

Consensus Building Method for Expert Crowdsourcing of Curriculum Topics

Mr. Brian Khoa Ngac, George Mason University

Brian Ngac is Deputy to the Vice President of Digital Engineering Research & Development Programs at Parsons Corporation's Defense & Intelligence Unit, and a PhD Candidate (ABD) at George Mason University's College of Engineering & Computing. He holds 12 internationally recognized cyber security and management certifications including the C|CISO, CISSP, ISSMP, CISM, and PMP. His areas of expertise are in cyber security, digital engineering (RDT&E), and business process improvement (solving business challenges with technology solutions). His research focus are in cyber executive management, expert crowdsourcing, and decision analytics.

Dr. Mihai Boicu, George Mason University

Mihai Boicu, Ph.D., is Assistant Professor of Information Technology at George Mason University. He published over 120 peer-reviewed publications, including 4 books. He performs theoretical and applied research in Artificial Intelligence, Machine Learning, Probabilistic Reasoning, Crowdsourcing and Engineering Education. He received more than 3M in funding from NSF, DARPA, IARPA, AFOSR, IC and other government agencies.

Title: Consensus Building Method for Expert Crowdsourcing of Curriculum Topics

Abstract

State of the art curriculum development efforts are done with a committee often consisting of two to four faculty members but are commonly undertaken by the assigned course instructor. However, the small number of faculty participants in the curriculum development effort can yield an out-of-date and insufficient curriculum for students entering the industry workforce [1], [2], [3], [4]. Crowdsourcing has been used to gather more input from domain experts consisting of faculty and industry professionals [2], [3], [5]. However, these efforts can yield large amounts of inputs from various crowd workers resulting in additional time required for the curriculum owner or committee to sort through all inputs, organize them into categories, identify similarities and determine which topics to include in the curriculum based on the crowd's consensus [6]. In a previous experiment performed, we obtained 1) that crowdsourcing efforts can be effective at gathering inputs which both confirm or reject existing topics and yield additional topics to be included; 2) continuing and dynamic expert crowdsourcing is helpful in building consensus of newly suggested topics for the curriculum; and 3) manual aggregation of crowdsourced inputs is an inefficient process [6]. This paper proposes the integration of a consensus building function into a crowdsourcing platform [7] to dynamically and automatically validate and integrate both structured inputs (topics with confirmed representation) and semi-structured inputs (newly suggested curriculum topics) from the expert crowd when developing a course curriculum. The topics are organized in an ontology, having a hierarchical relationship "subtopic of" and an associated consensus value representing the agreement between the crowd participants about the inclusion of the topic (and its representation). When a worker from the expert crowd is providing feedback about an existing topic, the consensus value for that topic will be updated (increasing or decreasing the agreement). To ensure confidence in the consensus value a minimal number of inputs is required, along with a named threshold confidence (which is established based on the number of available workers and the expected working time). A topic may be in five different states: accepted, rejected, tentatively accepted, tentatively rejected, and undetermined. While the crowd process focuses on the topics that are undetermined or have not been accepted or rejected, the consensus value continues to be updated and it is possible to have further refinement both of accepted and rejected topics. The process continues until consensus is reached on all topics. In this paper we describe an experiment performed using a dynamic crowdsourcing platform with consensus building function for a new cybersecurity curriculum and compare it with a previous experiment in which we use a typical crowdsourcing platform for another cybersecurity curriculum. The current cybersecurity curriculum focuses on the topic of managing network and data security, while the previous curriculum focuses on the topic of managing information security with vendors and partners.

Introduction

Currently, curriculum development efforts at many higher educational institutions are done with a small committee of faculty members or, more commonly, completed by the course instructor. Because of the small number of faculty participants, the curriculum development effort can yield an out-of-date and insufficient curriculum for students entering the industry workforce [1], [3], [4]. This not only impacts the students' ability to be competitive in the workforce, but also negatively impacts an already understaffed industry where there is a skills-gap [8].

The activity of crowdsourcing for curriculum development efforts has been used to gather more inputs from domain experts consisting of faculty and industry professionals [2], [3], [5]. Naturally, crowdsourcing efforts often results in yielding large amounts of inputs from various crowd workers. This in turn forces the curriculum owner or committee to manually filter through all the inputs from the crowd and determine which should be included in the curriculum [6]. This manual process is long and laborious because different crowd members will suggest different items and directions for the curriculum, so determining what a *correct* curriculum based on the consensus of the crowd can be a daunting task.

