

July 2018

Student Misconceptions about Cybersecurity Concepts: Analysis of Think-Aloud Interviews

Julia D. Thompson

San Jose State University, julia.thompson@sjsu.edu

Geoffrey L. Herman

University of Illinois, Urbana-Champaign, glherman@illinois.edu

Travis Scheponik

University of Maryland, Baltimore County, tscheponik@gmail.com

Linda Oliva


University of Maryland, Baltimore County, oliva@umbc.edu

Alan Sherman

University of Maryland, Baltimore County, sherman@umbc.edu

See next page for additional authors

Follow this and additional works at: <https://digitalcommons.kennesaw.edu/jcerp>

 Part of the [Educational Assessment, Evaluation, and Research Commons](#), [Information Security Commons](#), [Management Information Systems Commons](#), [Scholarship of Teaching and Learning Commons](#), and the [Technology and Innovation Commons](#)

Recommended Citation

Thompson, Julia D.; Herman, Geoffrey L.; Scheponik, Travis; Oliva, Linda; Sherman, Alan; Golaszewski, Ennis; Phatak, Dhananjay; and Patsourakos, Kostantinos (2018) "Student Misconceptions about Cybersecurity Concepts: Analysis of Think-Aloud Interviews," *Journal of Cybersecurity Education, Research and Practice*: Vol. 2018 : No. 1 , Article 5.

Available at: <https://digitalcommons.kennesaw.edu/jcerp/vol2018/iss1/5>

This Article is brought to you for free and open access by DigitalCommons@Kennesaw State University. It has been accepted for inclusion in Journal of Cybersecurity Education, Research and Practice by an authorized editor of DigitalCommons@Kennesaw State University. For more information, please contact digitalcommons@kennesaw.edu.

Student Misconceptions about Cybersecurity Concepts: Analysis of Think-Aloud Interviews

Abstract

We conducted an observational study to document student misconceptions about cybersecurity using thematic analysis of 25 think-aloud interviews. By understanding patterns in student misconceptions, we provide a basis for developing rigorous evidence-based recommendations for improving teaching and assessment methods in cybersecurity and inform future research. This study is the first to explore student cognition and reasoning about cybersecurity. We interviewed students from three diverse institutions. During these interviews, students grappled with security scenarios designed to probe their understanding of cybersecurity, especially adversarial thinking. We analyzed student statements using a structured qualitative method, novice-led paired thematic analysis, to document patterns in student misconceptions and problematic reasoning that transcend institutions, scenarios, or demographics. Themes generated from this analysis describe a taxonomy of misconceptions but not their causes or remedies. Four themes emerged: overgeneralizations, conflated concepts, biases, and incorrect assumptions. Together, these themes reveal that students generally failed to grasp the complexity and subtlety of possible vulnerabilities, threats, risks, and mitigations, suggesting a need for instructional methods that engage students in reasoning about complex scenarios with an adversarial mindset. These findings can guide teachers' attention during instruction and inform the development of cybersecurity assessment tools that enable cross-institutional assessments that measure the effectiveness of pedagogies.

Keywords

Adversarial thinking, Cybersecurity Assessment Tools (CATS), cybersecurity education, information assurance, misconceptions, novice-led thematic analysis

Cover Page Footnote

ACKNOWLEDGMENTS This work was supported in part by the U.S. Department of Defense under CAE-R grants H98230-15-1-0294 and H98230-17-1-0349 at UMBC, and H98230-15-1-0273 and H98230-17-1-03473 at the University of Illinois; and by the National Science Foundation under SFS grant 1241576 at UMBC.

Authors

Julia D. Thompson, Geoffrey L. Herman, Travis Scheponik, Linda Oliva, Alan Sherman, Ennis Golaszewski, Dhananjay Phatak, and Kostantinos Patsourakos

INTRODUCTION

As educators struggle to meet the growing demand for cybersecurity professionals (Frost & Sullivan 2015; Libicki, Senty, & Pollak, 2014), there is a corresponding awareness that we lack a rigorous research base that informs how to educate the workforce. There have been calls to develop rigorous assessment tools to measure how teaching practices help students learn cybersecurity (Burley, et al., 2014). In response to these calls, we initiated the *Cybersecurity Assessment Tools (CATS)* project (Sherman et al., 2017a). In this paper, we present a critical early step in creating rigorous assessment tools: documenting how students reason about cybersecurity concepts. No formal study previously explored student cognition and reasoning about cybersecurity.

Cybersecurity lies at the confluence of several disciplines, including computer science, engineering, information systems, networks, cryptography, human factors, and policy. Cybersecurity is an evolving field with new concepts and methods invented on an ongoing basis. Consequently, cybersecurity is a conceptually complex discipline with many facets and perspectives. As for any complex discipline, identifying a core conceptual framework can guide and inform studies about how students learn to think within the discipline. For example, in introductory mechanics courses, studying how students understand the core concept of *force*, alone, can provide critical insights into barriers to student learning (Hake, 1998; Hestenes, Wells, & Swackhamer, 1992). We postulate that *adversarial thinking* is such a core useful conceptual framework for reasoning about cybersecurity. Therefore, we focus our study on how students develop and use adversarial models to guide their reasoning about security scenarios and on what misconceptions students hold about fundamental cybersecurity concepts.

Adversarial thinking involves reasoning about actions and goals in a context in which there might be bad actors attempting to defeat those goals and carry out their own nefarious actions (Bodeau, Fabius-Green, & Graubart, 2010; Caltagirone, Pendergast, & Betz, 2013; Duggan, Thomas, Veitch, & Woodard, 2007; Mateski et al., 2012). Such reasoning requires an understanding of the goal requirements, as well as an understanding of who are the bad actors and what are their objectives, resources, access, capabilities, knowledge, motivations, and risk tolerance. It also requires a technical understanding of the computer systems and their potential vulnerabilities. Our Delphi processes revealed that adversarial thinking, and the associated management of trust and information in computer systems and networks, comprise the core concepts of cybersecurity (Parekh et al., 2017).

Expanding on Scheponik et al. (2016), this paper describes our second major step in a four-step effort to create CATS. In fall 2014, as the first step, we carried out two Delphi processes to identify core cybersecurity concepts (Parekh et al., 2017). In the second step, we developed twelve cybersecurity scenarios based on the Delphi results (Sherman et al., 2017b). In 2016, we also interviewed students who had just completed at least one course in cybersecurity as they reasoned about prompts derived from the twelve scenarios to identify problematic reasoning patterns and misconceptions. The findings from this research will inform the third step, developing questions for the CATS. In the fourth step, the authors will validate the concept inventory using expert review, cognitive interviews, and psychometric testing.

We assume that the reader is familiar with the basics of cybersecurity, information assurance, communications security, and cryptology, including the CIA triad (Confidentiality, Integrity, Availability) and authentication. For an introduction to these concepts, see Bishop (2003); Kim & Solomon (2014); Schneier (1996); Sherman et al. (2017b).

The main contribution of this paper is the analysis of student misconceptions about cybersecurity, leading to a taxonomy of misconceptions. The purpose of this study is to identify themes (or patterns) in student misconceptions and problematic reasonings. We do not examine the causes of, or potential remedies to, these student misconceptions as the causes may vary by categories or even particular misconceptions. Instead, we focus on documenting which categories of misconceptions appear across a variety of demographics, instructional methods, topics, concepts, and course objectives. The primary value of these overarching patterns is twofold: 1) the broad applicability of these categories can inform the design of assessment tools that should be broadly applicable and 2) all instructors can readily identify that students in their courses possess misconceptions or problematic reasonings that align with these categories, helping to guide their attention to students' problems. These categories can then be used to guide more targeted studies that aim to explore causality and remedies of particular misconceptions or categories of misconceptions.

