Cybersecurity Laboratory Education Research: A Lush Ecosystem or Elephant Graveyard?

Jason M. Pittman
Cybersecurity and Information Technology University of Maryland Global Campus
Largo, MD USA
jason.pittman@umgc.edu

Reilly Kobbe
Computer Science
High Point University
High Point, NC USA
rkobbe@hpu.edu

Taylor Lynch
Computer Science
High Point University
High Point, NC USA
tlynch1@hpu.edu

Helen G. Barker
Cybersecurity and Information Technology University of Maryland Global Campus
Largo, MD USA
helen.barker@umgc.edu

Abstract—Is cybersecurity laboratory education research a lush ecosystem or an elephant graveyard? The value of such a question cuts to the health of a research field. Further, the health of a research field stems from the lineage of work extending into the past and present. In other words, mature and robust fields of knowledge exhibit interlinked research with dense pockets of follow-up. In contrast, nascent or limited fields lack such linking or association measurable by the frequency of new research extended results. These interlinks and associations are indeed quantifiable through the meta-study of bibliometrics. In fact, prior research discovered that only thirty percent of computer science research - a strongly related field - are extended after publication. However, no work to date has examined cybersecurity laboratory education for the same phenomenon. To that end, this work evaluated 400 articles with the goal of ascertaining to what degree three operationalized follow-up categories occur in the literature. The results indicate 62.5% of articles do not extend existing research. The conclusions and recommendations included at the end of this work offer potential insights into why cybersecurity laboratory education literature exists in such a state.

Keywords—cybersecurity education, pedagogy, laboratories, bibliometrics

I. Introduction

The idea of standing on the shoulders of giants embodies a basic scientific principle. This principle is the foundation behind how a field of knowledge forms and, more importantly, grows. The evidence of the principle is observable when analyzing published literature, chiefly using *bibliometrics*. Put another way, it is reasonable to expect the body of knowledge to grow over time by extensions related to the established research topics, problems, purposes, and results.

Yet, a general problem exists insofar as only 30% of computer science papers are extended after publication [1]. Worse still, adjacent fields such as cybersecurity have not been analyzed at all. Indeed, one conclusion from the work [1] was a lack of similar investigation into literature follow-up within other but related scientific disciplines. Failure to extend research is a missed opportunity to address unanswered questions and to reduce potential gaps in a field of study [2]. Rarely will one study represent the absolute answer to a question in any given field of knowledge. Extending existing research may result in a total reversal of what is considered *true*. Simply asking a question in a new way has the potential of changing the way problems and results are viewed.

For this reason, Wainer and Valle [1] asked pointed questions related to tracking research. Particularly relevant to this work is "How many do effectively continue [existing research]?" (p. 104). Similarly, the interrogative motivation for the present study was, *to what are cybersecurity laboratory research papers extended after publication*. Accordingly, the purpose of this research was to measure whether cybersecurity laboratory research is extended in a statistically significant manner based on bibliometric structural indicators. The scope of the population included all cybersecurity education laboratory literature spanning the randomly selected years, 1997 to 2020. The sample was built by selecting specific articles from the population to create a manageable dataset. Data were classified into categories against which comparative data analysis was performed.

Before delving into the findings, however, the following section includes summaries of related work necessary to situate the results and recommendations. Operationally, an understanding of bibliometrics is critical and thus is presented in detail. Further, how the development or evolution of a field of knowledge can be traced is outlined. Lastly, some foundation is offered to illustrate the relation between computer science and cybersecurity as a body of work.

II. Related Work

Scientific progress and the operative health of a field of study are intertwined. Because of the tight coupling between these concepts, it is possible to infer the latter by the quantitative assessment of the former. This work is situated within a robust information sciences literature related to bibliometrics [3].

Accordingly, this section provides a background for bibliometric analyses. Such background information contextualizes the method and results which follow. Furthermore, operationalizing literature follow-up as a variable first requires establishing a conceptual framework. Thus, the related work is segmented into the following sections: (a) development of a field of study; (b) bibliometric

analysis of computer science; and (c) cybersecurity education publishing and citation bibliometrics.

A. Growth of a Field

For this work, the field of cybersecurity is taken to be a concept related to computer science focused on the assurance of privacy, confidentiality, and integrity of data [4] [5]. Critically, a prominent theme in the literature is that cybersecurity is a protoscience [6] with a set of emerging properties [7] that sets it apart from other adjacent fields. Further, cybersecurity is widely recognized as a nascent field of study [5] [8].

