIST Special Publication, 800-181. U.S. Department of Commerce, National Institute of Standards and Technology.
- Nieles, M., Dempsey, K., & Pillitteri, V. (2017). *An introduction to information security*, NIST Special Publication, 800-12. U.S. Department of Commerce, National Institute of Standards and Technology.
- Obama, B. (2017, September 25). *Presidential memorandum for the Secretary of Education*. <https://www.whitehouse.gov/presidential-actions/presidential-memorandum-secretary-education/>
- O’Gorman, G., & McDonald, G. (2012). *Ransomware: A growing menace*. Symantec.
- Oltsik, J. (2019, January 10). *The cybersecurity skills shortage is getting worse*. <https://www.csoonline.com/article/3331983/the-cybersecurity-skills-shortage-is-getting-worse.html>
- Oltsik, J. (2020, July 31). *The life and times of cybersecurity professionals 2020*. <https://www.esg-global.com/research/esg-research-report-the-life-and-times-of-cybersecurity-professionals-2020>
- Orlander, J. D., Barber, T. W., & Fincke, B. G. (2002). The morbidity and mortality conference: The delicate nature of learning from error. *Academic Medicine*, 77(10), 1001–1006.
- Palmer, G. (2001). *A road map for digital forensics research. Digital forensic research workshop (DFRWS) technical report (DTR) T001-01 final*. MITRE. <http://www.dfrws.org/2001/dfrws-rm-final.pdf>
- Pauley, K., Flin, R., Yule, S., & Youngson, G. (2011). Surgeons’ intraoperative decision making and risk management. *The American Journal of Surgery*, 202(4), 375–381.
- Pillay, H., Brownlee, J., & Wilss, L. (1999). Cognition and recreational computer games: Implications for educational technology. *Journal of Research on Computing in Education*, 32(1), 203–216.
- Pluralsight Cyber Security Courses (2019). <https://www.pluralsight.com/browse/information-cyber-security>
- Ponemon Institute. (2019). *The cost of cybercrime*. Accenture.
- Reigeluth, C. (1983). Current trends in task analysis: the integration of task analysis and instructional design. *Journal of Instructional Development*, 6(4), 24–35. <https://doi.org/10.1007/BF02906215>

- Riggle, J. D., Wadman, M. C., McCrory, B., Lowndes, B. R., Heald, E. A., Carstens, P. K., & Hallbeck, M. S. (2014). Task analysis method for procedural training curriculum development. *Perspectives on Medical Education*, 3(3), 204-218.
- Ryan, C. (2018). *Computer and internet use in the United States, 2016*. U.S. Department of Commerce, Economics and Statistics Administration, U.S. Census Bureau.
- Sanders, C. (2016a). *How analysts approach investigations with diagnostic inquiry*. <https://chrissanders.org/2016/05/how-analysts-approach-investigations/>
- Sanders, C. (2016b). *The effects of opening move selection on investigation speed*. <https://chrissanders.org/2016/09/effects-of-opening-move-investigation-speed/>
- Sanders, C. (2017). *A cognitive psychology approach to computer security investigations*. Presented at the Art into Science Conference, Austin, TX.
- Sanders, C., & Rand, S. (2019). *Creative choices: Developing a theory of divergence, convergence, and intuition in security analysts*. <https://chrissanders.org/2019/10/creative-choices-paper/>
- Sanders, C., & Smith, J. (2013). *Applied network security monitoring: Collection, detection, and analysis*. Elsevier.
- SANS Cyber Security Courses (2019). <https://www.sans.org/courses/>
- Schraagen, J. M., Chipman, S. F., & Shalin, V. L. (Eds.). (2000). *Cognitive task analysis*. Psychology Press.
- Seamster, T. L., & Redding, R. E. (2017). *Applied cognitive task analysis in aviation*. Routledge.
- Smith, M. (2016, March 23). *Three more hospitals hit with ransomware attacks*. <https://www.csoonline.com/article/3047180/three-more-hospitals-hit-with-ransomware-attacks.html>
- Sundaramurthy, S. C., Bardas, A. G., Case, J., Ou, X., Wesch, M., McHugh, J., & Rajagopalan, S. R. (2015). A human capital model for mitigating security analyst burnout. *Symposium on Usable Privacy and Security*, 347–359. USENIX.
- Sundaramurthy, S. C., McHugh, J., Ou, X. S., Rajagopalan, S. R., & Wesch, M. (2014). An anthropological approach to studying CSIRTs. *IEEE Security & Privacy*, 12(5), 52–60.
- Tofel-Grehl, C., & Feldon, D. (2013). Cognitive task analysis–based training: A meta-analysis of studies. *Journal of Cognitive Engineering and Decision Making*, 7(3), 293–304. <https://doi.org/10.1177/1555343412474821>

- Tucker, A., & Belford, G. (2019, March 14). *Computer science: definition, fields, & facts*. <https://www.britannica.com/science/computer-science>
- U.S. Cyber Command History. (n.d.). <https://www.cybercom.mil/About/History/>
- Wong, B. L. W. (2004). Data analysis for the critical decision method. In D. Diaper & N. Stanton (Eds.), *The handbook of task analysis for human-computer interaction* (pp. 327–346). CRC Press.
- Xia, F., Yang, L. T., Wang, L., & Vinel, A. (2012). Internet of things. *International Journal of Communication Systems*, 25(9), 1101–1102.
- Zetter, K. (2016, March 3). *Inside the cunning, unprecedented hack of Ukraine's power grid*. Wired. <https://www.wired.com/2016/03/inside-cunning-unprecedented-hack-ukraines-power-grid/>