In alignment with this purpose, we used a qualitative approach that let us richly describe the range of student misconceptions and reasonings to identify the patterns that transcended local idiosyncrasies. We sampled students from a variety of institutions with varying curricular goals and teaching methods. Our qualitative analysis methods enable richness of description and a basis for theory building to the domain of analysis (Borrego, Douglas, & Amelink, 2009). We begin by reviewing relevant background literature on student misconceptions and how students learn. Next, we more fully describe our research questions and methods. We then present four common themes (overgeneralizations, conflated concepts,

biases, and incorrect solutions) that describe the commonalities in student reasoning across institutions, scenarios, and topics. Finally, we suggest potential implications of our findings on teaching practices and suggest new avenues for research that could be explored. Appendix A provides our interview prompts.

BACKGROUND AND PREVIOUS WORK

In this section, we first provide a brief overview of our underlying theories of human cognition that guided our study design and then describe our current understanding of student misconceptions in cybersecurity, demonstrating the necessity of the current study.

This study is grounded in constructivist theory, which asserts that humans construct and interpret an understanding of reality that is different from reality based in the physical world (Guba & Lincoln, 1994). For example, Patton (2005) describes the physical sun as being “real,” while the perception of the sun is not real in an absolute sense. The experience of looking at the sun and feeling the heat on the skin is real to people, but the perception of the sun is influenced by individual life context and culture. It is through experiencing the natural world that humans develop conceptions, and with them, misconceptions. For example, children’s misconceptions are built on their intuitive interactions with the world, such as believing the world is flat or that heavier objects will fall faster than lighter objects (Chi, 2005; Vosniadou, 2007).

A core difference between experts and novices within a domain of knowledge is their ability to organize their knowledge around a core set of concepts (Bransford, Brown, & Cocking, 2000). While novices usually attend to the surface features of problems (e.g., the presence of an incline plane), experts attend to the deeper, abstract conceptual aspects of problems (e.g., Newton’s laws) (Chi, Feltovich, & Glaser, 1981). This conceptual knowledge enables experts to identify the important features of a problem more quickly, making them easier to solve, more easily transfer their knowledge to new situations, and accelerate their learning of new domain knowledge (Bransford, Brown, & Cocking, 2000). Misconceptions or a lack of conceptual knowledge hinders these critical skills.

Experts debate the cognitive structure of misconceptions, as some appear to be more robust or theory-like, while others appear to be more chaotic and unpredictable (Chi, 2005). For example, Özdemir & Clark (2007) note that it is commonly believed that students hold to a naïve theory about physics that resembles the now rejected Impetus Theory. In contrast, in computing contexts student conceptions of state or implication can vary dramatically from context to context and from problem to problem (Herman, Loui, & Zilles, 2012; Perkins & Martin, 1986). This fragmentation aligns with diSessa’s Knowledge in Pieces

theory (diSessa, Gillespie, & Esterly, 2004). Simon (1996) argued that student cognition is much more fragmented in computing contexts because students must wrestle with the “science of the artificial” rather than with the science of the physical world.

Because cybersecurity exists at the intersection of the physical and digital worlds, it is not immediately clear whether student understanding of cybersecurity concepts will be more coherent, drawing from years of experience with security concepts in the physical world, or incoherent, drawing from limited experiences with the digital world. This study begins to fill a void in the literature about the cognitive structure of concepts in cybersecurity.

Regardless of the cognitive structure of novices, experts in a field are known to be able to access their knowledge more effectively because it is organized into a coherent conceptual structure. In cybersecurity, we posit that “adversarial thinking,” including the ability to organize a scenario into an adversarial model, is one such conceptual structure.

No prior research has documented student misconceptions about cybersecurity concepts nor how students use adversarial models to guide their reasoning, necessitating careful and critical observation of patterns in students’ reasoning before grounded hypotheses and experimental studies can be made. The NICE framework (NIST, 2013) and professional certification tests, such as CISSP (CISSP, 2016), provide a basis for identifying standards in terminology, information, and notation, but they do not tell us about how students learn or reason about cybersecurity concepts. Similarly, the 2013 IEEE/ACM Computing Curriculum articulates some learning goals that institutions may want to adopt for their students, but this curriculum does not provide any guidance for how students learn those topics (Joint Task Force on Computing Curricula, 2013).

A variety of ongoing initiatives are taking place to help meet the demand for cybersecurity professionals: The National Initiative for Cybersecurity Education (NICE) framework articulated a common lexicon for cybersecurity education and jobs (NIST, 2013), and the IEEE/ACM Computing Curriculum added cybersecurity content to the undergraduate curriculum in computing (Joint Task Force on Computing Curricula, 2013; ACM/IEEE/AIS/IFIP Joint Task Force, 2017). Educators and stakeholders meet at conferences to discuss the challenges (e.g., the Colloquium for Information Systems Security Education (CISSE); National Cyber Summit; WISE; Bishop, Fletcher, Miloslavskaya, & Theocharidou, 2017). The National Science Foundation created the CyberCorps: Scholarship for Service (SFS) program, and the Department of Defense created the Information Assurance Scholarship Program (IASP).

METHODS

An observational study seeks to describe phenomena richly and rigorously. Our goal in this study is to identify patterns that transcend individual contexts yet are still identifiable and relatable (e.g., a taxonomy of marine animals should be drawn by observing animals from multiple bodies of water and generate descriptions that help others classify those animals). These observations can later be used to inform research questions that seek to examine causality or describe classes of phenomenon in more depth (e.g., research into how marine animals breathe would be confounded if whales and dolphins were incorrectly taxonomized as fish). Consequently, our methods focus on deeply observing the behaviors of a diverse group of students to identify patterns that emerge across that diversity.

Subjects

We recruited subjects ($n=26$), who had taken at least one course in cybersecurity and who ranged in experience and educational goals. Students in targeted cybersecurity courses were given recruitment fliers and those who agreed to participate were interviewed at their home institution. The subjects included five females and 21 males who ranged in experience and educational goals. Thirty-five percent were enrolled in a certificate program or Associates degree programs, 24 percent were enrolled in a Bachelors degree program, 31 percent were enrolled in a Masters degree program and 8% were enrolled in a Doctoral degree program. Ninety percent of subjects had relevant work or intern experience ranging from 6 months to 12 years.

The subjects attended one of three universities: University of Maryland Baltimore County (UMBC) ($n=12$), a public research institution (UMBC); Bowie State University (BSU) ($n=4$), a Historically Black University; and Prince Georges Community College (PGCC) ($n=9$), a community college that focuses on vocational training and preparing students for four-year institutions. These three institutions force our observations to hold true across students with different demographics and curricula with different priorities, teaching methods, and learning objectives.

The goal of sampling is to reach saturation of observations. This saturation occurs when additional interviews provide no new insights. This approach works for our study since we are concerned with the breadth of observations rather than their prevalence among different student demographics (e.g., gender, ethnicity). Previous similar studies by (Herman, Loui, & Zilles, 2011; Herman, Zilles, & Loui, 2012) found saturation occurred between eight and ten interviews. We can

determine whether saturation is reached by whether new themes emerged or were identified as we analyzed new interviews.

Data Collection

To create a diversity of observations of student knowledge across domain content, we developed 12 scenarios to reflect core concepts in cybersecurity identified in our Delphi process (Parekh et al., 2017). We designed the scenarios to be authentic, relevant, concise, and engaging, and we structured the prompts to allow students to have freedom of response in describing their approaches and ideas. We organized the scenarios into three protocols (Alpha, Bravo, Charlie), with each protocol including scenarios embodying a range of ideas and difficulty levels. Appendix A gives the interview prompts for each of our 12 scenarios and shows how they are organized into the three protocols. Our companion paper (Sherman et al., 2017b) gives exemplary responses to six of our favorite scenarios.

We presented students from a variety of institutions with four different cybersecurity scenarios each and interviewed each student for approximately one hour as they reasoned about those scenarios. We video- and audio-recorded the 26 interviews and then analyzed them using the novice-led thematic analysis method (Montfort, Herman, Brown, Matusovich, & Strevler, 2013). UMBC's Institutional Review Board approved the interview and analysis protocols.

