What Are Cybersecurity Education Papers About? A Systematic Literature Review of SIGCSE and ITiCSE Conferences

Valdemar Švábenský
Masaryk University
Czech Republic
svabensky@ics.muni.cz

Jan Vykopal
Masaryk University
Czech Republic
vykopal@ics.muni.cz

Pavel Čeleda
Masaryk University
Czech Republic
celeda@ics.muni.cz

ABSTRACT

Cybersecurity is now more important than ever, and so is education in this field. However, the cybersecurity domain encompasses an extensive set of concepts, which can be taught in different ways and contexts. To understand the state of the art of cybersecurity education and related research, we examine papers from the ACM SIGCSE and ACM ITiCSE conferences. From 2010 to 2019, a total of 1,748 papers were published at these conferences, and 71 of them focus on cybersecurity education. The papers discuss courses, tools, exercises, and teaching approaches. For each paper, we map the covered topics, teaching context, evaluation methods, impact, and the community of authors. We discovered that the technical topic areas are evenly covered (the most prominent being secure programming, network security, and offensive security), and human aspects, such as privacy and social engineering, are present as well. The interventions described in SIGCSE and ITiCSE papers predominantly focus on tertiary education in the USA. The subsequent evaluation mostly consists of collecting students' subjective perceptions via questionnaires. However, less than a third of the papers provide supplementary materials for other educators, and none of the authors published their dataset. Our results provide orientation in the area, a synthesis of trends, and implications for further research. Therefore, they are relevant for instructors, researchers, and anyone new in the field of cybersecurity education. The information we collected and synthesized from individual papers are organized in a publicly available dataset.

CCS CONCEPTS

• General and reference → Surveys and overviews; **• Social and professional topics → Computing education**; **• Security and privacy**;

KEYWORDS

cybersecurity education, systematic literature review, systematic mapping study, survey, SIGCSE community, ITiCSE community

ACM Reference Format:
Valdemar Švábenský, Jan Vykopal, and Pavel Čeleda. 2020. What Are Cybersecurity Education Papers About? A Systematic Literature Review of SIGCSE and ITiCSE Conferences. In *The 51st ACM Technical Symposium on Computer Science Education (SIGCSE '20), March 11–14, 2020, Portland, OR, USA.* ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/3328778.3366816

SIGCSE '20, March 11–14, 2020, Portland, OR, USA

1 INTRODUCTION

With the rising importance of cybersecurity and its global job vacancy increasing to 2.93 million [13], there is an imminent need for training more cybersecurity workers. As a result, a broad range of educational initiatives arose to address this need. In 2013, cybersecurity was included in ACM/IEEE computing curricula [15]. Four years later, the Joint Task Force on Cybersecurity Education (JTF) published comprehensive curricular guidance [16] to help educators design cybersecurity courses. In addition, informal methods of education, such as extracurricular events [7] and competitions [35] for learners of all ages and expertise are gaining popularity.

SIGCSE 2020 begins a new decade of computing education research. Therefore, we feel it is appropriate to review the research advancements and teaching methods we have seen over the last ten years. In this paper, we examine the development and state of the art of cybersecurity education as presented at ACM SIGCSE and ACM ITiCSE conferences from 2010 to 2019. We have chosen these two conferences because they represent the leading venues in the area of computing education. Apart from their rich history and a large number of quality submissions, they are currently the only two conferences in the field that rank as CORE A [30].

This literature review brings several contributions to various target groups. For cybersecurity instructors and educational managers, it shows what topics are taught and how. For researchers, it provides an overview of evaluation methods, implications for further research, and practical recommendations. Finally, for the SIGCSE/ITiCSE community as a whole, it serves as a snapshot of ten years of the latest development and synthesis of accepted papers.

2 RELATED WORK

This section maps related primary and secondary studies in two areas: cybersecurity education and computing education in general. We build upon and extend the below-mentioned works by focusing on the cybersecurity domain at SIGCSE and ITiCSE.

2.1 Cybersecurity Education

Although some papers reviewed or synthesized results of cybersecurity education efforts, none of them focused on SIGCSE and ITiCSE. Fujs et al. [9] performed a literature review on using qualitative methods in cybersecurity research, which includes research on cybersecurity education. Next, Cabaj et al. [3] examined 21 cybersecurity Master degree programs and their mapping to ACM/IEEE curricular guidelines [15]. Jones et al. [17] studied the core knowledge, skills, and abilities that cybersecurity professionals need. Parrish et al. [27] provided an overview of cybersecurity curricula.

Also, there are synthesis papers on the topic of Capture the Flag (CTF). This popular format of competitions or educational events

Table 1: Number of full papers at SIGCSE | ITiCSE conferences over the years 2010–2019 after each step of the literature review

SIGCSE | ITiCSE	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019	Total
Published papers	103 | 60	107 | 66	100 | 61	111 | 51	108 | 53	105 | 54	105 | 51	105 | 56	161 | 56	169 | 66	**1174 | 574**
Candidate papers	5 | 2	4 | 1	1 | 2	8 | 2	3 | 2	6 | 2	7 | 3	5 | 1	14 | 4	15 | 3	**68 | 22**
Selected papers	5 | 1	3 | 1	0 | 2	6 | 1	2 | 2	5 | 2	6 | 1	5 | 1	12 | 3	11 | 2	**55 | 16**

allows the participants to practice their cybersecurity skills. Taylor et al. [35] mapped the state of the art of 36 implementations of CTF. Next, Burns et al. [2] analyzed the solutions of 3,600 CTF challenges from 160 competitions to discover which challenge types are the most popular and what is their difficulty. Finally, Vigna et al. [38] presented a framework for organizing CTF events, building upon ten years of experience with running the largest educational CTF.

