Editors: Jelena Mirkovic, mirkovic@isi.edu | Bill Newhouse, newhouse@nist.gov

The CATS Hackathon: Creating and Refining Test Items for Cybersecurity Concept Inventories

Alan T. Sherman, Linda Oliva, Enis Golaszewski, Dhananjay Phatak, and Travis Scheponik | University of Maryland, Baltimore County
Geoffrey L. Herman, Dong San Choi, and Spencer E. Offenberger | University of Illinois at Urbana–Champaign
Peter Peterson | University of Minnesota Duluth
Josiah Dykstra | U.S. Department of Defense
Gregory V. Bard | University of Wisconsin, Stout
Ankur Chattopadhyay | University of Wisconsin, Green Bay
Filipo Sharevski | DePaul University
Rakesh Verma | University of Houston
Ryan Vrecenar | Texas A&M University

Presently, there is no rigorous, research-based method for measuring the quality of cybersecurity instruction. Validated assessment tools are needed so that cybersecurity educators have trusted methods for discerning whether efforts to improve student preparation are successful. The Cybersecurity Assessment Tools (CATS) Project[9] provides rigorous evidence-based instruments for assessing and evaluating educational practices (http://www.cisa.umbc.edu/cats/index.html). The first instrument is the Cybersecurity Concept Inventory (CCI), which measures how well students understand core concepts in cybersecurity (especially adversarial thinking) after a first course in the field. The second instrument, the Cybersecurity Curriculum Assessment (CCA), measures how well students understand the same core concepts after completing a full cybersecurity curriculum and whether they are ready to enter the workforce as cybersecurity professionals. These tools can identify pedagogies and content that are effective in teaching cybersecurity.

In February 2018, we hosted a two-day CATS Hackathon for 17 cybersecurity educators and professionals from across the nation to generate multiple-choice test items for the CCA, and to refine draft items for the CCI and CCA. The meeting was a "hackathon" in that participants collaborated on a common task in an informal setting.[7] Over the past several years, we developed a bank of approximately 36 questions and 12 draft questions for the CCI and CCA, respectively. Participants used these questions as a starting point, extending CCI questions to be CCA questions, refining draft CCA questions, and devising new CCA questions entirely. The intimate in-person event facilitated

Digital Object Identifier 10.1109/MSEC.2019.2929812
Date of current version: 29 October 2019

productive interactions among the participants, infusing fresh ideas into the project, promoting awareness of the tools, and enhancing the quality of the test items.

The CCI and CCA instruments described in this article can enable educators to assess how well their courses and approaches are helping students develop a strong, conceptual understanding of cybersecurity. Educators who would like to participate in our validation studies or use a preliminary version of these instruments are invited to contact us.

The CATS Project

Inspired by the Force Concept Inventory (FCI) of physics by Hestenes et al.,[4] we designed the CCI and CCA to be rigorous assessment tools relevant to a wide range of educational contexts. These assessment tools can provide a broadly accepted research instrument statistically analyzed by established methods, which can be used to determine how effective certain teaching practices are at helping students learn cybersecurity.[2,3,5]

Unlike the Certified Information Systems Security Professional information security certification, which is largely informational, our instruments assess conceptual understanding. We measured conceptual understanding because it is a critical transferable skill that accelerates future learning. Like the FCI, our new tests focus only on core concepts to maximize applicability to a variety of curricula and, thus, are intentionally not comprehensive. They do not measure general problem-solving, design, analytical, or interpersonal skills; rather, they are intended to compare teaching methods not individual instructors or individual students.

Each 50-min test comprises approximately 25 multiple-choice questions (MCQs), with five on each of the following five core cybersecurity concepts identified through our Delphi process.[6] Our Delphi process solicited the following input from a set of subject matter experts to create a consensus about contentious decisions, sharing comments without attributions:

1. Identify vulnerabilities and failures.
2. Identify attacks against the confidentiality, integrity, and availability triad and authentication.
3. Devise a defense.
4. Identify the security goals.
5. Identify potential targets and attackers.

Each test item has three parts: a scenario, a stem (question/prompt), and five answer choices (alternatives). Several items may share the same scenario, but each item has a unique stem and answer choices. Each stem focuses on one targeted concept, though scenarios may deal with multiple concepts. In addition, each stem has exactly one correct (best) choice and four distractors (incorrect answer choices). Test items should target these aforementioned timeless, fundamental concepts, not merely factual information that is memorized and recalled.