The development of a scientific field is germane to the topic of this work. There are many ways to build a scientific map of the development of a field, one way is by using category information. Mapping this way uses the words and phrases to find specific papers. These maps reveal the evolution of scientific fields at different resolutions [9].

Moreover, the study of scientific evolution is based on philosophies: cognitive view and social view. In the cognitive view, more emphasis is put on the importance of shared knowledge. In the social view, studies offer qualitative descriptions of stages for group formation [9]. Researchers can then construct the knowledge network of science-based on journal/conference paper citations.

B. Structure of Field

Like a geographical map, the knowledge network of sciences provides insight into the structure of science [10] [4]. The structure of a knowledge network is critical to the field because it identifies major ideas as well as similarities or differences across such ideas. Further, a sustainable scientific field should exhibit a well-formed mapped research cluster [11] [12]. But the extent to which sustainability science meets the requirements for the level to be maintained must be investigated. In this way, bibliometric data is useful to investigate the empirical validity of the requirements of sustainability. Overall, a key takeaway is that more publications imply more follow-up citations and higher impact [13].

An expected consequence then is after publication the results from higher impact work are extended by other researchers and cited as references in their articles. The number of citations a particular article has reflects its impact on the scientific community. As a result, this data can be used statistically and mathematically to also measure the importance of an article, known as bibliometrics [14].

C. Bibliometrics

There are three types of bibliometric indicators: quantity indicators, performance indicators, and structural indicators [14]. In this work, *structural* indicators are identified and measured by analyzing connections between publications. Bibliometrics techniques rely on counting citations to measure the impact of the work in the scientific community. Results of this process are presented in various forms so that they establish relationships with participants and expand the means by analysis, such as mapping [15].

Naturally, if there are issues to be addressed in the structure of the cybersecurity education field itself, the time to uncover and work towards solutions would be while the field is still developing. Thus, analyzing structural indicators – specifically, citation practices – is a primary methodology to uncover possible developmental issues in a growing scientific field [3]. For that reason, the mapping of a field is vital and beneficial as it reveals current topics and nascent authors. Moreover, future scientific impact can be discovered within emergent themes, not to mention valuable challenges, perspectives, and suggestions for follow-up [5].

Finally, while bibliometrics can utilize a variety of methodologies to trace the development of a field of study, trends in research (structural) keywords are relative to publications because of their importance in researching fields [16]. According to [17], "bibliometrics is to scientific papers what epidemiology is to patients" (p. 14). Such changes reflect shifts in trends within a field. Thus, with sufficiently large sample sizes, structural elements such as keywords may reveal underlying field development [18].

Relatedly, computer science is [19] "a well-established, dynamic, and still relatively new research field that made its breakthrough only some fifty years ago" (p. 1). Nowadays, it is a prestigious interdisciplinary scientific domain having significant interconnections with mathematics, physics, and even biology. Surprisingly, even though this field has grown into something highly respected, there still have not been many bibliometric studies measuring the published research citations of computer science (Fiala, 2012). In this context, the theoretical and practical significance of work [1] [20] are useful when familiarizing oneself with concepts from the field of bibliometrics.

Meanwhile, the more specific field of cybersecurity [21] [22] [23] has not yet begun to examine its publication or citation practices as structural indicators. Given the related work in computer science and the absence of similar investigation within cybersecurity, it seems a natural next step would be to apply bibliometric analyses to cybersecurity literature, specifically cybersecurity laboratory research targeting higher education.

III. Method

This work was motivated by a single research question: to what extent are cybersecurity laboratory research papers extended after publication. To answer this question, we adhered to the overall research methodology outlined in [1]. That is, we (a) searched for relevant literature to establish a population; (b) created a sample data set of articles; (c) classified articles into one of three categories; and (d) performed comparative data analysis against those categories. We introduced limited modifications where necessary to apply the method to the field of cybersecurity education.

A. Population

Whereas [1] analyzed ACM literature exclusively, we broadened the scope of the population to include all cybersecurity education laboratory literature spanning the

years 1997 to 2020. The start date was based on the first available research [24]. The end date was simply the most recent full year of literature available during this work (2020). Collectively, we felt this modification was reasonable given that cybersecurity education laboratory literature is published across a variety of journals, conferences, and preprint archives.