2.2 Other Areas of Computing Education

Many literature surveys in computing education focus on programming. ITiCSE Working Groups often create comprehensive reviews, such as the one by Luxton-Reilly et al. [22] on introductory programming education. The authors inspected 1,666 papers from 15 years. Ihantola et al. [12] surveyed a decade of 76 papers on educational data mining and learning analytics in programming. Next, Becker and Quille [1] reviewed topics of 481 papers about CS1 from 50 years of SIGCSE conferences. Finally, Keuning et al. [19] compared 101 tools that provide automated feedback to students who solve programming assignments. All these studies provide an excellent example of conducting literature reviews.

3 METHOD OF CONDUCTING THE REVIEW

There are two similar approaches for surveying research papers: a systematic literature review (SLR) and a systematic mapping study (SMS). Both approaches aim at summarizing the existing research, discovering research trends, and identifying gaps in current research. Although SLR and SMS slightly differ (see [20, 28, 29]), these differences are negligible for the purposes of this paper, which could be classified as either SLR or SMS. The methods we applied follow the well-established guidelines for SLR [20] and SMS [28, 29].

3.1 Research Questions

Our literature review examines five research questions:

(1) *What cybersecurity topics are discussed in the papers?*
(2) *How and in what context are the topics taught?*
(3) *How are the teaching interventions evaluated?*
(4) *What is the impact of the published papers?*
(5) *Who are the members of the cybersecurity education community at SIGCSE and ITiCSE?*

Our motivation for each of the research questions was:

(1) To discover which topics are taught and researched a lot and whether there are any underrepresented topics.
(2) To describe the common teaching practices.
(3) To examine research methods in cybersecurity education.
(4) To find out whether the community adopts the published teaching innovations and research results.
(5) To understand who forms the SIGCSE/ITiCSE community.

We address the research questions by extracting and analyzing data from 71 papers. The following Sections 3.2–3.4 describe the process. Table 1 provides context by showing how many papers were published within the selected time range. It also presents how the number of papers evolved at each step of our literature review.

3.2 Paper Search and Automated Filtering

We searched for papers in the ACM Digital Library by submitting the query: `cybersecur* OR secur*`. We used this broad search term to cover a wide area and avoid the risk of missing a relevant paper. We then refined the results to publications in conference proceedings since 2010. Further restriction on the SIGCSE conference yielded 209 results; the ITiCSE conference yielded 52 results. We searched for SIGCSE papers on May 30, 2019, and for ITiCSE papers on July 15, 2019. We also searched for ICER conference papers (on August 15, 2019), which surprisingly yielded 0 results.

As the next step, we removed 1- or 2-pages long submissions, which included, for example, poster abstracts and panel proposals. Although these submissions indicate a general topic of interest, they do not have enough space to discuss details relevant to our research questions. Afterward, we were left with 90 candidate full papers (68 from SIGCSE and 22 from ITiCSE), each 5–7 pages long.

3.3 Pilot Reading and Manual Filtering

Two authors preliminarily read the 90 papers to mark candidate false positives. The reading was independent to prevent bias. The resulting inter-rater agreement measured by Cohen's kappa [21] was 0.87, which is an "almost perfect agreement". The remaining discrepancies were only minor and were resolved by discussion. In the end, we eliminated 19 false positives, selecting 71 papers for the literature review. The most common reason for exclusion was that the paper mentioned the word "security", for example, in the general context that it is an essential concept, but did not deal with anything related to cybersecurity education.

3.4 Full Text Reading and Data Extraction

For each research question, we drafted several criteria that defined what kind of data we would extract from the papers. Subsequently, we performed a pilot test on 30 papers to see if these criteria are reasonable and if the data can indeed be extracted. Over the course of this action and several discussions among the authors, the criteria iteratively evolved. We present their final form in Section 4 along with the corresponding result to save space and improve readability.

Upon agreeing on the data extraction criteria and process, we thoroughly read the selected 71 papers and documented the results, creating our dataset [46]. When in doubt, we discussed unclear cases to reach agreement. Section 4 presents the synthesized results.

3.5 Limitations

This review is limited by narrowing its scope only to the SIGCSE, ITiCSE, and ICER conferences. However, other conferences, such as IEEE FIE and USENIX ASE/3GSE, also publish cybersecurity education papers. There are related journals as well, such as ACM TOCE, ACM Inroads, and Elsevier Computers and Education. Nevertheless, as mentioned in Section 1, the SIGCSE and ITiCSE conferences are the flagships in the field, and so we consider them representative of the trends within the cybersecurity education community.

We are confident that our broad search query captured all relevant papers. The filtering of false positives was double-checked. However, the data extraction from the papers might be problematic. Since it was done manually, the readers may have overlooked or misinterpreted something. Nevertheless, we performed cross-author discussion and validation to minimize the risk of incorrectness.

4 RESULTS

We now present the selected results that synthesize the data extracted from the 71 papers. The dataset of full records is publicly available [46]. Each Section 4.1–4.5 is mapped to the corresponding research question from Section 3.1. When needed, we include examples of representative papers within the category[1].

4.1 RQ1: What Topics Are Taught?

To address the first research question, we identified general topic areas as listed in the JTF Cybersecurity Curriculum [16], as well as specific topics discussed in the papers.

4.1.1 Which Knowledge Areas from the JTF Cybersecurity Curriculum do the papers focus on? The curriculum [16] consists of eight Knowledge Areas (KA), each including several Knowledge Units (KU). Figure 1 shows the distribution of how often each KA was present. The most frequent KUs (in 10 or more papers) were Network Defense (19 papers), Implementation (18 papers), System Control (15 papers), and Cryptography (14 papers). Nevertheless, performing this mapping was sometimes tricky. The papers differ in the level of detail when describing the teaching content, and only 10 papers reference a standardized curriculum to define their learning objectives (most frequently ACM/IEEE guidelines [15]).