It is our intent that, for each core concept, the five test items encompass a range of difficulty levels. We recognize, however, that experts tend to be poor judges of the difficulty of test items, so the actual difficulty of each item will not be reliably known until after student testing.

The CATS team developed draft test items using the following structured process to create scenarios, stems, and distractors. Building on the five core concepts identified in our Delphi process, we created scenarios and interview prompts, which we used to interview students to uncover their misconceptions.[10] It took significant planning, staff time, and effort to carry out, record, transcribe, and analyze these think-aloud interviews. Subsequently, in discussions held in a conference room or on Skype, we devised stems and answer choices. We based distractors mostly on the student misconceptions we uncovered during the interviews. Scenarios we developed for these interviews provide rich case studies for many learning activities.[8] To test draft questions, we used the PrairieLearn System developed at the University of Illinois at Urbana–Champaign (https://prairielearn.readthedocs.io/en/getting-started-docs/).

There is evidence that well-crafted MCQs can provide the same type of information as do Parson's problems (i.e., open-ended problems). MCQs are easy to grade and interpret, and there is a robust theory for creating and analyzing them. Seventy-six percent of our 36 CCI Delphi experts agreed or strongly agreed that "A carefully constructed multiple-choice assessment can provide valuable information for assessing the quality of instruction in a first course in cybersecurity." Other types of assessments (e.g., simulations, hands-on activities, and competitions) also have much to offer but are more complex to create, maintain, administer, and analyze.

It is essential that these tools have strong usability and validity and are widely implemented in diverse settings. Throughout the project, we benefited from inputs from a wide variety of experts, beginning with our Delphi experts.[6] We planned the hackathon to encourage and facilitate experts to collaborate on refining existing test items for the CCI and CCA and to develop additional test items for the CCA. The project will continue forward with expert reviews and pilot testing of draft test items.

The Hackathon

To generate new test items for the CCA, 17 participants were organized into several teams, each with three or four members. Each team

focused on one of the following tasks: 1) generating new scenarios and question stems; 2) extending CCI items into CCA items, and generating new answer choices for new scenarios and stems; and 3) reviewing and refining draft CCA test items. These substantial tasks kept each team fully engaged throughout the two-day hackathon. Each participant chose what team to join, based in part on their skill sets.

The event took place at an off-campus conference center, two days before the ACM Special Interest Group on Computer Science Education conference in Baltimore, Maryland. Thirteen experts were from universities and two each were from industry and government, respectively. Participants took the CCI at the beginning of the first day and the CCA at the beginning of the second day. In the following sections, we describe each task in more detail.

Task 1: Generate New Scenarios and Question Stems

These teams began by brainstorming potential scenarios. Team members shared their scenarios and developed a priority list of ones that needed further development. Members then refined each scenario by adding details, identifying critical assumptions, and drafting 1–4 candidate questions to probe student understanding of the scenario.

The guiding question for this task was, "Will the new CCA item probe one of the identified five core cybersecurity concepts?"

We strived to place complexity into the scenarios rather than into the stems. Doing so helped enable each stem to be as short and clear as possible and to focus directly on an important concept. This strategy also reduces the required time for students to complete the test because multiple stems may share a common scenario.

Participants found it helpful to build on life experiences and to introduce an artifact, such as a program fragment, log file, protocol, or architectural diagram.

To deemphasize the importance of information knowledge, instead of referring to an object (e.g., the SSL protocol), its name or acronym, we described the crucial properties of the object (e.g., a protocol that encrypts the transferred file using a key established by a key-agreement protocol between the sender and receiver.) To deemphasize vocabulary barriers, we included definitions of any terms that students found unfamiliar (e.g., masquerade) at the end of each test item.

Task 2: Extend CCI Items Into CCA Items, and Generate New Answer Choices

These teams focused first on extending existing CCI items to have greater technical detail, sophistication, and complexity. Participants focused on the differences between students who have taken only a single course versus students who have taken an entire curriculum in cybersecurity. Guiding questions for this task were, "What do students know?" and "What misconceptions might students have about this scenario?"

After extending the CCI items into CCA items, teams focused on developing correct answer choices and distractors. To ensure that distractors reflected student misconceptions, one member of the CATS team who had previously analyzed student misconceptions in cybersecurity contributed his expertise.[10] Teams exercised leeway to modify scenarios or stems as needed to generate compelling and clear correct answers and distractors.