Following that rationale, we constructed a set of five search strings with the intent to operationalize as many variants of the term cybersecurity as possible. The search strings included: (a) computer security AND lab OR laboratories; (b) cyber security AND lab OR laboratories; (c) information assurance AND lab OR laboratory; (d) information security AND lab OR laboratory; and (e) network security AND lab OR laboratory. The resulting searches were performed in Google Scholar.

B. Sampling

Rather than building a sample by selecting specific articles from the population, we created a tenable dataset for classification and analysis by randomly selecting four years from the population date range. Doing so ensured a completed set of original articles would be included, thus imparting rigor and integrity to the follow-up article links present in the population. Accordingly, we employed a simple random selection formula against a spreadsheet column consisting of rows inclusive of 1997 to 2020.

To that end, the formula selected 1997, 2001, 2014, and 2018 as target years for the sample. We reviewed the articles' abstracts for keywords aligned with the topic of this work to exclude irrelevant or unrelated literature. The final sample totaled 400 articles. 73 were found in 1997, 96 were found in 2001, 72 were found in 2014, and 159 were found in 2018.

C. Classification

We recorded the sample into a spreadsheet with sample years on separate worksheets. Each worksheet had five columns representing the search strings. These corresponded to the source *original* classification [1] and were coded as *A columns*. Additionally, we recorded the primary author's last name, publication year, and a number to indicate if the author's last name previously appeared.

Two additional columns represented the concept [1] of *follow-up*. These were coded as *B columns*. One of the B columns captured the primary author's last name and year published of all the articles which cited the original the corresponding A article. The other B column denoted which follow-up category the papers fit in: state of the art, extensions, or republication. We extended the classification model with a fourth category to denote articles that were not accessible.

Briefly, *state of the art* articles consisted of work which cited an original article as part of the background but otherwise does not continue the line of inquiry [1]. In contrast, work denoted as an *extension* grew the ideas in the original study or used results as foundational elements. Thirdly, *republications* represented a more-or-less duplication of the associated original. Then, finally, we formalized the categorization of *not accessible* for those studies that did not have at least an obtainable abstract.

D. Analysis

A two-phase analysis followed the classification of the sample literature. First, descriptive analysis revealed features and characteristics of the data to illustrate the extent to which cybersecurity education literature exists across the follow-up spectrum. Second, inferential analysis examined properties meta to the data and adhered to the protocol [1] used to estimate population parameters using multinomial confidence intervals.

With that said, this work deviated from the protocol insofar as this work does not differentiate between *conference* and *journal*. The deviation was necessary because the nature of the initial literature search, and thus population, encompassed a broad spectrum of literature sources as opposed to the single ACM literature source.

IV. Results

Data were drawn from four years selected randomly out of the total population space of 24 years (1997 to 2020). Fig. 1 demonstrates the stratification of 400 follow-up articles across the four-year sample. Reading clockwise from 1997, the initial three years in the sample demonstrated consistent breadth. However, the last year (2018) had roughly double the expanse with respect to the number of follow-ups.

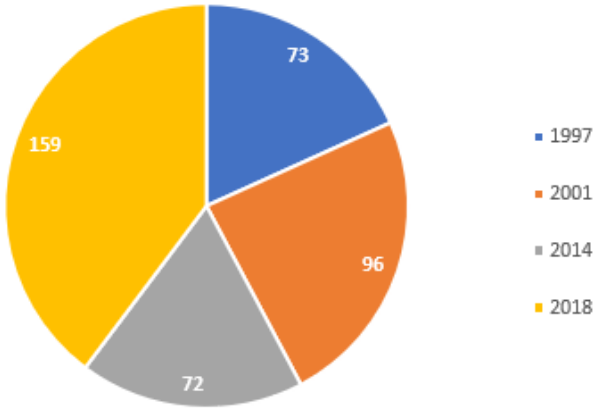

Fig. 1. The proportion of follow-ups across the four sample years.

The sample can be further described according to the foundational original studies grouped by search string. In doing so, original literature can be grouped according to related keywords which may establish a foundation for future analysis. To ease such descriptive analysis, Table I outlines a search string coding mechanism.