During the interviews, the interviewer read the prompt and provided it in written form to the subject, for some scenarios augmented with a diagram. We also prepared a detailed guide for the interviewer to help her respond to students answers that reflected any possible level of mastery. Specifically, we structured this guide as a hierarchy of suggestions for dealing with “hits” (correct answers) and “misses” (incorrect answers).

The lead interviewer was a professor with expertise in educational research, extensive experience with conducting cognitive interviews, and minimal experience with cybersecurity. Because interviews are a didactic process, the interactions between the interviewer and the student may bias the student responses. Since our goal was to encourage students to engage in extended descriptions of their thinking even if they were uncertain of the correctness of their answers. To minimize the likelihood that students might feel embarrassed about an answer or be fearful of being judged, the interviewer shared her lack of experience in cybersecurity with the subject at the beginning of the interview and reassured students that the focus of the student was not an evaluation of their individual performance. This method also minimized the likelihood that the student would try to “fish” for correct answers from the interviewer or suspect

that follow-up questions implied that the student gave a wrong answer. In classroom settings, experts often only ask follow-up questions if students answered a question incorrectly, so positioning the interviewer as a novice suggests to the student that follow-up questions are sincere requests for more information rather than Socratic, evaluative questions meant to highlight an error.

While this positioning may encourage more active dialogue between the interviewer and interviewee, it poses a few critical limitations. First, the interviewer may not realize that certain nuances in student understanding need to be further explored. Second, the interviewee may “dumb down” their answers to avoid confusing the novice interviewer. To compensate for these limitations, other members of the research team with expertise in cybersecurity were also present during the interviews. These members were allowed to ask questions only when invited by the lead interviewer at the end of each scenario. They could follow-up on any answers that an expert perceived needed further exploration and engage students in more sophisticated dialogue that could reveal the extent of student knowledge.

Each interview lasted approximately one hour. We conducted the interviews using a think-aloud format (Ericsson & Simon, 1993). After the subject provided their initial response to a scenario, the lead interviewer prompted them to provide more detail and explanation of their approaches. After each prompt, the lead interviewer allowed plenty of time for subjects to articulate their thought process and provide a rationale for their approaches. The prompting continued until the subject had no further clarification or justification for their response.

We audio- and video-recorded the interviews, photographed notes taken by the students, and transcribed the audio recording (with help by undergraduate assistants). We then analyzed the transcriptions, audio files, and notes as described in the next subsection. Table 1 details the number of students receiving each interview protocol at each school.

School	Interview Protocol			
	Alpha	Bravo	Charlie	Total
UMBC	4	4	4	12
PGCC	2	3	4	9
BSU	1	1	2	4
Total	7	8	10	25

Table 1: Number of students interviewed by school and interview protocol.

Novice-Led Paired Thematic Analysis

Three researchers analyzed the interviews using a novice-led thematic analysis (Montfort et al., 2013): two with expertise in the content area and one with relatively minimal knowledge in the content area and expertise in qualitative research methods. The “novice” led the analysis and learned the content while interpreting and analyzing the data. Having the novice lead the analysis provides two primary benefits that enable richer observations and descriptions of student misconceptions. First, because the novice is likely to be more familiar with how cybersecurity concepts and terms are used in colloquial contexts, they will be better able to identify the conceptual and linguistic knowledge that students are using to build their knowledge of cybersecurity. Second, experts frequently suffer from what is called “expert blind-spot,” becoming so familiar with content knowledge that they fail to recognize when their explanations or justifications make assumptions a novice does not understand (Nathan, Alibali, & Koedinger, 2005; Nathan & Petrosino, 2003). This blind-spot can be problematic when analyzing student responses to technical conceptual questions because the expert may add missing information into a question or scenario without realizing they are doing so. These benefits of novice-led thematic analysis enable the team to perceive more nuances in the data while the presence of experts on the analysis team ensures that findings reflect accepted disciplinary knowledge.

In accordance with this process, the researchers familiarized themselves and calibrated their understandings of the data through an iterative cycle of independent analysis and collaborative discussion. To orient the content novice toward the appropriateness of student responses without overly influencing her interpretation, the content experts read through the interviews and labeled sections coarsely as either “correct” or “incorrect,” and provided a few words of explanation as needed. The novice began her analysis with these initial codes in mind. The researchers then discussed and came to a consensus on interpretations of each excerpt of student incorrect answers, more fully fleshing out why, in what ways, and to what extent responses were incorrect, and tracking interpretations in a spreadsheet. The process allowed the researchers, with differing perspectives, to challenge each other’s assumptions.

From these interpretations, themes of misconceptions and student reasoning emerged. In accordance to Braun & Clarke (2006), initial patterns and themes focused on a subset of the data (Alpha interview responses) and then expanded to include the remaining two interview protocols (Bravo and Charlie). The team reviewed the themes based on relevance to the coded data set, defined with substantive examples, and checked through conversations among co-authors.

These steps validated and added richness and context to the themes. Because no new themes were identified when analyzing the Bravo and Charlie interviews, we believe that we indeed reached saturation of observations with our sample.

FINDINGS

Four themes of student misconceptions emerge from the interviews: overgeneralizations, conflation, biases, and incorrect assumptions (including limiting the adversary and failure to see vulnerabilities). We now explain each of these themes and illustrate them with examples.

Theme 1: Overgeneralization

An overgeneralization is a type of misconception in which a person takes a concept that is valid in one context and inappropriately applies it to a new context. For example, a student who is learning decimal notation may think that 0.09 is greater than 0.1 because numbers with more digits are always greater, or that 9 is always greater than 1 (Sackur-Grisvard & Leonard, 1985; Steinle, 2004). While these conceptions may be true for integers or single digit integers, respectively, these conceptions of numbers are misconceptions when applied to decimal representations.

We observed students make inappropriate generalizations about most technologies and security protocols that they discussed. We illustrate this misconception in three examples: encryption, biometrics, and the security of the Internet. Some other concepts students generalized include multi-factor authentication, passwords, security badges, security questions, text messages, hash functions, and man-in-the-middle attacks.

The first example involves encryption, which is a process that mixes plaintext in a complicated way with a sequence of bits known as the key. The resulting ciphertext looks like gibberish. With knowledge of the key, a recipient can reverse the process to yield the plaintext. If the process is secure, the adversary (who is not given the key) cannot recover the key or plaintext.

Many students overgeneralize and form misconceptions by assuming that encryption achieves additional properties beyond confidentiality: preventing manipulation, protecting against theft, and ensuring availability. For example, when asked how encryption can support the security of a drone, a student incorrectly stated that encryption prevented alteration and theft of the data, in addition to protecting confidentiality:

[There is] high reliability that something encrypted is going to remain untampered. It increases that reliability and the assurance that something that has not been altered or stolen or read or whatever.

The notion that encryption prevents data from being “read” is correct, since it focuses on confidentiality, but connecting encryption to alteration and theft is incorrect. Integrity is the term used to ensure that a file has not been altered. Integrity can be established by computing a hash-function or a Message Authentication Code, such as keyed-Hash Message Authentication Code, along with or without encryption. It is possible that the student’s statement about reliability also reflects confusion between encryption and error-correcting coding.

The second example involves biometrics, which involve authentication protocols that are based in human biology, such as fingerprints, DNA samples, or retinal scans. Students overgeneralized by considering biometrics to be a panacea for improving security of any system. While biometrics may increase authentication strength in some systems, they do not always do so. They also have their own weaknesses and limitations (whether ethical or practical) that need to be considered.

For example, biometrics can be forged as a password can be guessed or stolen. However, whereas a user can change a compromised password, a user cannot change their compromised biometrics. Storing biometrics can also pose ethical dilemmas, requiring users to surrender otherwise private personal information. Students failed to consider these types of vulnerabilities and limitations. For example, when asked about increasing the security of a voting system, one student suggested including a DNA sample:

Interviewer: What would that form look like? What might that method be that ensures that the person showing up to vote is who he or she says they are?