Task 3: Review and Refine Draft CCA Test Items

These teams refined and prioritized draft items and made notes about the scenarios, stems, and alternatives for future work. Teams first reviewed draft CCA items that the CATS team had previously created and then reviewed draft CCA items generated by Task 2 teams. Members also kept track of how many test items covered each core concept, and they estimated the approximate difficulty of each item. The guiding question for this task was, "Which scenarios and stems are worthy of inclusion in the CCA?"

Teams focused on quality control, making sure that all wording was precise, concise, and clear. One member of the CATS team experienced in crafting MCQs participated. Members ensured that each test item stated all critical assumptions. Team members answered each draft item and verified that everyone agreed on the correct answer.

Members applied best practices in writing effective MCQs, including advice offered by the Vanderbilt Center for Teaching.[1] Each stem should be meaningful by itself, and alternatives should be plausible, homogeneous, and nonoverlapping. Each test item should be easy for experts to answer but challenging for students with poor or incomplete conceptual understanding to answer.

Many difficulties could be resolved by adding more detail, especially about the assumptions and adversarial model. Whenever possible, we preferred to insert such details into the scenario rather than into the stem.

An Example: Forensic Analysis of a Network Log File

We presented a sample CCA test item that originated from Josiah Dykstra at the hackathon and evolved through several discussions and refinements, both during the hackathon and afterward by the CATS team. Dykstra is a government employee who brought significant knowledge and experience in forensics, networks, cybersecurity, and cloud computing to the hackathon.

Scenario H2. Consider the following log of corporate user activity. The corporation issues each employee a work PC and a smartphone.

Day	Time	User	Action	Device	Data Volume [kilobytes]	
21 May	20:22:28	Bob	Local login	Work PC	0 UP	0 DOWN
21 May	20:23:01	Bob	Connection to local server	Work PC	6,702 UP	244,328 DOWN
21 May	20:25:12	Bob	Access to acmeshare.com	Work PC	122,164 UP	3,456 DOWN
21 May	20:26:35	Bob	USB drive connected	Work PC	122,164 UP	0 DOWN
22 May	08:28:12	Alice	Connection to remote host	Work PC	122,164 UP	2,378 DOWN
22 May	08:32:12	Charlie	VPN login to network	Smartphone	2,490 UP	4,566 DOWN
22 May	08:38:55	Charlie	Access to acmeshare.com	Smartphone	0 UP	125,620 DOWN

Notes:
1) *acmeshare.com is a fictional, free file-sharing service.*
2) *UP* and *DOWN* data transfer volumes are given from the perspective of the device specified in the Device column. For example, in Line 2, User Bob transferred 6,702 KB from a Work PC to the local server, and User Bob transferred 244,328 KB from the local server to that PC.

Question H2-1. What is the most serious malicious activity possibly suggested by this log?

A. Bob, Alice, and Charlie cooperated to exfiltrate data.
B. Alice sent corporate secrets to some unspecified remote host.
C. Bob connected a USB drive and wrote sensitive data to it from his corporate work PC.
D. Charlie and Bob shared a malicious file via acmeshare.com.
E. Bob logged in from work at 20:22:28, after the authorized access times.

Figure 1. An example of a CCA test item that evolved from the hackathon.

Figure 1 gives the current polished version of test item H2-1. H2-1 depends on scenario H2, which introduces an artifact that is a network log file of corporate user activity. Stem H2-1 probes Core Concept 5 (i.e., identify targets and attackers) by asking the student to identify the most serious malicious activity. Note that the alternatives are plausible, homogeneous, and nonoverlapping. We suggest that the reader now pause to answer the question.

The CATS team estimates the difficulty of this test item to be medium. We consider this test item to be more appropriate for CCA than for CCI because it requires the student to understand a somewhat technical log file, however modest the technical aspects may be.

To answer this test item, the student must read and understand the log file and make inferences about it. The student must determine who the adversary or adversaries are and what they have done. To make these inferences, the student must demonstrate some technical ability to analyze a log file, common sense, and adversarial thinking in a corporate network environment.

To help us keep track of our test items and their status, for each test item, we assigned a line of metadata summarizing the item's difficulty, status, core concept, and secondary topic. The metadata for test item H2-1 are medium, ready, identify targets and attackers, and log analysis.