4.1.2 What are the primary cybersecurity topics? We performed open coding of the topics that the individual papers focused on. The most commonly assigned codes (to 10 or more papers) were:

- Secure programming and software development, including reverse engineering (24 papers).
- Network security and monitoring (23 papers).
- Cyber attacks, malware, hacking, offensive security, and exploitation (17 papers).
- Human aspects, including privacy, social engineering, law, ethics, and societal impact (17 papers).
- Cryptography (15 papers).
- Authentication and authorization (13 papers).

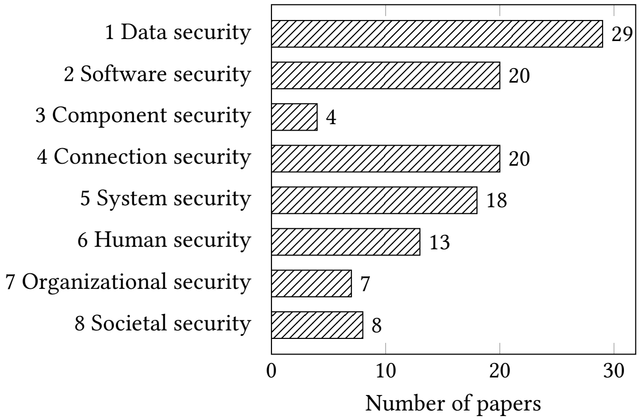

Figure 1: The distribution of how often each of the eight JTF Cybersecurity Curriculum Knowledge Areas [16] was discussed in the selected papers.

While this coding was performed independently from the mapping to the JTF Curriculum, the findings are similar to the most frequent KUs mentioned earlier, supporting the validity of the results.

4.2 RQ2: How Are the Topics Taught?

Next, we examined the target groups, teaching methods, time frames, and student population sizes.

4.2.1 Whom are the teaching interventions aimed at? The most prominent target group for teaching interventions in 54 papers are university or college students (undergraduates or graduates[2]). Other target groups are in the minority: instructors and educational managers in 7 papers, K-12 students (middle or high school) in 7 papers, and professional learners in 3 papers. Lastly, 8 papers do not explicitly report their target group, although in most cases, the paper content suggests university students.

4.2.2 Which teaching methods are used? The subject of 64 papers (90%) is something that was applied in teaching practice, such as a tool, an exercise, or a full course. The remaining 7 papers only propose or discuss an idea without having it used in actual teaching. Within these 64 papers that describe a teaching intervention, the most common teaching method mentioned in 51 papers is some form of hands-on learning during class time or self-study. This includes labs, exercises, practical assignments, educational games, and other activities for practicing the selected topic. Other common teaching methods are lectures (24 papers), long-term projects that typically span the whole semester (10 papers), discussion (8 papers), and writing (7 papers). Exactly half of the 64 teaching papers mentions involving students in pairwork or groupwork.

4.2.3 What is the time frame? 4 papers study an intervention only lasting up to one hour, typically a single short assignment. Next, 8 papers deal with the time frame of several hours: typically an exercise or a workshop within a single day. 4 papers study a period of several days up to one week. 9 papers are within the time window of several weeks or one month. The most papers – 23 – study an

[1] A single paper may correspond to multiple categories. Therefore, the counts reported in respective categories sometimes add up to more than the total of 71 papers. Also note that some data, such as the citation counts or the availability of hyperlinks, are bound to the time of writing this section (the second half of August 2019).

[2] Including midshipmen or cadets in military academies in 3 papers.

intervention lasting more than 1 month, typically a semester-long course. Related to that, 27 papers describe experience or data from a single run of the intervention, 12 papers from two runs, and 17 papers from three or more runs.

4.2.4 How big are the participant populations? Out of the 64 papers that discuss practical teaching interventions, we looked at how many students participated in it. Since the reporting was largely heterogeneous, we performed several simplifications. First, if the paper reported a repeated intervention, we considered the total sum of the participants. Second, if the paper reported a range, we considered the average number. Third, if the paper reported only the sample size present in the evaluation, we picked this number as the lower bound, although the real number of participants could have been higher. So, we report only what was discernible from the paper. The median number of participants was 62.5, and the distribution was as follows: 3 papers had 1–19 participants (minimum 17), 19 papers had 20–49 participants, 15 papers had 50–99 participants, and 17 papers had 100 or more participants (maximum 14,000). 9 papers did not report the number, and for [5], this was not applicable, as it dealt with ethical policies, not human participants.

4.3 RQ3: How Is Evaluation Performed?

This section examines whether the papers present any evaluation. If they do, we investigate what was evaluated, what data were collected, and how the data were analyzed. We finish the section by looking at sample sizes and publishing of datasets.

4.3.1 What is the goal of the evaluation? We examined what aspects related to teaching and learning were evaluated in the papers. By doing so, we synthesized four evaluation goals (EG):

EG1 *Subjective perception of participants*, which includes learners' or teachers' attitudes and opinions of the subject of the paper (for example, a teaching intervention, course, or a cybersecurity exercise). The evaluation focuses on whether the participants reported self-perceived learning or found the subject useful, effective, understandable, easy/difficult, enjoyable, motivational, or supporting their career interest. There were 50 papers in this category, such as [24].

EG2 *Objective learning of students*, which includes measuring learners' performance with summative assessment and computing test scores or grades to determine learning. There were 21 papers in this category, such as [34].

EG3 *Artifacts produced by students*, which includes examining submissions of assignments [11], screen recordings from a learning environment [33], or logs from using a tool [43]. The focus was usually on better understanding students and their interactions, not necessarily measuring their learning as in EG2. There were 8 papers in this category.

EG4 *Other artifacts*, which includes examining textbooks [36], reviewing of teaching materials by teachers [10], or analyzing ethical policies [5]. There were 5 papers in this category.