Subject: Signatures and DNA testing. You can put a signature, but that may be referring to someone else instead. That’s why you need DNA clarification to ensure that it is you who is voting, not someone else.

This student’s response is troubling because they fail to consider the limitations of DNA testing and they describe DNA testing with language (i.e., “need” and “ensure”) that implies that their suggestion is fail proof. The origin of this overgeneralization is unclear, whether students developed these impressions from media and movies or their lack of experience with forgery of biometrics. This misconception also emphasizes student failures to examine vulnerabilities as we will discuss further in Theme 4.

The third example involves insecurities of the Internet. Some students have heard that the Internet has vulnerabilities, which is true, but they take this fact and generalize that nothing can be carried out securely on the Internet, which is incorrect. For example, one student stated:

I don't trust any network, anywhere, unless I am the one that controls what is going on. I would never do anything that is not on my own because you never know who is reading it.

While the Internet has many vulnerabilities, the student's response is unreasonably extreme for some applications. For example, cryptography provides tools (e.g., encryption, hashing, digital signatures, virtual private networks) for sending messages over insecure channels with confidentiality, integrity, and authentication.

This response has additional problematic aspects. Even on a private network, one will likely rely on other networks to access information. There is a bias in believing that systems under one's control are necessarily more secure. For many situations, never using a public network would be highly inconvenient. Nevertheless, there is some merit to the student's paranoia; their distrust of all networks has considerable validity.

Theme 2: Conflating Concepts

Conflation is a type of misconception in which a person treats two things that are different as the same. Evidence of conflation includes when a student uses names or words associated with a concept interchangeably. For example, in digital logic contexts, students frequently conflate the concepts of multiplexers and decoders, confusing their topological features (numbers of inputs and outputs) and their functions (Herman, Loui, & Zilles, 2011). Consequently, when asked to describe a multiplexer, some students will describe a decoder instead. Further, these conflations revealed that students could not articulate the purposes or functions that multiplexers and decoders play inside a computer.

In our interviews, students conflated many different concepts such as integrity and availability, hashing and encryption, authorization and authentication, authentication and passwords, risk and threat, and certificate authority and certificate of authenticity. We present three student quotations to illustrate this theme.

In the context of cybersecurity, the term *threat* refers to a potential violation of security, while *risk* refers to the likelihood that an adversary will exploit that threat coupled with the potential resulting loss (Bishop, 2003). By contrast, in colloquial contexts, the idea of risk denotes any dangerous situation—someone may state that they are putting themselves at risk if they drive without a seat belt. This everyday concept of risk excludes the idea of considering the likelihood of a threat occurring where a threat in this driving context might include a drunk driver.

When discussing the security of a drone delivering a package, a student conflated the everyday concept of risk with the cybersecurity concept of risk:

Interviewer: So if you are working for FedEx, how would you design your drones to make sure that doesn't happen?

Subject: ... We need to identify the risks, what all a person can do. [The attacker] can send illegal things, he can send drugs, and he can hack control. He can control to another address. Once we have identified all of the risks, we can create a security protocol for each of those risks.

While the student's reasoning might be correct in everyday conversation, the student's reasoning fails because he is actually describing threats and not risks in formal cybersecurity rhetoric. In our analysis, we similarly identified the terms confidentiality and threat being used in a context true to their definitions outside of the cybersecurity domain, but incorrect within cybersecurity.

A different student also conflated threat and risk, using threat to describe risk:

If the average person has a drone for whatever purpose, they're probably just playing with it for a few moments and it's unlikely that anybody is going to take it from them. If it's something that is going to be carrying anything of value that anyone else might want, that increases the threat.

This student uses the term "threat" instead of the appropriate term "risk" (the increased value of the package increases the likelihood that an adversary would exploit the threat of theft). In these two quotations, students used the word "risk" to describe threats and the word "threat" to describe "risks." This bidirectional confusion of the words to concepts reveals that students likely struggle to create and maintain distinctions between these two concepts. This conflation may, in particular, reveal a lack of adversarial thinking. While threats focus on the characteristics of the system being secured, risk focuses on the motivation and capabilities of adversaries relative to the threats of the system.

Students also conflated passwords and authentication. A password is typically a string of characters used to authenticate a person in conjunction with a username, perhaps as part of an authorization process to gain access to a system. By contrast, authentication is the act of verifying who a person or entity is. We, for example, found that some students conflate passwords with authentication, perhaps because passwords are a common form of authentication. In the following quotation, the student identified scanning fingerprints as a type of "password," instead of a form of authentication:

You could do other forms of passwords, where you don't want to do just standard text, or whatever. You can do something like biometric as a way to take that next step to sync it to that individual, that somebody else

couldn't have access to, such as a fingerprint. It's a fingerprint, a safer measure to make sure that somebody isn't doing it—to tap into your personal information and access it under yourself.

By using the term “password” to refer to the broader concept of authentication, the student reveals a lack of conceptual distinction between passwords and authentication. This type of conflation is slightly different in origin than the conflation of threat and risk. Rather than the conflation resulting from differences between colloquial and technical language, this misconception comes from failing to create an appropriate hierarchy of concepts and specific instantiations of those concepts. This type of misconception is prevalent in other disciplines, such as conflating speed (a scalar) and velocity (a vector whose magnitude is a speed) (Trowbridge & McDermott, 1980). Notably, this student's statement also reflects the overgeneralization of the value of biometrics, as discussed in Theme 1.

Theme 3: Biases

Biases can result in misconceptions when students inappropriately project their assumptions and beliefs on a concept. These beliefs and assumptions are often formed through existing mental frameworks based on life experiences. We identified three biases: user, physical, and personal.

User Bias

Prior to formal instruction in cybersecurity, students primarily experience cybersecurity systems as users—for example, shopping at online stores or paying bills through websites. During interviews, students relied heavily on these prior experiences to orient and guide their thought processes, creating a bias toward thinking about concerns of a user rather than about threats to the data system. For example, for a scenario in which the students imagined that they were the security engineer responsible for an online shopping site, students focused on issues of identity fraud rather than on security of the database. In response to this prompt, one student said:

I think where this question is trying to go, is to see if someone is trying to register for the same account...to get access to their account information.

This student reveals a concern about identity theft (a user concern), rather than SQL injections or denial-of-service attacks to the system. Identity theft is a valid concern, and the security engineer for the shopping site has a responsibility to protect personal information that it collects (e.g., as part of payments). However, the security engineer must also address a wide range of other issues.

A similar example includes students arguing that a system would be secure as long as the latest software and patches were installed and updated, reflecting an emphasis on what the actions a typical user can do to try to protect their system. Additionally, when asked about increasing database security, some students focused on selecting complex passwords, an action that users can take to protect their data but which would not necessarily make data on the server more secure. These examples highlight that students tend to focus the description of security to the user side, likely based on their previous experiences.

Physical Bias

Cybersecurity requires coordination of physical and cyber systems. For example, there can be physical locks on a door to a room that holds a computer with digitally encrypted files. Both the physical locks and digital encryption are part of the security of the system and both can be potentially compromised. Students, however, demonstrated biases toward the physical elements of the systems, discussing them more readily, believing that they were more secure, and inappropriately analogizing physical systems to cyber systems. Since most of life is experience within a physical domain, it is not surprising that students build their conceptions of security based on these physical systems.

For example, when asked about making a computer system in a hotel business center more secure, one student focused on the characteristics of the room and did not discuss cyber threats:

Interviewer: What if the hotel has gotten a complaint because something insecure happened in their business office. They hired you and said, come on in and make our business office more secure for our guests. What are some things that you might be able to offer to them to make those transactions more secure for their guests?