Previous manual crowdsourcing experiment

Recently, we performed a crowdsourcing experiment for curriculum development of a graduate course titled "Managing Information Security with Vendors & Partners" which utilized the Curriculum Development using Crowdsourcing Framework (CDC-F) [6]. Using CDC-F, the experiment ran through two rounds of crowdsourcing consisting of a baseline curriculum developed by the course's owner; 19 cyber security experts across the commercial, government, and academic sectors; and yielded three unique results. The results included: 1 – one additional curriculum topic was added; 2 – two of the curriculum topics were merged into one; and 3 – three new subtopics were formalized.

While the crowdsourcing effort was automated in terms of inputs from the crowd and the results were positive, the analysis of the crowdsourced inputs was manual, slow, and painful. The confirmation of whether a topic or subtopic should be included was simple due to the closed-end format of the question. The suggestion of topics and subtopics request was in an open-ended format. This resulted in a large amount of semi-structure responses from the crowd. In round one, there were 72 suggested topics and in round two, there were over 80 new subtopic suggestions from the crowd.

To improve this experiment, in this paper the researchers decided to do a follow-on with a new course that would integrate the automation of analyzing the crowdsourced inputs to build consensus in determining new topics and subtopics for the curriculum.

Building consensus on the crowd

The consensus building method is a process to aid the curriculum owner in achieving agreement among the participants efficiently. This contribution is important in a variety of ways including 1 – offering a process in achieving consensus asynchronously; 2 – reducing the probability for a

dominant participant in the effort to influence or “strong-arm” other participants into leaning a particular direction; and 3 – the democratic voting nature of the process objectifies all topics, subtopics, and inputs. There were two main types of operations performed by the crowd, and the consensus building function was applied to both: 1 – validating the existing list of topics and subtopics, and 2 – suggesting additional topics and subtopics to be added to the current curriculum.

Operation 1: Validating the existing list of topics and subtopics. The goal of this process was to validate a hierarchical organization of topics (topics having subtopics, which may have their own sub-subtopics). This hierarchy of topics contains both topics proposed by the curriculum owner or dynamically accepted from the crowd. The method aims to build consensus toward which topics and subtopics to be considered and which is their relative importance. The main problem is that the list may be quite large, and to maintain expert engagement, it will not be effective and accurate to present the entire list to all members of the crowd. Rather, the method selects the parts that will have the best chance to clarify the consensus on the proposed list. To do so, we developed and evaluated this consensus measure which aggregates the votes and applies a weighted sum. The measure also has a confidence level based on the number of respondents. This algorithm is explained in the consensus building method section below. When a worker from the expert crowd is providing feedback about an existing topic, the consensus value for that topic will be updated (increasing or decreasing the agreement). To ensure confidence in the consensus value a minimal number of inputs is required, named threshold confidence, (which is established based on the number of available workers and the expected working time). While the crowd process focuses on the topics under review, the consensus value continues to be updated and it is possible to have further refinement both of accepted and rejected topics. The process continues until consensus is reached on all topics. The process with a crowd participant is as follows: 1 – Identify the topics with low consensus or low confidence for each category; 2 – Create a contextual presentation of the topics (top-down starting with the main topics and going into the more specific subtopics) that will maximize the number of identified topics presented while minimizing the number of contexts; and 3 – Ask the participant to validate the identified topics (while offering the possibility to validate the other suggested topics presented as well).

A topic and subtopic may be in four main states: 1 – undecided (or under review after it is proposed, and before reaching consensus); 2 – accepted (or tentatively accepted, if the consensus value is above the acceptance threshold with the required confidence); 3 – rejected (or tentatively rejected, if the consensus value is below the rejection threshold with the required confidence) and 4 – undetermined (if there is no consensus after enough votes were casted). When consensus was reached (i.e., the measure was over a given threshold), the topic will be marked with the consensus state and no longer proposed as a crowd task. A validated topic will continue to be included in the context presented, but the new participants are no longer asked to validate it (however, they are still able to accept or reject the topic). A rejected topic was not included in the context presented, but the participants were able to see it when they tried to add new topics and subtopics, as described in the next process.