Finally, 8 papers did not present any evaluation, such as [18] that only described a tool and its possible usage scenarios.

4.3.2 What types of data are collected? We synthesized six data types (DT) collected for the evaluation:

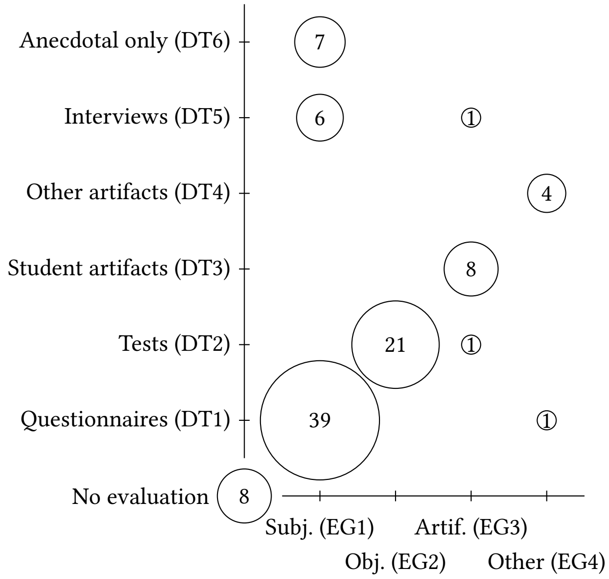

Figure 2: The mapping of evaluation goals (EG) to the collected data types (DT) and the corresponding frequency.

DT1 *Questionnaire data*, most often about participants' subjective perceptions (EG1). Out of 40 papers in this category, 26 employed post-intervention questionnaires only (such as [37]), while the remaining 14 used pre- and post-questionnaire study design (such as [8]).

DT2 *Test data*, the results of summative assessment, such as grades or the number of assignments solved, usually to measure objective learning (EG2). Out of 22 papers in this category, 10 employed post-test only (for example, [26], although the test includes just two questions), and the remaining 12 used pre- and post-test study design (such as [34]).

DT3 *Student artifacts* (see EG3 for examples). There were 8 papers in this category.

DT4 *Other artifacts* (see EG4 for examples). There were 4 papers in this category.

DT5 *Interviews* and focus group discussions to most often examine participants' subjective perceptions (EG1). There were 7 papers in this category, such as [32].

DT6 *Anecdotal only*, if the paper did not present any of the above evidence types (DT1–DT5), but only reported the authors' or participants' experience, for example, from course feedback, not backed up with any research tool. There were 7 papers in this category, for example, [31].

Figure 2 shows the mapping of EG to DT. The 8 papers that did not present any evaluation also did not collect any evidence.

4.3.3 How are the collected data analyzed? We identified three analysis methods (AM) for interpreting the collected data:

AM1 *Descriptive statistics*, which were most often used to analyze questionnaire data (DT1) and test data (DT2). A majority of 54 papers employed this evaluation method, although the depth of the statistics varied from a single histogram [44] to detailed tables including confidence intervals [7].

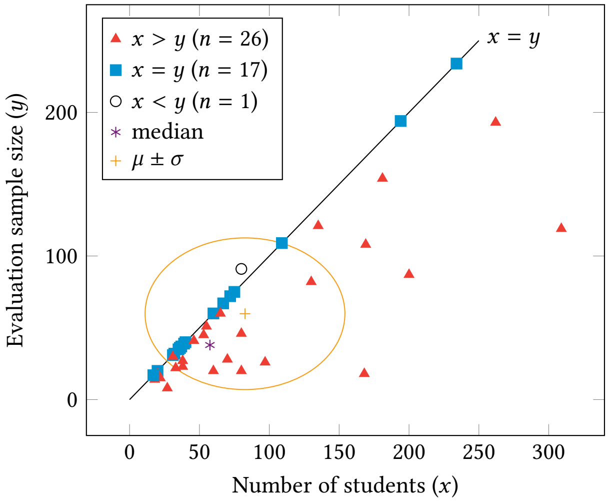

Figure 3: The number of students participating in a teaching intervention compared to the size of the sample that participated in the subsequent evaluation (44 data points).

AM2 *Inferential statistics*, which involved formulating a hypothesis, applying an appropriate statistical test, and reporting the results, including p-value and sample size. There were 19 papers in this category, and although 4 more had attempted to use inferential statistics, the reporting had been insufficient. For example, although [4] presents a p-value, information about which test was applied is missing. Another transgression is present in [14]: t-test was performed, but the statistic value and the p-value are not reported. On the contrary, [8] is a great example of a thorough statistical evaluation.

AM3 *Qualitative analysis*, which included expert reviews of both student and non-student artifacts (DT3, DT4), their visual analysis, or thematic coding of interviews (DT5). There were 10 papers in this category.

15 papers did not employ any analysis method, since 8 of them did not collect any evidence, and the remaining 7 collected only anecdotal evidence, for which no special analysis was performed.

4.3.4 What is the sample size? We also examined the sample size in papers that evaluated students. The median sample size was 40.5, and the distribution was as follows: 6 papers had 1–19 participants (minimum 8), 21 papers had 20–49 participants, 10 papers had 50–99 participants, and 11 papers had 100 or more participants (maximum 1,974). 11 papers did not report the number, and for 12 papers this was not applicable, as they either performed no evaluation or did not evaluate people. Figure 3 puts this data[3] in context to the participant population reported in Section 4.2.4. There is a visible trend that not all students participate in the evaluation.

4.3.5 Can the evaluation be replicated? Surprisingly, none of the examined papers provides a link to the corresponding dataset that was analyzed. Although one paper [6] includes a link to a Gitlab repository with the data, the repository is unavailable. Nevertheless, papers such as [45] present exemplary practice by including the full wording of the questions that the evaluators asked.