Subject: That is a very interesting question. I am saying that I am trying to picture the locks on most of the business rooms in the hotels. Some of them are very small, they are just little tiny small space there with two or three computers. So, when they sit there, if all of those computers are taken, those two or three, then everybody—somebody next to you can just see exactly what you are doing. So, perhaps one thing that I recommend to them is to-If that is the only space that they have, then maybe only have a single computer in that small space.

Having a room may provide some security from a person walking behind and overseeing transactions; however, a key logger, camera, or malware could be on the machine, and the physical system might not address those issues. When asked

to consider interventions beyond physical ones to increase security, the student provided only vague descriptions of “Internet security:”

Subject: Okay, then you just wonder how they are going to secure their Internet, so that they make sure the security of their Internet is trusted. They give you this Internet facility and you use it, if your data is messed up, it is going to come back to them. So, make sure their Internet security is correct.

The student recognizes that security and interventions can go beyond the physical systems, and he knows that the online actions need to be “secure” and “trusted.” However, this student did not provide evidence of understanding of how to create security or trust within the digital domain.

Students demonstrated further biases toward the security of physical systems, asserting that human beings cannot be hacked, overestimating the strength of having a physical key on a system, or believing that the physical presence of a security engineer would be needed to remove viruses from a network. Additionally, students sometimes inappropriately compared the security of physical systems to make sense of the security of cyber systems. For example, one student analogized that (correctly-addressed) email messages might be misdelivered as postal mail is sometimes mistakenly delivered to a nearby building. Making analogies with physical systems (e.g., sending messages in locked boxes) can help students understand cybersecurity concepts, but analogies can be imperfect and care must be taken not to assume that a physical or cyber system is necessarily more secure than the other type.

Personal Bias

Beliefs or stereotypes about countries and cultures was an additional source of bias in student reasoning about cybersecurity. For example, one student’s beliefs about the security of file transfer depended on which country he was in:

Let’s say I am sending a file in America itself. Then there is a good chance my file won’t be hacked. But let’s say I am overseas—I am going to countries where hacking is very common. Let’s say Syria or China, or somewhere like that.

This response is incorrect for two reasons: hackers reside in the US as well as in other countries, and theft of data can be accomplished remotely (revealing an additional physical bias in the student’s reasoning). This student’s belief about countries and security, revealing implicit trust in certain individuals without verification, is likely based in personal beliefs and stereotypes. Further, this student’s unverified trust in individuals and systems betrays a lack of adversarial thinking. Nevertheless, there may be some truth to the belief that the risk of cyber attack to visitors is greater in some countries than others. While no other students

explicitly revealed similar beliefs during the interviews, this type of bias is important because it may support ways in which students artificially limit their adversaries or overlook vulnerabilities—our next theme.

Theme 4: Incorrect Assumptions - Limiting the Adversary and Failure to See Vulnerabilities

The incorrect assumptions theme focuses on the process by which students make assumptions about systems, technologies, or adversaries without questioning the basis or validity of those assumptions. These assumptions betray a lack of adversarial thinking as students do not adequately consider the range of possible motives or capabilities of their adversaries before proposing and justifying security measures. These incorrect assumptions and a failure to question them, led to a range of misconceptions and problems in student reasoning. Some examples that fall into this category include: students relying heavily on policies or a lack of knowledge rather than physical systems to deter an adversary, to assume some groups of people such as military personnel are all trustworthy, and that a system which is under attack by an adversary will exhibit detectable and anomalous behavior. We provide three examples of this process to illustrate the dangers of incorrect assumptions.

The first example focuses on the security of a military voting system. A student identified one possible motive for an adversary, and failed to consider other motives. The proposed solution had significant vulnerabilities, including creating a new mechanism for an attack on the voting system:

There should be a feature for when you enter a password, more than a certain number, a wrong one, it scrambles the data and makes it unusable and destroys [the data]. Makes the data unusable so that, in the worst-case scenario, the bad guy wouldn't get any useful information out of it.

A significant vulnerability of this proposed system is that an attacker may not focus on theft of the information, but rather on destruction of data. In the proposed solution, the attacker can intentionally choose to enter in false passwords to achieve their goal. This student consequently proposed a solution that may aid an attacker, because he assumed the attacker's goal must be theft.

In the second example, a student incorrectly assumed malware cannot be stored on keyboards or mice. Malware is a type of software intended for some nefarious end, such as to damage or disable a computer system, or to steal information. It can impact a system instantaneously and can be loaded onto a computer through a variety of means, including through keyboards and mice. When asked how to mitigate malware from a USB port the student stated:

You could have an alarm or disable USB ports on the machine, but the problem with that is that most machines nowadays only have USB ports for the keyboard [and] mouse. Unlike older machines that have the older PS2 connectors for the keyboard and mouse. You would have to be careful about how you are going to enable USB for the keyboard or mouse, but throw an alarm if it something else. It's difficult. It's not an easy answer.

This answer is problematic because the student incorrectly assumed that a keyboard or mouse could not be used to install malware. Further, the student failed to consider that malware can affect a network faster than a person can respond to a physical alarm (another example of physical bias). This suggests that the student does not understand how malware can be installed onto a system, and the speed at which it can infect the system. This misconception resulted in incorrect reasoning on a proposed mitigation of a malware attack.

In the third example, a student role-plays a penetration tester attempting to exfiltrate a USB stick hidden under a floor tile in a top-secret government facility. In response to this challenge, the student suggests a social engineering solution. This response would be correct by itself; however, the student limited the adversary and made the claim that it is impossible to break into a top-secret governmental building, without providing adequate evidence or defending the claim. As a result, the student's solution is incorrect because the student situates the solution as the *only* possible solution.

Subject: ... it is a top-secret government facility so there is really no way to break in. I can only see paying somebody off or getting somebody to get it for you. Really, that is the only way to get something that is that secretive, to me. I really don't see another alternative of getting in there.

Interviewer: Okay, well, think about it for a second. Think outside of the box.

Subject: I have a question. Does Dave just leave the USB stick under there at all times? Does he leave it there when he leaves work, plugs it back in and uses it whenever he comes to work?

Interviewer: We don't know that.

Subject: We don't know that? Okay. I don't know any of these things. First things first. I'm outside of the fence. Am I able to see inside the building? That's a big thing. It has no window, right? It is government. Okay, that is eliminated. The architecture of the building—so, we know it has no windows. I can get some shady Russians to make me a fake identity. Get a government clearance with it. Work at the same building he does. Actually, the only thing I can think of, honestly, is to literally pay somebody off, or pay Dave off, himself, to bring the USB stick out to me. Oh, I am a penetration tester. I am

not a hacker, or anything. I honestly don't know how to get the USB drive. I am really stumped on that. I really don't know. I would have an inside man.

Interviewer: Any other thoughts? It's hard to do, right, because it is a secret government facility.

Subject: Yeah, I don't know how I will be able to do that. I can't walk in there. It's not like it is on the system where I could see and attack.

Even though the proposed social-engineering attack would be plausible, the adversarial mindset requires a more open-minded approach. The student shows that they are stuck, and is trying to think through ideas, but each time they propose a solution, the student immediately rejects it as an option. A determined adversary might have substantial resources and patience. The student's response is flawed because the student is constrained by the belief that there is *only* one solution.

DISCUSSION

The four themes of student misconceptions we identified reveal that many students did not understand the functional and technical components of cybersecurity, such as encryption, and they did not adequately consider the complexity of the concepts and scenarios. Adversarial thinking requires a reflective mindset that carefully considers not only one's own actions but also the motivations, capabilities, and opportunities of adversaries who desire to harm the systems defenders seek to protect.

The first theme is *overgeneralizing* the usefulness or applicability of one concept into inappropriate contexts. Overgeneralizations reveal that students believe that because a solution or idea was useful in one context it must have some general utility.