At the hackathon, and considering that Dykstra is an expert in forensics, we suggested that he create a scenario involving forensics. Needing more questions involving "identify targets and attackers," we encouraged him to focus on that concept. We also suggested that he introduce a technical artifact; for forensics, using a log file was a natural choice.

Originally, Dykstra proposed three stems for Scenario H2, which we shall call *H2-1a*, *H2-2a*, and *H2-3a* (see Figure 2). In the ensuing discussions, we settled on only one stem. H2-3a did not seem to exercise a very important concept, and H2-1a and H2-2a are overly similar, so the answer to one gives a major hint of the answer to the other.

We also modified the stem to focus more directly on the important targeted concept of identifying what malicious activity took place and by whom. As stems should be, Stem H2-1 is a meaningful question by itself.

After multiple meetings, the team spent significant time and effort polishing the test item. Much of that effort went into improving the clarity of the item. It is our experience that many students become confused about various details, including ones that team members had considered to be clear. Small changes in wording can affect how students perceive a test item. Our instruments should not be tests of intelligence or reading comprehension; rather, each test item should challenge a student's conceptual understanding of the targeted concept.

Edits included making the log file more uniform, inserting additional information in the log file, and clarifying the meaning of data uploads and downloads. We added clarifying details about the file-sharing service and who issued the workstations and smartphones.

Question H2-1a. Imagine you are an insider stealing corporate secrets. What change would you make in this log to cover your tracks?

A. Modify all of the data volume entries with random values.
B. Delete the records of login actions.
C. Change all the timestamps to 00:00:00.
D. Erase the action field from all records.
E. Append 500 fake records to the log.

Question H2-2a. Which inference can you draw about the attack?

A. Alice, Bob, and Charlie are colluding in the attack.
B. The attack originated from a remote, external hacker.
C. The firewall is misconfigured.
D. Bob cannot be the attacker.
E. [to be written].

Question H2-3a. What other forensic data would implicate the insider(s)?

A. Network traffic captures.
B. Intrusion detection logs.
C. Firewall logs.
D. Browser history.
E. List of deleted files.

Figure 2. The original three stems and their proposed answer choices for Scenario H2.

We also finely edited the wording, e.g., replacing the strong verb *colluded* with the softer and less suggestive term *cooperated*. Although making such edits may seem simple, our experience is that it is difficult and time consuming to construct quality test items.

In case the reader is uncertain, we note that answer choice A is the best alternative for each of the stems mentioned in this section. The sizes of the data flows provided useful clues.

The two-day hackathon resulted in four promising new CCA test items and useful feedback on all 36 draft CCI questions and 12 draft CCA questions. It also increased awareness about our project, infused new ideas into it, and established connections for possible future collaborations.

We learned that our choices for the event, including its size, length, and structure, worked well. The diversity of participants as well as their interactions contributed greatly to the event's success. Asking the participants to bring some of their favorite questions (e.g., from final exams) is an effective way to involve everyone from the beginning.

The greatest challenge to running our hackathon was finding time in the busy experts' schedules. Scheduling the hackathon in close proximity to and immediately prior to a major relevant conference made it more convenient for participants to attend. Supporting their travel also helped.

We encountered many challenges in developing quality MCQs. The process takes a significant amount of time and effort. Some appealing open-ended questions (e.g., devising or comparing a design or attack) are difficult to formulate as an MCQ without depreciating the most attractive aspects of the question. Often we found it difficult to generate more than three appealing distractors. We endeavored to make the test items as timeless as possible, but this goal was challenging to achieve, especially for the more technical CCA.

Although most experts liked the majority of our draft questions, some experts disagreed with some of our answer choices. The reason usually involved either relative weights placed on various considerations or that the expert made a hidden assumption. In such cases, we edited the test item to add details and clarify assumptions.

We are beginning to study and experiment with a new method of generating distractors: online crowdsourcing. Using Amazon Mechanical Turk, we have collected open-ended responses to draft stems.[11] We find it especially helpful to use a specific stem. Team members analyzed the responses, observed groups of similar incorrect answers, and noted whether the incorrect responses are consistent with misconceptions that we had expected. This method is fast and efficient: for US$.25 per response, we can easily collect 50–100 responses overnight. Whereas crowdsourcing is unmoderated with no collaboration and limited control of subjects, the hackathon was moderated and facilitated collaboration among carefully selected participants. Although a high percentage of the responses

appeared to lack genuine effort, we found that there was enough quality data to make this method very promising.