TABLE I. CODING FOR LITERATURE SEARCH STRINGS

Code	Search Strings
S1	”computer security” AND ”lab” OR ”laboratory”
S2	”cyber security” AND ”lab” OR ”laboratory”
S3	”information assurance” AND ”lab” OR ”laboratory”
S4	”information security” AND ”lab” OR ”laboratory”
S5	”network security” AND ”lab” OR ”laboratory”

Regarding the frequency of original work compared to the number of follow-ups, for the year of 1997, there were eight unique original articles. The eight originals came from just two of the searches; seven from S1 and one from S3. The remaining years contained duplicate results from the two operative searches. Likewise, for the year 2001 there were also eight unique original articles. In contrast, these eight originated from the S1, S3, S5 search strings. The S2 did not include any original studies and S4 only contained duplicate originals present in other years. The sample year 2014 contained 12 original works. These 12 studies were found using the S1, S2, S3, and S5 search strings. The last sample year of 2018 yielded 72 original studies which represented a significant increase from the prior three years. Further, each of the search strings contributed to this collection with 33, 24, two, four, and nine, respectively. Table II summarizes the comparison of these descriptive values against the total number of follow-ups for each sampled year.

TABLE II. FREQUENCY OF ORIGINALS BY SEARCH STRING

Year	Frequency					Follow-ups
1997	**S1**	**S2**	**S3**	**S4**	**S5**	73
	7	0	1	0	0	
2001	**S1**	**S2**	**S3**	**S4**	**S5**	96
	4	N	1	0	3	
2014	**S1**	**S2**	**S3**	**S4**	**S5**	72
	6	2	3	0	1	
2018	**S1**	**S2**	**S3**	**S4**	**S5**	159
	33	24	2	4	9	

Note: N represents null results in the search and 0 represents zero unique

A. Follow-up

Analysis of the 400 follow-up articles in the sample demonstrated that 250 represented state of the art research (62.5%), 79 (19.7%) in the extension category, two in the republication category (0.005%), 51 (12.7%) were not extended, and 18 (0.045%) were not accessible. Table III summarizes the frequencies underlying how these values breakdown by sample year in each follow-up category.

TABLE III. FREQUENCY OF CATEGORIZED FOLLOW-UPS BY SAMPLE YEAR

	1997	2001	2014	2018
N	73	96	72	159
state of the art	44	69	51	86
extension	25	21	14	19
republication	0	0	0	2
not extended	0	0	0	51
not accessible	4	6	7	1

More specifically, the year 1997 yielded a total of seventy-three follow-ups. Forty-four of those articles were state of the art (60.2%), twenty-five of them extended (34.2%) on the authors' previous idea. However, in the year of 1997, there were zero republications, and zero articles were not extended. Four articles were also not accessible. In 2001, there was a total of ninety-six articles which proved to be relevant. Of those, sixty-nine were state of the art (71.8%), twenty-one were extensions (21.8%), zero were republicans, and zero were not extended. There were also five articles which were not accessible. Similar to the two prior sample years, 2014 had a total of seventy-two articles. Fifty-one of those were state of the art (70.8%), fourteen extended the original idea (19.4%), zero were republications, and zero were not extended. For the same stratum, seven articles were not accessible. Lastly, in the year 2018 the total of articles was 159. Of those, 86 were state of the art (54%), nineteen were extensions (11.9%), two republicans, fifty-one were not extended (32%) and one article was not accessible.

Supporting the sample descriptive analysis, we inferentially measured to what extent the described parameters extrapolate to the broader population. In adhering to the source analyses [1], Table IV reveals a series of confidence intervals for proportions of categorized follow-up literature in the cybersecurity (education) laboratory domain.

TABLE IV. CONFIDENCE INTERVALS FOR PROPORTIONS OF CATEGORIZED FOLLOW-UPS

Confidence Intervals					
	Count	Proportion	Low CI	Upper CI	Alpha
state of the art	250	0.62500	0.5852	0.6648	0.10

Confidence Intervals					
	Count	Proportion	Low CI	Upper CI	Alpha
extension	79	0.19750	0.1648	0.2302	0.10
republication	2	0.00500	0	0.0108	0.10
not extended	51	0.12750	0.1001	0.1549	0.10
not accessible	18	0.04500	0.028	0.062	0.10

V. Conclusions

This work examined to what extent cybersecurity (education) laboratory research generated a follow-up study. In doing so, the intent was to uncover features and characteristics associated with the overarching maturity of the field of knowledge. While this research did not seek to evaluate the quality of researcher recommendations it is important to note that there may be issues such as citation practices, the quality of recommendations, and a failure to establish a standard recommendation model that contributes to the lack of extension of existing research [2].