In the second theme, students *conflated* two concepts into one concept. Conflations reveal that students do not distinguish between similar but distinct concepts, simplifying their understanding of cybersecurity, again revealing a lack of nuance in understanding the nature of problems or the limitations of solutions. While this type of shallow reasoning can be detrimental in any context, it is particularly disconcerting in cybersecurity. The adversarial mindset needed for cybersecurity is inherently complex, requiring consideration of multiple actors and systems and each of their capabilities, weaknesses, and motives, and considering the interactions among these myriad factors.

The third theme reveals that *biases* guided and directed student engagement with the scenarios. For example, students frequently considered only how they

would act during a given situation or inappropriately analogized with familiar physical systems.

In the fourth theme, students applied *incorrect assumptions*, often inappropriately limiting or underestimating the adversary. Students did not adequately explore and question the motives and capabilities of the actors, sometimes projecting overconfidence in their answers. Students approached scenarios without a reflective stance, instead seeking to identify solutions immediately without fully considering the vulnerabilities of the systems or potential threats.

These types of biases are normative and to be expected given that students will construct their knowledge based on prior experiences as described by our framing in constructivist theory. Therefore, with limited experience in a discipline, students will often rely on attempting to analogize to prior experiences or to rely on patterns of reasoning that worked in prior contexts. Notably, shallow approaches to problem solving are common among novices across many disciplines. Novices frequently employ means-end analysis while solving problems, attempting to find the shortest path from a problem statement to a desired solution. For example, if a student is asked to find the position of a particle given its velocity, the student might pull up an equation sheet looking for equations that contain variables for position and velocity rather than identify the concepts (e.g., conservation of momentum) that inform the solution process.

Alignment of Findings with Prior Work

Our findings provide evidence of fragility or context-dependence in student knowledge/beliefs. For example, in the context of the online shopping site, a student explicitly stated they did not want to trust the end user in the design. They used this concept as the basis of describing a design that would prevent an SQL injection attack, which was an appropriate design for the scenario. In contrast, in the next scenario, the same student implicitly trusted the user while describing how to design a connected private and public network by having the system depend on the user's compliance to the policy. When the interviewer addressed the notion of not trusting the user, the student defended the proposed design, stating that it will deter the user because the user could face criminal charges. While the student had some understanding that a security engineer should design a system that prevents malicious attacks, even from the user, they applied this knowledge only in specific contexts.

Robustness of Findings Across Institutions

We observed some differences in how students from the various institutions responded. For example, students at community college tended to propose solving cybersecurity problems by applying specific corporate products rather than describing generic solutions. Students pursuing M.S. degrees tended to project more self-confidence in their answers than did their undergraduate counterparts, regardless of the quality of their responses. We did not observe differences in misconceptions across the three schools.

As we had expected, analyzing 25 interview transcripts was a significant amount of work.

Overall, our results suggest that students do not adequately comprehend core concepts in cybersecurity. The misconceptions themes (over generalizations, conflated concepts, biases, and inappropriate assumptions or solutions) suggest that students use a form of “*satisficing*” in their reasoning (Brown, 2004), becoming too easily satisfied that a system is secure after identifying only one possible source of security rather than seeking to explore the adversarial space more thoroughly.

Implications, Open Problems, and Future Research

The misconceptions we uncovered can be viewed as aspects of two central problems: they lack a useful framework around which to organize their thoughts and they do not comprehend the complexity of cybersecurity challenges. Based on these findings, we offer tentative suggestions for how these findings may be applied to the classroom and suggest future research that could validate or challenge these suggestions.

To address the first challenge, we posit that adversarial thinking may be a useful framework to help students organize their knowledge. If true, all cybersecurity learning activities should require students to think deeply about their adversary as well as about potential ways to counter that adversary.

To address the second challenge, we suggest using teaching methods that are designed to engage students in complex scenarios. Model Eliciting Activities (MEAs) and case studies are two such teaching methods documented in the research literature. We draw upon the research literature further to suggest why MEAs and case studies may be helpful and how to employ them in the classroom effectively.

Summary of MEAs and Case Studies

We recommend that instructors embrace complexity and seek to engage students in scenarios that require adversarial reflections and careful distinctions between core concepts. Beginning instruction with complexity is counter-intuitive to many instructors as science instruction frequently employs simplifying assumptions (e.g., frictionless environments in physics or teaching mathematical equations before introducing word problems). However, teaching techniques that embrace complexity can help students learn fundamental concepts and better understand the limits of when concepts apply.

In MEAs, instructors give students real world datasets that include information that may not lead to simple perfect mathematical solutions (Zawojewski, Diefes-Dux, & Bowman, 2008). Students must determine how to extract information from these datasets to create mathematical models (e.g., compute averages, perform regressions to model the data) that can inform decisions based on the data or aid in interpretation (Bostic, 2013). Because the data are complicated, students must make reflective decisions about the appropriateness of different models. For example, students need to decide which measure of central tendency (e.g., mean, median, mode) is most appropriate, requiring students to develop conceptual distinctions between these measures (Bostic, 2013).

A similar approach may help students develop conceptual distinctions between concepts such as confidentiality and integrity. The emphasis in grading should be not on the numerical correctness of answers but on the reasoning behind the chosen model. Similarly, cybersecurity instruction may benefit from engaging students in complex scenarios that do not have clear “best” solutions. Students would need to engage in creating models of their adversaries and their systems that justify those solutions.

Case studies, like MEAs, embrace complexity by designing instruction around real-world scenarios that often lack a single clear right answer (Davis & Yadav, 2014). Cases are constructed by carefully documenting facts about instructive real-world examples (Davis & Yadav, 2014) such as security breaches or documentation of existing security plans. The instructor provides supporting lecture content that students need to understand the case, but students are responsible for carrying out their own tasks (Henderson, Bellman, & Furman, 1983). These tasks could include explaining failures or successes of security plans in cases, devising attacks against a plan, or devising defenses against a known adversary. Cases are often viewed as the best way to help students understand the complexity of topics in business and law, but are relatively unexplored in engineering and computer science topics (Davis & Yadav, 2014). These types of

approaches generally improve students' ability to monitor their own learning and may help students learn concepts better too (Gallucci, 2006; Yadav, Shaver, & Meckl, 2010).

Suggestions for Designing Complex, Yet Tractable Instruction

It is important for students to comprehend the complexity of cybersecurity issues, yet immersing students in unrestricted cybersecurity scenarios can be excessively challenging. To address this conflict, we suggest that it can be useful to find ways to instantiate core concepts in more limited contexts, to ground thinking, while still enabling students to experience the full range of concepts. An example of this approach is to use the slightly more limited context of cryptography and communications security for introducing core concepts, including the CIA triad and authentication, while maintaining a strong emphasis on adversarial thinking. From such an anchor, students can then build off complexity for more general cybersecurity challenges.

Understanding one core concept can help students learn other related concepts more accurately and faster than if they had misconceptions about that one core concept. For example, teaching students about the idea of emergent processes (discernable macro-level phenomena that occur through the uncoordinated actions of many independent individuals) through an example such as ants marching in a line can help those same students learn about other emergent processes such as diffusion (Slotta & Chi, 2006).

The example of using cryptography for this approach is attractive because of the pervasive misconceptions we observed about cryptography and the pivotal nature of its core concepts. In our study, 17 of the 25 students exhibited misconceptions about cryptography, while most of the other students did not describe cryptography in sufficient detail in their interviews. We categorized most of these misconceptions as overgeneralizations.

When using the example approach of cryptography, care must be taken to promote awareness that cybersecurity is more than cryptography. For example, there are many potential ways to defeat a cryptographic communications system without breaking any cryptographic operation. Malware on a host computer might leak private keys and messages without breaking any cipher. Discussing such potential vulnerabilities provides important learning moments to help students appreciate the complexity of cybersecurity.

Open Problems

Our work motivates additional research into understanding cognitive patterns students develop while learning cybersecurity. For example, it would be

interesting and useful to explore cybersecurity, cryptography, and student fragility in thinking. Also, it would be interesting to study the causes of the observed misconceptions.