It would be helpful to use a suitable integrated test-development system that supports version control, collaborative test item development, record keeping, comments, expert review, cognitive interviews, pilot testing, and psychometric testing. Unfortunately, we are not aware of any such system. Instead of building and maintaining our own system, we used a variety of existing tools, especially ones that support real-time collaboration, e.g., GitHub, Google Docs (including with the Edity HTML plugin), Skype, SurveyMonkey, Excel, and PrairieLearn. When developing test items, we found it especially helpful to engage in a remote conference, during which the participants could simultaneously edit a common file. For each test item, it is highly desirable to maintain exactly one authentic source file to avoid the inevitable errors that result from copying or converting test items. Edity helped us achieve this goal, albeit imperfectly, given that PrairieLearn inputs test items as HTML files.

As evidenced from feedback submitted via a SurveyMonkey questionnaire, the majority of participants found the hackathon fruitful and that it produced valuable products. Participants stated that the collaboration with diverse stakeholders was particularly valuable in addressing the diverse and evolving field of cybersecurity education. All of the participants indicated that they would be willing to continue contributing to the development of the instruments and that they would administer pilot versions of the tests.

In 2019, we will complete development of the draft CCA while beginning validation studies of the CCI. These validation studies include cognitive interviews, expert review, small-scale pilot testing, and large-scale psychometric testing. We also plan to carry out several half-day "minihackathons" associated with various cybersecurity educational conferences. We welcome participation in these and future studies for the CCA.

Our experience with the hackathon demonstrates that this type of collaborative workshop is an effective way to generate and improve test items and raise awareness about the project. We hope that the resulting instruments will help identify effective strategies for teaching and learning cybersecurity concepts. ■

Acknowledgments

We would like to thank all of the hackathon participants. This work was supported in part by the U.S. Department of Defense under Centers for Academic Excellence grants H98230-15-1-0294, H98230-15-1-0273, H98230-17-1-0349, and H98230-17-1-0347; the National Science Foundation under Scholarship for Service grant 1241576; and the Department of Graduate Education grant 1820531.

References

1. C. J. Brame, "Writing good multiple choice test questions." Accessed on: Jan. 19, 2019. [Online]. Available: https://cft.vanderbilt.edu/guides-sub-pages/writing-good-multiple-choice-test-questions/
2. G. L. Herman, C. C. Zilles, and M. C. Loui, "A psychometric evaluation of the digital logic concept inventory," *Comput. Sci. Educ.*, vol. 24, no. 4, pp. 277–303, 2014.
3. R. Hambleton and R. J. Jones, "Comparison of classical test theory and item response theory and their applications to test development," *Educational Meas.: Issues Practice*, vol. 12, no. 3, pp. 38–47, Sept. 1993. doi: 10.1111/j.1745-3992.1993.tb00543.x.
4. D. Hestenes, M. Wells, and G. Swackhamer, "Force concept inventory," *Phys. Teacher*, vol. 30, no. 3, pp. 141–166, 1992.
5. N. Jorion, B. D. Gane, K. James, L. Schroeder, L. V. DiBello, and J. W. Pellegrino, "An analytic framework for evaluating the validity of concept inventory claims," *J. Eng. Educ.*, vol. 104, no. 4, pp. 454–496, 1995.
6. G. Parekh et al., "Identifying core concepts of cybersecurity: Results of two Delphi processes," *IEEE Trans. Edu.*, vol. 61, no. 1, pp. 11–20, 2016.
7. D. H. Pink, *Drive: The Surprising Truth About What Motivates Us*. New York: Riverhead Trade, 2011.
8. A. T. Sherman et al., "Cybersecurity: Exploring core concepts through six scenarios," *Cryptologia*, vol. 42, no. 4, pp. 337–377, 2018.
9. A. T. Sherman et al., "Creating a cybersecurity concept inventory: A status report on the CATS Project," in *Proc. Nat. Cyber Summit*, 2017. [Online]. Available: https://arxiv.org/abs/1706.05092
10. J. Thompson et al., "Student misconceptions about cybersecurity concepts: Analysis of think-aloud interviews with students," *J. Cybersecurity Educ.*, vol. 1, no. 5, pp. 1–29, 2018.
11. T. Scheponik et al., "Investigating crowdsourcing to generate distractors for multiple-choice assessments," in *Proc. Nat. Cyber Summit*, June 2019.