Overall, having less than 25% of the field's findings extended reveals a lack of proper development in the cybersecurity laboratory education literature field and the lack of understanding of the importance of extending on previous studies. Conjointly, more than half of the articles in this sample have never been replicated or furthered with a proper scientific mechanism. The outcome is a sizable hole; without the extension of articles, the field's future growth will be problematic. This is especially true when more than half of the articles studied fall into the state of the art category. A total of 12.7% of articles did not have proper citations or follow-up citations. These articles had ideas but no proof of where they originated from. Four percent could not be accessed. These findings raise attention on the issue of properly extending research, halting the growth of the field. By bringing attention to these issues, we hope to influence an increase in the rigor of scientific follow-up.

To that end, like the conclusions offered by [1], the data show the body of cybersecurity laboratory research contains roughly one-fifth (19.7%) extension follow-up. In contrast, a significant majority of work (62.5%) uses existing research for conceptual or theoretical context only. In comparison, the source study [1] found between 23% and 30% of research received an extension. This work exhibits a lower state of the art proportion comparably but cybersecurity laboratory education, in this research, also has a higher incidence of inaccessibility and indeterminable citation. In this context, many if not all the same conclusions can be reached here as were inferred in the source study. Yet, in the continued spirit of research extension, the following additional conclusions are offered.

Because of stratifying the sample by year, an interesting find is observable in the frequency of follow-ups. The count is largely stable across the initial three strata. Specifically, it starts at 73 in 1997, rises to 96 in 2001, and then lessens back down to practically the same number in 2014 as begun with. However, the frequency of research follow-up exploded to 159 in 2018. Yet, there is no appreciable increase in extension; only state of the art and not extended follow-ups.

The explosion in research from 2018 with a concomitant increase in work not extended and a lower frequency of extension may be indicative of an overall trend. That is, while more work is being produced, such literature is not rooted in or based on existing research. Likewise, there is a conspicuous absence of republications. Are these effects related to a cause in the research itself or something external?

For example, externally the Y2K bug that was crippling telecommunications going into 2000 or the rise of hacktivism in 2004 may have been social issues driving diverse, linearly disconnected research pursuits. In this way, reasonable future work may be fruitful when examining the potential correlations between popular cultural cybersecurity events and the research literature. More concretely, perhaps a means to quantify such speculation rests in determining to what extent cybersecurity laboratory education research constitutes action research versus empirical research based on analysis of published research methods.

In the vein of this research itself serving as a causative agent, this work points towards two possible alternative explanations for the lack of direct follow-up in cybersecurity laboratory literature. First, existing bibliometric research (including this study) assumes work-in-progress is able to find existing literature. However, the metadata associated with published research (i.e., PDF attributes) or keywords contained in the articles may be flawed. Likewise, the quality of search strings is vital to cultivating relevant literature datasets. Second, there is a potentially weak assumption insofar as the published literature is readable. Readability might be related to the grammar and semantics of the work or readability may apply to the requisite technical knowledge to meaningfully synthesize results and findings.

Future exploration into these ideas as possible factors is suggested. To be sure, the field of cybersecurity laboratory education research is expanding. At the same time, the findings suggest the field is nascent and growing laterally just as much as it is vertically. Thus, there is no better time to apply serious bibliometric analyses to the literature than now.

Likewise, it is notable that future work towards understanding why those works categorized as extensions are not more visible within the cybersecurity laboratory education literature. Research into metadata standards, journal and conference indexing practices, as well as author self-promotion through social media may be of benefit in this context.

REFERENCES

[1] J. Wainer and E. Valle, "What happens to computer science research after it is published? tracking cs research lines," *Journal of the American Society for Information Science and Technology*, vol. 64, no. 6, pp. 1104–1111, 2013.

[2] P. Brown, K. Brunnhuber, K. Chalkidou, I. Chalmers, M. Clarke, M. Fenton, C. Forbes, J. Glanville, N. J. Hicks, J. Moody et al., "How to formulate research recommendations," *Bmj*, vol. 333, no. 7572, pp. 804–806, 2006.

[3] W. Glanzel, *Bibliometrics as a research field a course on theory and application of bibliometric indicators*, 2003.

[4] M. C. Pham, M. Derntl, and R. Klamma, "Development patterns of scientific communities in technology enhanced learning," *Journal of Educational Technology & Society*, vol. 15, no. 3, pp. 323–335, 2012.