Operation 2: Suggesting additional topics and subtopics to be added to the current curriculum. While the validation of the predefined curriculum was an important operation, the capability to add additional topics and subtopics suggested by the expert crowd is also necessary. However,

there are many issues that may arise when the crowd proposes new topics or subtopics. One example is that the new topic proposed allows for open-ended interpretation which can result in misunderstandings and incorrect voting. Another example is that new input suggestions may be synonym-equivalents of current topics and subtopics, or previously suggested ones from other respondents which can result in unnecessary and unintentional duplications. And finally, responses provided may be more detailed (or less detailed) than requested (e.g., a sub-subtopic is given versus a subtopic). To address such issues, we extended the current elicitation method in our dynamic crowdsourcing platform with the following validation process: 1 – when presented with a list of topics and subtopics, a crowd user indicates their desire to add a new suggestion; 2 – if they so desire, the user is asked to briefly formulate the new topic; 3 – the user is then presented with the current proposed topics that are in the validation process and must identify one of the following situations: the proposed topic is (a) better reformulation of another proposed topic; (b) a subtopic of another proposed topic; (c) a super-topic of another proposed topic; (d) overlapped (related) with another proposed topic; or (e) a new topic (not related with the previous proposed topics); 4 – Other expert crowd participants are then asked to confirm the categorization of the proposed topic (similar with the options in step 3); 5 – If a crowd acceptance level was obtained, the topic is proposed for global validation; If not, the topic is maintained in the proposed topics list but not included in the global validation list.

Consensus building method

We created a dynamic crowdsourcing method to elicit curriculum elements using a collective intelligence approach. We describe below the method applied to the identification of a new topic to be included in the course. While one expert is presented with a list of subtopics for a given topic, there is the option to propose a new subtopic. Once a new subtopic is created, a crowdsourcing task is added for the validation of this new subtopic. The same process is done for the initial curriculum proposed by the curriculum owner. The new task will be proposed to all experts investigating that portion of the curriculum. There are three possible answers for this new task: accept the topic, reject the topic, or skip the task. Before we explain the method, we will introduce a few constant values used.

The following threshold values were identified and experimentally tested with respect to the number of analyses to be performed by the experts. These numbers were determined assuming a crowd between 20-40 experts. The numbers will be adapted for a smaller or larger crowd following a logarithmic growth.

- MinNA: Minimum (Required) number of analyses = 5
- RecNA: Recommended number of analyses = 10

Based on the number of analyses the following types of decision will be taken:

NA = Number of analyses	Type of decision
$1 \leq NA < \text{MinNA}$	No decision will be taken
$\text{MinNA} \leq NA < \text{RecNA}$	High confidence decisions may be taken
$\text{RecNA} \leq NA$	Low confidence decisions may be taken

Table 1. Type of decision to be taken based on number of crowd analyses

The following thresholds for acceptance rates (defined as the number of answers accepting the suggestion per the number of answers accepting or rejecting the suggestion) were identified and experimentally tested with respect to the task classification:

- AHC = Accept with high confidence = 80%
- ALC = Accept with low confidence = 60%
- RLC = Reject with low confidence = 40%
- RHC = Reject with high confidence = 20%

Based on the answers received the following classifications of a crowd task are possible:

Number of answers (NA = accept+reject)	Acceptance rate = AR = accept / (accept+reject)	Status
NA < MinNA		undecided
MinNA <= NA	AR >= AHC	accepted
	AR <= RHC	rejected
MinNA <= NA < RecNA	RHC < AR < AHC	undecided
RecNA <= NA	AR >= ALC	tentatively accepted
	AR <= RLC	tentatively rejected
	RLC < AR < ALC	undetermined

Table 2. Crowd task status based on number of answers received and task acceptance rate

The next table presents the interpretation of the previously identified crowd task states with their rationale:

<u>Topic/Subtopic</u> <u>Suggestion Status</u>	<u>Description</u>
Accepted	<i>We received enough answers, and the vast majority of experts accepted the topic.</i>
Rejected	<i>We received enough answers, and the vast majority of experts rejected the topic.</i>
Tentatively Accepted	<i>We received many answers, and the majority of experts accepted the topic.</i>
Tentatively Rejected	<i>We received many answers, and the majority of experts rejected the topic.</i>
Undetermined	<i>We received many answers, and there is no consensus between experts.</i>

Table 3. Interpretation of crowd task status

The equal distribution of rate based on confidence is appropriate for typical crowd tasks. These thresholds were set based on a previous analysis when performing the previous crowdsourcing experiment, where a 60% (a three-fifths supermajority) agreement minimum was needed for confirmation of a topic or subtopic [6]. However, they may lead to topic inflation (many new topics to be added). To mitigate such situation, we plan to test more conservative rates (e.g., 85%, 70%, 30%, 15%). Of course, more experiments will need to be performed to refine and validate the above thresholds for a general situation.