4.4 RQ4: What Is the Impact of the Papers?

Our next question was whether the papers influenced other researchers and practitioners. However, as Malmi [25] argues, it is difficult to measure the impact of computing education research. He discusses counts of citations and paper downloads as two possible but imperfect metrics. For tools, he suggests download count as well, but this information is usually private and thus inapplicable. Therefore, we considered three metrics: providing a publicly available output, paper download count, and citation count.

4.4.1 Do the papers provide output usable by other educators? 30 of the 71 examined papers reference output for other educators. This most often includes course materials, lab modules/exercises, or software tools. The authors use dedicated websites, their institutional websites, or public repositories to share the content. However, out of the 30 linked websites, only 22 are still available, leaving 49 of the 71 papers (69%) without any supplementary materials.

4.4.2 How much are the papers downloaded and cited? The median cumulative number of downloads from the ACM Digital Library is 176 (min = 19, max = 1120, $\mu = 223.1$, $\sigma = 207.3$). The citation analysis on Scopus[4] showed that the median citation count is 2 (min = 0, max = 18, $\mu = 3.3$, $\sigma = 4.4$). After removing self-citations, the median dropped to 1 (min = 0, max = 17, $\mu = 2.5$, $\sigma = 3.8$). We also looked at how many of the non-self-citations are from the SIGCSE or ITiCSE community. The median is 0 (min = 0, max = 5, $\mu = 0.5$, $\sigma = 1$). This shows that the papers are rarely cited, not even within the community. Nevertheless, these metrics are biased toward older papers, which have a higher chance of being downloaded or cited compared to the recently released papers. Lastly, to complement the point of view, the examined papers themselves include a median of 18 references (min = 4, max = 39, $\mu = 18.1$, $\sigma = 7$).

4.5 RQ5: Who Forms the SIGCSE and ITiCSE Cybersecurity Community?

Finally, we examined the people that publish cybersecurity education papers and their affiliations.

4.5.1 Who publishes cybersecurity education research? Within the selected papers, there were 251 author listings and 202 unique authors[5], out of which 175 – a vast majority – appeared only in one paper. This implies a lack of cybersecurity education researchers dedicated to publishing at SIGCSE and ITiCSE. However, 14 authors appeared in two papers, and the remaining 13 authors appeared in three or more papers, which suggests that there is a small but

[3] Two outliers with huge sample sizes were removed from the plot for readability. The remaining 25 papers either did not report x or y values, or this distinction was not applicable. The one point above the $x = y$ line is the paper [26], in which the authors reported that 80 students participated in the presented exercise but then report evaluation results from a test taken by 91 students.

[4] We also considered Web of Science and Google Scholar databases, but disregarded them, since Web of Science does not index all years 2010–2019 of both conferences, and Google Scholar indexes a lot of lower-quality citations, such as bachelor's theses.

[5] We manually double-checked the automatic analysis of authors' names to account for minor differences in how the same authors list their names in different papers, for example, Heather Richter Lipford [33] vs. Heather Lipford-Richter [43].

stable community of cybersecurity educators. The most prolific authors with five papers were Ching-Kuang Shene and Jean Mayo with their series of visualization papers on ITiCSE [23, 39–42].

4.5.2 What are the authors' affiliations? We looked at the authors' affiliations and countries to better understand the demographics of the community. Out of the 251 author listings, the vast majority were affiliated to universities (190) and colleges (31), following with military institutions (22), research centers (7), and private companies (2)[6]. The most represented country was the USA (203), followed by Canada (17) and the Czech Republic (10). This corresponds to the fact that most SIGCSE/ITiCSE papers examine higher education interventions within the context of the USA.

5 DISCUSSION AND CONCLUSIONS

We performed a systematic literature review of 71 cybersecurity education papers from SIGCSE and ITiCSE, the leading conferences in the field of computing education, over the period from 2010 to 2019. Our dataset is publicly available as supplementary material in the ACM Digital Library and also on Zenodo [46]. Apart from reviewing the current literature, we also provided a framework for future researchers by listing the possible evaluation goals, evidence types, and analysis methods. We now summarize the most commonly observed trends in the examined papers and provide recommendations for both research and practice.

5.1 Summary of the Observed Trends

A typical SIGCSE/ITiCSE cybersecurity education paper deals with topics such as secure software development, network security, cyber attacks, cryptography, or privacy. It describes a course, hands-on exercise, or a tool applied in teaching practice, in the context of a North American university. It usually reports data and teaching experience from a period of one semester, with a population of several dozens of undergraduate students.

Considering the research goals, the typical evaluation examines subjective experiences and perceptions of students, using questionnaires as the most common research tool. Also, pre- and post-test study designs are standard to examine learning gains after the applied teaching intervention. The evaluation is performed on a subset of the student population. The results are presented with descriptive statistics; sometimes, inferential statistics are used to confirm relationships within the data.

Even though most papers mention creating new tools or teaching materials, only 31% of the papers provide an output available to other educators and researchers. Surprisingly, no paper includes a dataset as supplementary material. Finally, a small number of citations of the papers may suggest that cybersecurity education research is fragmented. A possible explanation is that the researchers explore disjoint topics and rarely use others' results. What is more, almost 87% of the unique authors who contributed to SIGCSE/ITiCSE cybersecurity education research did so only once.

5.2 Implications of this Literature Review

Several research ideas stem from this review. Since K-12 education was underrepresented, it may be worthwhile to examine teaching interventions with younger learners. Next, since not all students participated in the evaluation, exploring how to motivate them to take part in education research can be valuable. Moreover, as most papers used questionnaires or tests for evaluation, researchers may consider employing approaches of educational data mining or learning analytics to better understand students' learning processes. Lastly, it would be interesting to compare research trends in cybersecurity conferences with computing education conferences.