Alan T. Sherman is a professor of computer science at the University of Maryland, Baltimore County. His research interests include secure voting, applied cryptography, and cybersecurity education. He is a Senior Member of the IEEE. Contact him at sherman@umbc.edu.

Linda Oliva is an assistant professor in the Education Department at the University of Maryland, Baltimore County. Her research interests include cybersecurity

education, teacher efficacy, domain-specific expertise, and program evaluation. Contact her at oliva@umbc.edu.

Enis Golaszewski is a Ph.D. student at the University of Maryland, Baltimore County (UMBC). Golaszewski received a B.S. in computer science from UMBC. Contact him at golaszewski@umbc.edu.

Dhananjay Phatak is an associate professor in the Computer Science and Electrical Engineering (CSEE) Department at the University of Maryland, Baltimore County. His research interests include computer arithmetic algorithms and implementations, number theory, cryptology and efficient realization of cryptographic algorithms and primitives, and cybersecurity. Phatek received a Ph.D. in CSEE, an M.S. in microwave planar integrated circuits, both from the University of Massachusetts, and a B.Tech. in electrical engineering from IIT Bombay, Mumbai. Contact him at phatak@umbc.edu.

Travis Scheponik is a Ph.D. student in the Computer Science and Engineering Department at the University of Maryland, Baltimore County, and a software engineer at Analysis, Computing & Engineering Solutions. His main research interest is cybersecurity. Contact him at tschep1@umbc.edu.

Geoffrey L. Herman is a teaching assistant professor of computer science at the University of Illinois at Urbana–Champaign. His research focuses on studying how students learn engineering and computer science concepts and assessing students' knowledge of these concepts. Herman received a Ph.D. in electrical and computer engineering. Contact him at glherman@illinois.edu.

Dong San Choi is a university lecturer at University of Illinois at Urbana–Champaign. Choi received a master's degree in computer engineering. Contact him at choi88@illinois.edu.

Spencer E. Offenberger is a graduate student in electrical and computer engineering at the University of Illinois at Urbana–Champaign. His research interests include robotics, vision, and artificial intelligence. Contact him at so@illinois.edu.

Peter Peterson is an assistant professor of computer science at Swenson College of Science and Engineering, Duluth, Minnesota. Peterson received a Ph.D. in computer science. Contact him at pahp@d.umn.edu.

Josiah Dykstra is a researcher and technical lead at the U.S. Department of Defense. He received his Ph.D. in computer science at University of Maryland, Baltimore County, researching the technical and legal challenges of digital forensics for cloud computing. His research interests include network security, intrusion detection, malware analysis, digital forensics, and cloud computing. He is active in the academic research community, serving on conference committees including Usenix Security and the Digital Forensics Research Workshop. Josiah is a member of ACM, IEEE, American Academy of Forensic Sciences, Cloud Security Alliance, NIST Cloud Forensics Working Group, IFIP Working Group 11.9 on Digital Forensics, and American Bar Association E-Discovery and Digital Evidence Committee. Contact him at josiahdykstra@acm.org.

Gregory V. Bard is a faculty member in the Department of Mathematics, Statistics and Computer Science, at the University of Wisconsin–Stout. His current focus is computer algebra. Contact him at bardg@uwstout.edu.

Ankur Chattopadhyay is an assistant professor in the Department of Information and Computing Sciences at the University of Wisconsin, Green Bay. He received a Ph.D. in computer science from the University of Colorado, Colorado Springs, with a graduate certification of accomplishment in information assurance. Contact him at chattopa@uwgb.edu.

Filipo Sharevski is an assistant professor with the College of Computing and Digital Media at DePaul University. His research interests include computer security, human–computer interaction, and telecommunication and networking. He received a Ph.D. in interdisciplinary cybersecurity from Purdue University, West Lafayette, Indiana. Contact him at fsharevs@cdm.depaul.edu.

Rakesh Verma is a professor of computer science at the University of Houston and director of its Reasoning and Data Analytics for Security Laboratory. His current research interests are formal methods and data analytics applied to natural language understanding and cybersecurity. Contact him at rverma@uh.edu.

Ryan Vrecenar is a computer and electrical engineering student at Texas A&M University. Vrecenar received a master's degree in computer engineering. Contact him at ryanvrecenar@email.tamu.edu.