When validating a topic, we developed and utilized the following consensus building algorithm:

Global Data Structures: Crowd Task: {Description, Status, Exposure, Accepted, Rejected} Open tasks: a hierarchical structure of crowd tasks sorted based on priority
Propose new crowd task (description) • <i>crowd task</i> = {description, status=undecided, exposure=1, accepted=1, rejected=0} • Add <i>crowd task</i> to <i>open tasks</i>
Dynamic Crowd Control (expert <i>e</i>): • Initialize expert open tasks with global open tasks • For each crowd task <i>t</i> in <i>expert open tasks</i> ○ If <i>t</i> has status undecided ▪ Elicit decision <i>d</i> from expert <i>e</i> ▪ Update task <i>t</i> for expert decision <i>d</i>
Update task <i>t</i> for expert decision: • If <i>d</i> = skip: increment <i>t.exposure</i>; return • If <i>d</i> = accept: increment <i>t.accepted</i> and <i>t.exposure</i> • Else: increment <i>t.rejected</i> and <i>t.exposure</i> • Update task <i>t</i> status (based on Table 2 above)

Table 4. Dynamic crowdsourcing method

When a new topic (*t1* for example) is proposed by the curriculum owner or expert, the topic will have a new Crowd Task (CT1 for example) created for validation of the topic. When this occurs,

the following variables are created and assigned a value: Task exposure = 0; Task analysis = 0; Topic accepted = 1; and Topic rejected = 0. The expert crowd is then called on to validate CT1. The task will be added to the expert open tasks and each time an expert will log to perform analysis will be added to the expert's open tasks in a priority hierarchical order. Then if assigned with the task, an expert will be able to either accept, reject or skip the CT1. If they accept, then the Topic Accepted variable is increased by one. If they decline, then the Topic Rejected variable is increased by one. Either decision will also increase the Topic Analysis and Task Exposure variables each by one. The skip will just increase the task exposure. The status of the task will be updated as described in Table 2. Therefore, the task may or may not be further prioritized for the experts, depending on its new status.

Design of experiment

The goal of this research is to evaluate the use of a dynamic crowdsourcing platform with our consensus building method, in order to create a curriculum outline that is industry-relevant while maintaining academic rigor for a newly approved graduate course in the Information Security Management program titled ISM 650 – Managing Network and Data Security.

The experiment has three distinct phases. Phase one focused on creation of a baseline curriculum using the current practice, being performed by the curriculum owner who is usually the individual course instructor or course coordinator. Phase two focused on utilizing the current state of the art curriculum development method, the curriculum being developed by two to four faculty members in a committee setting, starting with the initial baseline created in phase one. Phase three focused on building asynchronous consensus using independent experts through a dynamic crowdsourcing effort.

To begin each phase of the experiment, the ISM 650 course title, description, prerequisites, corequisite, and duration of the course were presented to the participants for contextual purposes. These documents are available in Appendix A and will be referred as *Course Requirements*.

The baseline curriculum

The initial baseline curriculum was developed by the curriculum owner – who is usually the course's primary instructor. In this case, the curriculum owner was a Term Instructor and Dean's Teaching Fellow who currently teaches a graduate executive cyber security course titled TECM 747 – Information Assurance & Security Management. To create the initial baseline curriculum, the curriculum owner reviewed the approved course's description and attributes presented in Appendix A. From there, the curriculum owner began outlining the curriculum based on their current knowledge of the subject matter. This was followed by performing additional research using authoritative resources such as industry guidelines on best practices, emerging technology watchlists, informational articles from reputable security organizations and companies, and articles on relevant cyber security compromises. The curriculum owner also looked at the prerequisites and corequisite courses mentioned in Appendix A to ensure minimum duplication of content. The effort by the curriculum owner took about five hours in total and resulted in the baseline curriculum outline for ISM 650, which consisted of eight distinct topics along with a

total of 44 subtopics. The resulted outline is shown in Appendix B and will be referred as *Baseline Curriculum*.

The academic committee's curriculum

An academic committee was then formed specifically to create a curriculum outline for ISM 650. The committee's composition included three faculty from the department offering the new course. The first was the curriculum owner who developed the baseline curriculum. The second was a Tenured Full Professor, Dean's Scholar, and former Associate Dean who currently teaches an undergraduate course titled MIS 420 – Information Security & Assurance. And the third was a Term Assistant Professor who currently teaches an undergraduate course titled MIS 320 – Networks and Security. MIS 320 is also one of the required prerequisite courses of the ISM 650 course.