To support high-quality research, we recommend future authors to familiarize themselves with exemplary papers and subsequently perform more rigorous evaluations while sharing their datasets. The community would also benefit from more thorough reporting of the research methods. In some papers, the description of the methods was unclear or incomplete, complicating the extraction of data for this review. To support teaching practitioners, we suggest using standardized documents such as the JTF Curriculum [16] to precisely define learning outcomes and addressed topics. Also, using stable public repositories to share content would be beneficial, since the tools published in 8 out of 30 papers are no longer accessible.

ACKNOWLEDGMENTS

This research was supported by the ERDF project *CyberSecurity, CyberCrime and Critical Information Infrastructures Center of Excellence* (No. CZ.02.1.01/0.0/0.0/16_019/0000822). We also thank Pavel Šmerk for his assistance with the citation analysis.

REFERENCES

[1] Brett A. Becker and Keith Quille. 2019. 50 Years of CS1 at SIGCSE: A Review of the Evolution of Introductory Programming Education Research. In *Proceedings of the 50th ACM Technical Symposium on Computer Science Education (SIGCSE '19)*. ACM, New York, NY, USA, 338–344. https://doi.org/10.1145/3287324.3287432

[2] Tanner J. Burns, Samuel C. Rios, Thomas K. Jordan, Qijun Gu, and Trevor Underwood. 2017. Analysis and Exercises for Engaging Beginners in Online CTF Competitions for Security Education. In *2017 USENIX Workshop on Advances in Security Education (ASE 17)*. USENIX Association, Vancouver, BC, 9. https://www.usenix.org/conference/ase17/workshop-program/presentation/burns

[3] Krzysztof Cabaj, Dulce Domingos, Zbigniew Kotulski, and Ana Respício. 2018. Cybersecurity education: Evolution of the discipline and analysis of master programs. *Computers & Security* 75 (2018), 24–35. https://doi.org/10.1016/j.cose.2018.01.015

[4] Justin Cappos and Richard Weiss. 2014. Teaching the Security Mindset with Reference Monitors. In *Proceedings of the 45th ACM Technical Symposium on Computer Science Education (SIGCSE '14)*. ACM, New York, NY, USA, 523–528. https://doi.org/10.1145/2538862.2538939

[5] Benedict Chukuka and Michael Locasto. 2016. A Survey of Ethical Agreements in Information Security Courses. In *Proceedings of the 47th ACM Technical Symposium on Computing Science Education (SIGCSE '16)*. ACM, New York, NY, USA, 479–484. https://doi.org/10.1145/2839509.2844580

[6] Pranita Deshpande and Irfan Ahmed. 2019. Topological Scoring of Concept Maps for Cybersecurity Education. In *Proceedings of the 50th ACM Technical Symposium on Computer Science Education (SIGCSE '19)*. ACM, New York, NY, USA, 731–737. https://doi.org/10.1145/3287324.3287495

[7] Michael H. Dunn and Laurence D. Merkle. 2018. Assessing the Impact of a National Cybersecurity Competition on Students' Career Interests. In *Proceedings of the 49th ACM Technical Symposium on Computer Science Education (SIGCSE '18)*. ACM, New York, NY, USA, 62–67. https://doi.org/10.1145/3159450.3159462

[8] Serge Egelman, Julia Bernd, Gerald Friedland, and Dan Garcia. 2016. The Teaching Privacy Curriculum. In *Proceedings of the 47th ACM Technical Symposium on Computing Science Education (SIGCSE '16)*. ACM, New York, NY, USA, 591–596. https://doi.org/10.1145/2839509.2844619

[9] Damjan Fujs, Anže Mihelič, and Simon L. R. Vrhovec. 2019. The Power of Interpretation: Qualitative Methods in Cybersecurity Research. In *Proceedings of the 14th International Conference on Availability, Reliability and Security (ARES '19)*. ACM, New York, NY, USA, 92:1–92:10. https://doi.org/10.1145/3339252.3341479

[10] Binto George, Martha Klems, and Anna Valeva. 2013. A Method for Incorporating Usable Security into Computer Security Courses. In *Proceeding of the 44th ACM*

[6] One author in [8] had two affiliations, therefore, the sum of the affiliations is 252.

Technical Symposium on Computer Science Education (SIGCSE '13). ACM, New York, NY, USA, 681–686. https://doi.org/10.1145/2445196.2445395

[11] Sara Hooshangi, Richard Weiss, and Justin Cappos. 2015. Can the Security Mindset Make Students Better Testers?. In *Proceedings of the 46th ACM Technical Symposium on Computer Science Education (SIGCSE '15).* ACM, New York, NY, USA, 404–409. https://doi.org/10.1145/2676723.2677268

[12] Petri Ihantola, Arto Vihavainen, Alireza Ahadi, Matthew Butler, Jürgen Börstler, Stephen H. Edwards, Essi Isohanni, Ari Korhonen, Andrew Petersen, Kelly Rivers, Miguel Ángel Rubio, Judy Sheard, Bronius Skupas, Jaime Spacco, Claudia Szabo, and Daniel Toll. 2015. Educational Data Mining and Learning Analytics in Programming: Literature Review and Case Studies. In *Proceedings of the 2015 ITiCSE on Working Group Reports (ITICSE-WGR '15).* ACM, New York, NY, USA, 41–63. https://doi.org/10.1145/2858796.2858798

[13] (ISC)^2. 2018. *Cybersecurity Professionals Focus on Developing New Skills as Workforce Gap Widens.* Technical Report. Cybersecurity Workforce Study.