It remains an open problem how educators can apply our findings to develop more effective and efficient strategies for teaching and learning cybersecurity. It might be interesting to explore what computer textual analysis of our interview transcripts may reveal.

Future Research

We are using our findings to develop two educational cybersecurity assessment tools: the Cybersecurity Concept Inventory (CCI), and the Cybersecurity Curriculum Assessment (CCA). The CCI measures conceptual understanding of students completing a first course in cybersecurity, and the CCA assesses how well a college cybersecurity curriculum prepares graduates for a career in cybersecurity. In particular, we base incorrect choices in multiple-choice assessment items in the CCI on misconceptions uncovered in this study.

Our next steps are to validate the draft CCI with expert review, cognitive interviews, and psychometric testing. We also plan to apply the CCI to identify and assess effective ways to teach and learn cybersecurity.

CONCLUSION

Using novice-led paired thematic analysis, we analyzed 25 student interviews of cybersecurity students to document student misconceptions and problematic reasoning. Four themes emerged: over generalizations, conflated concepts, biases, and incorrect assumptions. Furthermore, while reasoning about the cybersecurity scenarios presented to them, most students failed to comprehend the complexity and subtlety of possible vulnerabilities, threats, risks, and mitigations. Our work is the first formal study of student misconceptions about cybersecurity.

We carried out this study as part of our Cybersecurity Assessment Tools (CATS) project, which is developing educational assessments including a Cybersecurity Concept Inventory (CCI) to assess student understanding of core concepts in cybersecurity. We plan to apply the CCI to identify and compare effective ways to teach and learn cybersecurity. Our study of student misconceptions can help shape cybersecurity content, assessment, and pedagogy to improve student learning.

Our work will be useful in developing curricula, educational techniques and materials for cybersecurity. Knowledge of student misconceptions can lead to strategies addressing and possibly avoiding them. While it was not the purpose of

this study to suggest specific teaching strategies, our findings support the conclusion that students should learn cybersecurity in a way that encourages them to recognize and appreciate its considerable complexity. Case study is one approach that can achieve this goal. Our companion paper (Sherman et al., 2017b), which answers six of the interview prompts, provides material that can be used in case studies and other learning activities. We encourage educators to consider the misconceptions identified in our study to develop and refine more effective approaches to teaching cybersecurity.

References

- ACM/IEEE/AIS/IFIP Joint Task Force. (2017). Cybersecurity curricula 2017: Curriculum guidelines for Post-Secondary degree programs in cybersecurity. New York City. Available at: <http://www.csec2017.org>
- Bishop, M. A. (2003). The art and science of Computer Security. Boston, MA: Addison-Wesley Longman Publishing Co. Inc.
- Bishop, M., Fitcher, L., Miloslavskaya, N. & Theocharidou, M. (2017), Eds. Information Security Education for a global digital society, Proceedings of WISE 10, Rome, Italy, May 29-31, 2017, Springer.
- Bodeau, D., Fabius-Greene, J. & Graubart, R. (August 2010). How do you assess your organization's cyber threat level? The MITRE Corporation.
- Borrego, M., Douglas, E. P., & Amelink, C.T. (2009). Quantitative, qualitative, and mixed research methods in engineering education. *Journal of Engineering Education*, 98(1), 53–66.
- Bostic, J. D. (2013). Model-eliciting activities for teaching mathematics, *Mathematics Teaching in the Middle School*, 18(5), 262-266.
- Bransford, J. D., Brown, A. L., Cocking, R. R. (2000). How people learn: Brain, mind, experience, and school. Washington, D.C.: National Academies Press.
- Braun, V. & Clarke V. (2006). Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3(2), 77–101.
- Brown, R. (2004). Consideration of the origin of Herbert Simon's theory of 'Satisficing' (1933-1947). *Management Decision*, 42(10), 1240–1256.
- Burley, D., Buck, S., Dark, M., Fitzgerald, S., Hawthorne, E., Kono, T., & Portz, S. (2014) Cybersecurity education workforce final report. Arlington, VA: George Washington University Arlington Center.
- Caltagirone, S., Pendergast, A. & Betz, C. (May 7, 2013). The Diamond Model of intrusion analysis, Department of Defense.
- Chi, M.T.H., Feltovich, P.J. & Glaser R. (1981). Categorization and representation of physics problems by experts and novices. *Cognitive Science*, 5, 121–152.
- Chi, M. T. H. (2005). Commonsense conceptions of emergent processes: Why some misconceptions are robust. *The Journal of the Learning Sciences*, 14(2), 161–199.
- CISSP (2016). CISSP. Retrieved from <https://www.isc2.org/CISSP/Default.aspx>

- Davis, C. & Yadav, A. (2014). Case studies in Engineering. In A. Johri & B. M. Olds (Eds.), *Cambridge Handbook on Engineering Education* (pp. 161-173). Cambridge: Cambridge University Press.
- diSessa, A. A., Gillespie, N. M. & Esterly, J. B. (2004). Coherence versus fragmentation in the development of the concept of force. *Cognitive Science*, 28: 843–900.
doi:10.1207/s15516709cog2806_1
- Duggan, D. P., Thomas, S.R., Veitch, C. K. & Woodard, L. (September 2007). Categorizing threat building and using a generic threat matrix, SAND2007-5791, Sandia National Laboratories.
- Ericsson, K. A. & Simon, H. A. (1993). *Protocol Analysis*. Cambridge, MA: MIT Press.
- Frost & Sullivan (2015). The 2015 (ISC)² global information security workforce study:
[https://www.isc2cares.org/uploadedFiles/wwwisc2caresorg/Content/GISWS/FrostSullivan-\(ISC\)%C2%B2-Global-Information-Security-Workforce-Study-2015.pdf](https://www.isc2cares.org/uploadedFiles/wwwisc2caresorg/Content/GISWS/FrostSullivan-(ISC)%C2%B2-Global-Information-Security-Workforce-Study-2015.pdf)
- Gallucci, K. (2006). Learning concepts with cases. *Journal of College Science Teaching*, 36(2), 16-20.
- Guba, E. G. & Lincoln, Y.S. (1994). Competing paradigms in qualitative research. *Handbook of Qualitative Research*, 2 (163–194), 105. Thousand Oaks, CA: Sage.
- Hake, R. (1998). Interactive-engagement vs. traditional methods: A six-thousand-student survey of mechanics test data for Introductory Physics. Arlington, VA: National Science Foundation.
- Henderson, J. M., Bellman, L. E. & Furman, B. J. (1983). A case for teaching engineering with cases. *Journal of Engineering Education*, 288-292.
- Herman, G. L., Loui, M.C. & Zilles, C. (2011). Students' misconceptions about medium-scale integrated circuits. *IEEE Transactions on Education*, 54 (4), 637–645
- Herman, G. L., Zilles, C. & Loui, M.C. (2012). Flip-flops in students' conceptions of state. *IEEE Transactions on Education*, 55(1), 88–98. doi:10.1109/TE.2011.2140372
- Hestenes, D., Wells, M., & Swackhamer, G. 1992. "Force Concept Inventory." *The Physics Teacher*, 30, 141–166.
- The Joint Task Force on Computing Curricula (2013). *Computer Science curricula 2013: Curriculum guidelines for undergraduate degree programs in Computer Science*. Retrieved from <https://www.acm.org/education/CS2013-final-report.pdf>
- Kim, D. & Solomon, M. G. (2014). *Fundamentals of Information Systems security*, second edition. Burlington, MA: Jones & Bartlett Learning.
- Libicki, M. C., Senty, D., & Pollak, J. (2014). Hackers wanted: An examination of the cybersecurity labor market. Retrieved from Santa Monica, CA:
http://www.rand.org/pubs/research_reports/RR430.html
- Mateski, M., Trevino, C.M., Veitch, C. K., Michalski, J., Harris, J. M., Maruoka, S. & Frye, J. (2012). Cyber threat metrics. Sandia National Laboratories.
- Montfort, D. B., Herman, G.L., Brown, S.A., Matusovich, H.M. & Streveler, R.A. (2013). Novice-led paired thematic analysis: A method for conceptual change in engineering. *Paper presented in the American Society of Engineering Education*, Atlanta, GA.
- Nathan, M. J., Alibali, M. W., & Koedinger, K.R. (2005). Expert blind spot: When content knowledge & pedagogical content knowledge collide. Institute of Cognitive Science: University of Colorado, Boulder.
- Nathan, M. J. & Petrosino, A. (2003). Expert blind spot among preservice teachers. *American Educational Research Journal*. 40(4): p. 905–928.
- NIST. (2013). *The National Cybersecurity Workforce Framework*. Retrieved from Washington, DC: <http://csrc.nist.gov/nice/framework/>
- Özdemir, G., & Clark, D.B. (2007). An overview of conceptual change theories. *Eurasia Journal of Mathematics, Science & Technology Education*, 3(4), 351-361.