[5] L. B. Furstenau, M. K. Sott, A. J. O. Homrich, L. M. Kipper, A. A. Al Abri, T. F. Cardoso, J. R. Lopez-Robles, and M. J. Cobo, "20 years of scientific evolution of cyber security: A science mapping," in *International Conference on Industrial Engineering and Operations Management*. IEOM Society International, 2020, pp. 314–325.

[6] E. N. Hatleback, "The protoscience of cybersecurity," *The Journal of Defense Modeling and Simulation*, vol. 15, no. 1, pp. 5–12, 2018.

[7] A. Roque, "Validating computer security methods: Meta-methodology for an adversarial science," *arXiv preprint arXiv:1710.01367*, 2017.

[8] K. M. Roberts, "Addressing current and future resource deficiencies within the field of cybersecurity: A generic qualitative inquiry," Ph.D. dissertation, Capella University, 2020.

[9] X. Sun, K. Ding, and Y. Lin, "Mapping the evolution of scientific fields based on cross-field authors," *Journal of Informetrics*, vol. 10, no. 3, pp. 750–761, 2016.

[10] M. C. Pham, R. Klamma, and M. Jarke, "Development of computer science disciplines: a social network analysis approach," Social Network *Analysis and Mining*, vol. 1, no. 4, pp. 321–340, 2011.

[11] S. Wendzel, C. Levy-Bencheton, and L. Caviglione, "Not all areas are equal: analysis of citations in information security research," *Scientometrics*, vol. 122, no. 1, pp. 267–286, 2020.

[12] C. Cancino, J. M. Merigo´, F. Coronado, Y. Dessouky, and M. Dessouky, "Forty years of computers & industrial engineering: A bibliometric analysis," *Computers & Industrial Engineering*, vol. 113, pp. 614–629, 2017.

[13] E. D. Schoolman, J. S. Guest, K. F. Bush, and A. R. Bell, "How interdisciplinary is sustainability research? analyzing the structure of an emerging scientific field," *Sustainability Science*, vol. 7, no. 1, pp. 67–80, 2012.

[14] V. Durieux and P. A. Gevenois, "Bibliometric indicators: quality measurements of scientific publication," *Radiology*, vol. 255, no. 2, pp. 342– 351, 2010.

[15] A. A. El-Maamiry and M. A. Ghauri, "Measuring information quality: Concerns on the use of bibliometric studies," *International Journal of Information Dissemination and Technology*, vol. 3, no. 4, pp. 274–278, 2013.

[16] M. d. C. Gimenez-Espert and V. J. Prado-Gasco, "Bibliometric analysis of six nursing journals from the web of science, 2012–2017," *Journal of advanced nursing*, vol. 75, no. 3, pp. 543–554, 2019.

[17] G. Lewison and M. Devey, "Bibliometric methods for the evaluation of arthritis research." *Rheumatology (Oxford, England)*, vol. 38, no. 1, pp. 13–20, 1999.

[18] Y. Song, X. Chen, T. Hao, Z. Liu, and Z. Lan, "Exploring two decades of research on classroom dialogue by using bibliometric analysis," *Computers & Education*, vol. 137, pp. 12–31, 2019.

[19] D. Fiala and G. Tutoky, "Computer science papers in web of science: A bibliometric analysis," *Publications*, vol. 5, no. 4, p. 23, 2017.

[20] O. Mubin, M. Arsalan, and A. Al Mahmud, "Tracking the follow-up of work in progress papers," *Scientometrics*, vol. 114, no. 3, pp. 1159– 1174, 2018.

[21] S. L. Vrhovec, D. Fujs, L. Jelovcan, and A. Mihelic, "Evaluating case study and action research reports: Real-world research in cybersecurity." *J. Univers. Comput. Sci.*, vol. 26, no. 7, pp. 827–853, 2020.

[22] D. Fujs, S. L. Vrhovec, and D. Vavpotic, "Bibliometric mapping of research on user training for secure use of information systems." *J. Univers. Comput. Sci.*, vol. 26, no. 7, pp. 764–782, 2020.

[23] R. Heradio, L. De La Torre, D. Galan, F. J. Cabrerizo, E. Herrera-Viedma, and S. Dormido, "Virtual and remote labs in education: A bibliometric analysis," *Computers & Education*, vol. 98, pp. 14–38, 2016.

[24] C. E. Irvine, D. F. Warren, and P. C. Clark, "The nps cisr graduate program in infosec: Six years of experience," NAVAL POSTGRADUATE SCHOOL MONTEREY CA DEPT OF COMPUTER SCIENCE, Tech. Rep., 1997.