[14] Ge Jin, Manghui Tu, Tae-Hoon Kim, Justin Heffron, and Jonathan White. 2018. Game Based Cybersecurity Training for High School Students. In *Proceedings of the 49th ACM Technical Symposium on Computer Science Education (SIGCSE '18).* ACM, New York, NY, USA, 68–73. https://doi.org/10.1145/3159450.3159591

[15] Association for Computing Machinery (ACM) Joint Task Force on Computing Curricula and IEEE Computer Society. 2013. *Computer Science Curricula 2013: Curriculum Guidelines for Undergraduate Degree Programs in Computer Science.* ACM, New York, NY, USA. https://doi.org/10.1145/2534860

[16] Joint Task Force on Cybersecurity Education. 2017. *Cybersecurity Curricular Guideline.* Retrieved November 25, 2019 from http://cybered.acm.org/

[17] Keith S. Jones, Akbar Siami Namin, and Miriam E. Armstrong. 2018. The Core Cyber-Defense Knowledge, Skills, and Abilities That Cybersecurity Students Should Learn in School: Results from Interviews with Cybersecurity Professionals. *ACM Trans. Comput. Educ.* 18 (2018), 11:1–11:12. https://doi.org/10.1145/3152893

[18] Niakam Kazemi and Shiva Azadegan. 2010. IPsecLite: A Tool for Teaching Security Concepts. In *Proceedings of the 41st ACM Technical Symposium on Computer Science Education (SIGCSE '10).* ACM, New York, NY, USA, 138–142. https://doi.org/10.1145/1734263.1734312

[19] Hieke Keuning, Johan Jeuring, and Bastiaan Heeren. 2018. A Systematic Literature Review of Automated Feedback Generation for Programming Exercises. *ACM Trans. Comput. Educ.* 19, 1 (Sept. 2018), 3:1–3:43. https://doi.org/10.1145/3231711

[20] Barbara Kitchenham and Stuart Charters. 2007. *Guidelines for performing Systematic Literature Reviews in Software Engineering.* Technical Report. EBSE.

[21] J Richard Landis and Gary G Koch. 1977. The measurement of observer agreement for categorical data. *Biometrics* 33, 1 (1977), 159–174.

[22] Andrew Luxton-Reilly, Simon, Ibrahim Albluwi, Brett A. Becker, Michail Giannakos, Amruth N. Kumar, Linda Ott, James Paterson, Michael James Scott, Judy Sheard, and Claudia Szabo. 2018. Introductory Programming: A Systematic Literature Review. In *Proceedings Companion of the 23rd Annual ACM Conference on Innovation and Technology in Computer Science Education (ITiCSE 2018 Companion).* ACM, New York, NY, USA, 55–106. https://doi.org/10.1145/3293881.3295779

[23] Jun Ma, Jun Tao, Jean Mayo, Ching-Kuang Shene, Melissa Keranen, and Chaoli Wang. 2016. AESvisual: A Visualization Tool for the AES Cipher. In *Proceedings of the 2016 ACM Conference on Innovation and Technology in Computer Science Education (ITiCSE '16).* ACM, New York, NY, USA, 230–235. https://doi.org/10.1145/2899415.2899425

[24] Naja A. Mack, Kevin Womack, Earl W. Huff Jr., Robert Cummings, Negus Dowling, and Kinnis Gosha. 2019. From Midshipmen to Cyber Pros: Training Minority Naval Reserve Officer Training Corp Students for Cybersecurity. In *Proceedings of the 50th ACM Technical Symposium on Computer Science Education (SIGCSE '19).* ACM, New York, NY, USA, 726–730. https://doi.org/10.1145/3287324.3287500

[25] Lauri Malmi. 2015. Can We Show an Impact? *ACM Inroads* 6, 1 (Feb. 2015), 30–31. https://doi.org/10.1145/2727129

[26] Monique Mezher and Ahmed Ibrahim. 2019. Introducing Practical SHA-1 Collisions to the Classroom. In *Proceedings of the 50th ACM Technical Symposium on Computer Science Education (SIGCSE '19).* ACM, New York, NY, USA, 879–884. https://doi.org/10.1145/3287324.3287446

[27] Allen Parrish, John Impagliazzo, Rajendra K. Raj, Henrique Santos, Muhammad Rizwan Asghar, Audun Jøsang, Teresa Pereira, and Eliana Stavrou. 2018. Global Perspectives on Cybersecurity Education for 2030: A Case for a Meta-discipline. In *Proceedings Companion of the 23rd Annual ACM Conference on Innovation and Technology in Computer Science Education (ITiCSE 2018 Companion).* ACM, New York, NY, USA, 36–54. https://doi.org/10.1145/3293881.3295778

[28] Kai Petersen, Robert Feldt, Shahid Mujtaba, and Michael Mattsson. 2008. Systematic Mapping Studies in Software Engineering. In *Proceedings of the 12th International Conference on Evaluation and Assessment in Software Engineering (EASE'08).* BCS Learning & Development Ltd., Swindon, UK, 68–77. http://dl.acm.org/citation.cfm?id=2227115.2227123

[29] Kai Petersen, Sairam Vakkalanka, and Ludwik Kuzniarz. 2015. Guidelines for conducting systematic mapping studies in software engineering: An update. *Information and Software Technology* 64 (2015), 1 – 18. https://doi.org/10.1016/j.infsof.2015.03.007

[30] Computing Research and Education Association of Australasia. 2016. *CORE.* Retrieved November 25, 2019 from http://www.core.edu.au/

[31] Khaled Salah. 2014. Harnessing the Cloud for Teaching Cybersecurity. In *Proceedings of the 45th ACM Technical Symposium on Computer Science Education (SIGCSE '14).* ACM, New York, NY, USA, 529–534. https://doi.org/10.1145/2538862.2538880