- Parekh, G., Scheponik, T., DeLatte, D., Herman, G. L., Oliva, L., Phatak, D. & Sherman, A. T. (2018). Identifying core concepts of cybersecurity: Results of two Delphi processes. *IEEE Transactions on Education* vol. 61, number 1, pp 11-20.
- Patton, M. Q. (2005). *Qualitative research*. NY: John Wiley & Sons.
- Perkins, D. N. & Martin, F. (1986). *Fragile knowledge and neglected strategies in novice programmers*. Paper presented at the First Workshop on Empirical Studies of Programmers.
- Sackur-Grisvard, C., & Leonard, F. (1985). Intermediate cognitive organization in the process of learning a mathematical concept: The order of positive decimal numbers. *Cognition and Instruction*, 2, 157–174.
- Scheponik, T., Sherman, A. T., DeLatte, D., Phatak, D., Oliva, L., Thompson, J. & Herman, G. L. (October 2016). How students reason about Cybersecurity concepts. In *Proceedings of the Frontiers in Education Conference (FIE)*.
- Schneier, B. (1996). *Applied Cryptography: Protocols, Algorithms, and Source Code in C*, second edition. Hoboken, NJ: Wiley & Sons.
- Sherman, A. T., Oliva, L., DeLatte, D., Golaszewski, E., Neary, M. Phatak, D. Scheponik, T., Herman, G. & Thompson, J. (2017a). Creating a cybersecurity concept inventory: A status report on the CATS project, *National Cyber Summit*, June 6-8, 2017, Huntsville, AL: <http://arxiv.org/abs/1706.05092>
- Sherman, A. T., DeLatte, D., Herman, G. L., Neary, M., Oliva, L., Phatak, D., Scheponik T., & Thompson, J. (2017b). Cybersecurity: Exploring core concepts through six scenarios, *Cryptologia*, Sept 2017 (online).
- Simon, H. A. (1996). *The Sciences of the artificial*. Cambridge, MA: MIT Press.
- Slotta, J. D. & Chi M. T. (2006). Helping students understand challenging topics in science through ontology training. *Cognition and Instruction*, 24(2), p. 261–289.
- Steinle, V. (2004). Changes with age in students' misconception of decimal numbers. Melbourne, Australia: Department of Science and Mathematics Education, University of Melbourne. Retrieved from <http://eprints.unimelb.edu.au/archive/00001531/>
- Trowbridge, D.E. & McDermott, L.C. (1980). Investigation of student understanding of the concept of velocity in one dimension, *Am. J. Phys.* 48 (12) 1020.
- Vosniadou, S. (2007). The conceptual change approach and its reframing. In S. Vosniadou, A. Baltas, & Vamvakoussi, Z. (Eds.), *Reframing the Conceptual Change Approach in Learning and Instruction* (pp. 1–15). Amsterdam: Elsevier.
- Yadav, A., Shaver G. M., & Meckl, P. (2010). Lessons learned: Implementing the case teaching method in a mechanical engineering cours. *Journal of Engineering Education*, 99(1), 55-69.
- Zawojewski, J. S., Diefes-Dux, H. A. & Bowman, K. J. (2008) (Eds). Models and modeling in Engineering Education. Rotterdam, The Netherlands: Sense.

Appendix A: Twelve Interview Scenarios

This appendix gives our twelve interview prompts. Table 2 shows how we organized these prompts into three protocols, each comprising four scenarios. We hope that educators will find these scenarios useful for a variety of learning activities.

Scenario	Alpha	Bravo	Charlie
1	package delivery by drones	client-database design	lost luggage
2	file transfer	online shopping	precinct voting
3	database input error	protecting trade secrets	two-channel authentication
4	private network design	Nuclear Test Ban Treaty	exfiltrating a USB stick

Table 2. The 12 interview scenarios organized into three protocols Alpha, Bravo, and Charlie.

PROTOCOL ALPHA

Scenario 1: Package Delivery by Drones

Consider how a company might deliver packages by drones. As a security engineer for the company, what threats can you identify?

Scenario 2: File Transfer

Alice and Bob want to share a file over the Internet. What are some of the cybersecurity issues that may arise? Sketch a diagram to help explain your answer.

Scenario 3: Database Input Error

When a user Mike O'Brien registered a new account for an online shopping site, he was required to provide his username, address, first and last name. Immediately after Mike submitted his request, you—as the security engineer—receive a database input error message in the logs. What might you infer from this error message?

Scenario 4: Private Network Design

Bob wants you to design a secure network that allows him to have a segment on the public Internet and a private segment with no public access. Traffic must also be able to be taken from the public segment to the private segment, but no data must ever go from private to public. As the security architect sketch a design that meets Bob's requirements.

Protocol Bravo

Scenario 1: Client-Database Design

Sketch a diagram that illustrates a secure communication between a customer and a database.

Scenario 2: Online Shopping

While Mary is traveling, she decides to do some shopping online. She is connecting from a computer in a hotel business office. What are some of the cyber security issues that might arise? Sketch a figure to illustrate your explanation.

Scenario 3: Protecting Trade Secrets

There is a server that holds a company's trade secrets. As the chief security officer devise a comprehensive security strategy to protect these trade secrets from corporate enemies.

Scenario 4: Nuclear Test Ban Treaty

To comply with the terms of the Nuclear Test Ban Treaty, Country *A* would like to implant a seismic sensor under Country *B*'s soil to monitor underground weapons testing. Country *A* fears that *B* will try to falsify the signals of the sensor, and Country *B* fears that *A* will try to exfiltrate spy information embedded in the seismic data. Neither party trusts the other. Requirements of the system include each of the following:

1. Country *A* wants assurance that the seismic data it receives came from its sensor and were not modified.
2. Country *B* wants to be able to monitor the signals transmitted from the sensor in real time. It too wants assurance that the signals were not modified.
3. The design should be fair to both parties.

How would you design a system that complies with these requirements? Draw a sketch to illustrate your design.

Protocol Charlie

Scenario 1: Lost Luggage

Bob's manager Alice is traveling abroad to give a sales presentation. Bob receives an email with the following message: "Bob, I just arrived and the airline lost my luggage. Would you please send me the technical specifications for our new product? Thanks, Alice." What should Bob do?

Scenario 2: Precinct Voting

A security company is designing a precinct voting system for military voters overseas. The system must provide voter authentication, ballot confidentiality, integrity of marked ballots during transportation, and assured operations. How would you design such a system?

Scenario 3: Two-Channel Authentication

Alice is logging onto a server. She sends her username and password over the Internet. The server instructs the security computer to send a challenge to Alice's cell phone. For example, the challenge is the name of Alice's pet, and the response is 'Skippy'. If the security computer deems the response is valid, then the security computer signals the server to accept Alice's login. Comment on the security of this system.

Scenario 4: Exfiltrating A USB Stick

Alice works in a top-secret government facility where she has hidden a USB memory stick, with critical information, under a floor tile in her workspace. Starting from outside the fence of the building, how would you, as a penetration tester, retrieve the USB stick?