Starting from the curriculum owner's initial baseline curriculum, the academic committee was instructed to propose its own version of the curriculum using Appendix A as a starting point, and Appendix B as an optional resource. The committee chose to use Appendix A and B as a starting point and started with an initial synchronous meeting over Microsoft TEAMS to set goals and expectations; share initial ideas about the curriculum and its potential; and ended with setting a meeting to finalize their version of the ISM 650 curriculum outline. A Microsoft TEAMS site was created to allow for an asynchronous working collaborative effort on the curriculum by the committee members for a week. The initial goal was to give the committee a week before meeting again, but actually met an additional week later synchronously through Microsoft TEAMS after rescheduling due to some unexpected conflicts. During the second meeting, the academic committee reviewed their curriculum outline of topics and subtopics together. There were some clarifications and minor edits before the academic committee determined they were satisfied with their version of the ISM 650 curriculum outline – which is shown in Appendix C – and adjourned. When comparing the academic committee's curriculum with the baseline curriculum, the summarized changes are as follows: 0 new topics were added; 2 out of 8 topics were revised; 4 out of the 44 subtopics were revised; 6 new subtopics were added (44 originally, 47 now); 1 subtopic was removed; and 2 subtopics were moved/combined. Overall, the academic committee had two meetings at one hour each, spent a total of five hours combined of additional time to revise the curriculum, and the whole process took about two weeks in duration.

Crowdsourcing experts with the consensus building function

In the third phase, a crowd of experts was leveraged on the professional social media platform, LinkedIn, to develop phase three's curriculum outline of topics and subtopics for ISM 650. To target cyber security experts on LinkedIn, the researchers used LinkedIn's Campaign Ad feature which allows for target-audience specification through defining location and current job position parameters. This feature allowed the researchers to request participations from a specific group of experts who are currently working in the United States and who are in the positions of security engineer, senior information security analyst, information security specialist, network security engineer, data security engineer, network security administrator, data security administrator, network security analyst, head of information security, data security manager, information security manager, cyber security manager, director of cyber security, chief information security

officer, network security manager, director of information security, vice president of information security, information technology professor, cyber security instructor, and more. When a potential participant sees the request for response on their LinkedIn page and chooses to respond, they were sent to the researcher's google forms site to provide their consent to participate, email, and other professional demographic information. Once ready, the participants were sent an email which linked to the researchers' dynamic crowdsourcing platform's site. This site was a customization of the crowd sourcing platform *Argupedia* [7] for eliciting an ontology of topics and was published on Google Cloud (<https://topics-ontology.uc.r.appspot.com/public/welcome>).

Once at the crowdsourcing site, the respondent was able to view the course's title, description and other attributes in Appendix A. The role of each respondent of the expert crowd was to go through and confirm (or reject) the initial proposed list of topics and subtopics (initially determined by the curriculum owner's baseline curriculum) as relevant to the course and industry. They were also able to suggest modifications to and new additions of topics and subtopics as they feel fit. When responding, the expert crowd provided three main inputs: 1 – their professional demographics; 2 – their confirmation or rejection of the presented topics and subtopics; and 3 – their optional input on the modifications to and newly suggestions of topics and subtopics.

As inputs came in through the dynamic crowdsourcing platform, the ontology platform was used to categorize and organize them into a topical hierarchy which was used to visually represent the hierarchical relationship of the topics, subtopics, and sub-subtopics. The inputs were moderated in order to avoid spam or out-of-scope inputs.

Results and discussion

While the experiment is currently being conducted, the experiment has received inputs from 15 expert participants across the cyber industry. In total, these 15 expert participants provided 996 different data points through confirming (753), revising (145), rejecting (45), and suggesting new curriculum topics and subtopics (42), while 11 decisions were postponed. Of the 8 main topics from the baseline curriculum, all 8 were able to be confirmed with high confidence based on the 75 confirmations and 42 revisions inputted by the 15 expert participants. The researchers noted revisions as confirmations because the experts were given the direction to rename a topic or subtopic only when they agree that the topic or subtopic is important for the course but could be stated in a better way in the curriculum. Otherwise, the topic or subtopic would be rejected. There was only one input of "reject" out of the 120 inputs received for main topics.

Of the 44 subtopics from the baseline curriculum, 42 were confirmed with high confidence and 2 were accepted with low confidence based on the 479 confirmations and 67 revisions inputted by the 15 expert participants. There were an additional 24 "reject" inputs as well, but the number of rejects did not bring the confidence level below the acceptance threshold.