[32] Michael Skirpan, Jacqueline Cameron, and Tom Yeh. 2018. Quantified Self: An Interdisciplinary Immersive Theater Project Supporting a Collaborative Learning Environment for CS Ethics. In *Proceedings of the 49th ACM Technical Symposium on Computer Science Education (SIGCSE '18).* ACM, New York, NY, USA, 946–951. https://doi.org/10.1145/3159450.3159574

[33] Madiha Tabassum, Stacey Watson, Bill Chu, and Heather Richter Lipford. 2018. Evaluating Two Methods for Integrating Secure Programming Education. In *Proceedings of the 49th ACM Technical Symposium on Computer Science Education (SIGCSE '18).* ACM, New York, NY, USA, 390–395. https://doi.org/10.1145/3159450.3159511

[34] Blair Taylor and Siddharth Kaza. 2011. Security Injections: Modules to Help Students Remember, Understand, and Apply Secure Coding Techniques. In *Proceedings of the 16th Annual Joint Conference on Innovation and Technology in Computer Science Education (ITiCSE '11).* ACM, New York, NY, USA, 3–7. https://doi.org/10.1145/1999747.1999752

[35] Clark Taylor, Pablo Arias, Jim Klopchic, Celeste Matarazzo, and Evi Dube. 2017. CTF: State-of-the-Art and Building the Next Generation. In *2017 USENIX Workshop on Advances in Security Education (ASE 17).* USENIX Association, Vancouver, BC, 11. https://www.usenix.org/conference/ase17/workshop-program/presentation/taylor

[36] Cynthia Taylor and Saheel Sakharkar. 2019. ');DROP TABLE Textbooks;–: An Argument for SQL Injection Coverage in Database Textbooks. In *Proceedings of the 50th ACM Technical Symposium on Computer Science Education (SIGCSE '19).* ACM, New York, NY, USA, 191–197. https://doi.org/10.1145/3287324.3287429

[37] Maxim Timchenko and David Starobinski. 2015. A Simple Laboratory Environment for Real-World Offensive Security Education. In *Proceedings of the 46th ACM Technical Symposium on Computer Science Education (SIGCSE '15).* ACM, New York, NY, USA, 657–662. https://doi.org/10.1145/2676723.2677225

[38] Giovanni Vigna, Kevin Borgolte, Jacopo Corbetta, Adam Doupé, Yanick Fratantonio, Luca Invernizzi, Dhilung Kirat, and Yan Shoshitaishvili. 2014. Ten Years of iCTF: The Good, The Bad, and The Ugly. In *2014 USENIX Summit on Gaming, Games, and Gamification in Security Education (3GSE 14).* USENIX Association, San Diego, CA, 7. https://www.usenix.org/conference/3gse14/summit-program/presentation/vigna

[39] James Walker, Man Wang, Steven Carr, Jean Mayo, and Ching-Kuang Shene. 2019. Teaching Integer Security Using Simple Visualizations. In *Proceedings of the 2019 ACM Conference on Innovation and Technology in Computer Science Education (ITiCSE '19).* ACM, New York, NY, USA, 513–519. https://doi.org/10.1145/3304221.3319760

[40] Man Wang, Steve Carr, Jean Mayo, Ching-Kuang Shene, and Chaoli Wang. 2014. MLSvisual: A Visualization Tool for Teaching Access Control Using Multi-level Security. In *Proceedings of the 2014 Conference on Innovation and Technology in Computer Science Education (ITiCSE '14).* ACM, New York, NY, USA, 93–98. https://doi.org/10.1145/2591708.2591730

[41] Man Wang, Jean Mayo, Ching-Kuang Shene, Steve Carr, and Chaoli Wang. 2017. UNIXvisual: A Visualization Tool for Teaching UNIX Permissions. In *Proceedings of the 2017 ACM Conference on Innovation and Technology in Computer Science Education (ITiCSE '17).* ACM, New York, NY, USA, 194–199. https://doi.org/10.1145/3059009.3059031

[42] Man Wang, Jean Mayo, Ching-Kuang Shene, Thomas Lake, Steve Carr, and Chaoli Wang. 2015. RBACvisual: A Visualization Tool for Teaching Access Control Using Role-based Access Control. In *Proceedings of the 2015 ACM Conference on Innovation and Technology in Computer Science Education (ITiCSE '15).* ACM, New York, NY, USA, 141–146. https://doi.org/10.1145/2729094.2742627

[43] Michael Whitney, Heather Lipford-Richter, Bill Chu, and Jun Zhu. 2015. Embedding Secure Coding Instruction into the IDE: A Field Study in an Advanced CS Course. In *Proceedings of the 46th ACM Technical Symposium on Computer Science Education (SIGCSE '15).* ACM, New York, NY, USA, 60–65. https://doi.org/10.1145/2676723.2677280

[44] Le Xu, Dijiang Huang, and Wei-Tek Tsai. 2012. V-lab: A Cloud-based Virtual Laboratory Platform for Hands-on Networking Courses. In *Proceedings of the 17th ACM Annual Conference on Innovation and Technology in Computer Science Education (ITiCSE '12).* ACM, New York, NY, USA, 256–261. https://doi.org/10.1145/2325296.2325357

[45] Maximilian Zinkus, Oliver Curry, Marina Moore, Zachary Peterson, and Zoë J. Wood. 2019. Fakesbook: A Social Networking Platform for Teaching Security and Privacy Concepts to Secondary School Students. In *Proceedings of the 50th ACM Technical Symposium on Computer Science Education (SIGCSE '19).* ACM, New York, NY, USA, 892–898. https://doi.org/10.1145/3287324.3287486

[46] Valdemar Švábenský, Jan Vykopal, and Pavel Čeleda. 2019. *Dataset: What Are Cybersecurity Education Papers About? A Systematic Literature Review of SIGCSE and ITiCSE Conferences.* Zenodo. https://doi.org/10.5281/zenodo.3506640