Moreover, there were also 40 newly suggested subtopics provided by the expert crowd, 10 received the recommended number of analyses (RecNA), 9 received the minimum number of analysis (MinNA) while the remaining 21 are recent provided topics with insufficient analyses provided. For the 10 newly suggested subtopics that received the recommended number of

analyses (RecNA), 6 are accepted with high confidence, 2 are accepted with low confidence, and 2 remain undetermined. For the newly suggested subtopics that received only the MinNA, 1 was accepted with low confidence, 3 are still undetermined, and 5 are currently rejected with low confidence. The remaining 21 have not yet received the required minimum number of analysis (MinNA) as required by the consensus building method, so they are currently in the state of “undecided.”

The results show that expert crowdsourcing and the consensus building method not only allow for the curriculum owner to asynchronously confirm that their curriculum topics and subtopics are relevant to the industry, but can also be used to receive and achieve agreement among the expert crowd on additional topics and subtopics that are also relevant to the industry and should be included in the curriculum for the graduates to be more competitive and ready for the job market.

When compared to the previous experiment [6], the amount of newly suggested subtopics was 50% more than this current experiment. This is probably due to the fact that in the previous experiment, there was no way to revise a topic or subtopic so many experts suggested a new subtopic that was a revision of a proposed subtopic. However, in this experiment, the revision feature was implemented and used 67 times which helped decrease the amount of newly suggested subtopics. In this experiment, the confirmation rates were higher per subtopic ($\geq 80\%$) than the previous experiment ($\geq 60\%$).

Limitations

While the initial results show that expert crowdsourcing with the consensus building method can yield impactful findings for the curriculum owner, it should be noted that our crowd was limited in size. While the current size of crowd experts (15) is greater than a typical curriculum committee, it will be interesting to view the results when 50+ crowd experts provide their inputs by the end of May 2023. The research also only focused on one course’s curriculum within an academic program. It will be interesting to see how the expert crowd responds and achieves consensus when another course’s curriculum is in question. Another limitation was of the LinkedIn Campaign Ad features. When targeting users for their crowd participation, the job positions for security professionals and security management were abundant. However, there were only three cyber security academic related roles that could be targeted on LinkedIn Campaign Ad – and two of them were more generic information technology teaching roles. This may have impacted the amount of academics that participated in the experiment.

Path forward

There is still much to be done including enhancing the crowdsourcing platform on which the consensus building method is being utilized. This includes creating features that allow for the ordering of the topics and subtopics as desired by the expert crowd; allowing for experts to provide comments for each topic and subtopic to clarify their thoughts if necessary; allowing for the moving of subtopics from one topic to another if desired; and some other usability features as suggested by the expert crowd such as a progress bar.

The next step in the research effort will be considering and integrating the professional diversity factor into the consensus building process. This will allow us to account for each crowd worker’s

years of experience, area of expertise, sector, and more. Their inputs will be tied to their professional demographics which will have different weight factors on the consensus building function depending on the number of respondents within each area of measure.

Comparing the expert crowd's curriculum development outcomes to the curriculum owner's baseline and the curriculum committee's curriculum outcome is also on the to-do list. An evaluation committee will be put together to rank the three curriculums to determine which seems to be the most industry relevant.

References

1. Gupta, U. (2016). Crowd sourcing decisions: an uphill task. *Human Capital Magazine*, 52-55.
2. Satterfield, J.M., Adler, S.R., Chen, H.C., Hauer, K.E., Saba, G.W., & Salazar, R. (2010). Creating an ideal social and behavioural sciences curriculum for medical students. *Medical Education*, 44(12), 1194-1202.
3. Nakayama, S. (2012). SH&E curriculum: involving practicing safety professionals in its development. *Professional Safety*, 57(5), 68-73.
4. Thompson, T.A., & Purdy, J.M. (2009). When a Good Idea Isn't Enough: Curricular Innovation as a Political Process. *Academy of Management Learning & Education*, 8(2), 188-207.
5. Woodward, B., Imboden, T., & Martin, N.L. (2013). An undergraduate information security program: more than a curriculum. *Journal of Information Systems Education*, 24(1), 63-69.
6. Ngac, B.K. & Boicu, M. (2022). Crowdsourcing Cyber Experts to Determine Relevant Topics During Cyber Curriculum Development Efforts. *Journal of Innovations in Education and Teaching International*, 59(2), 1-11.
7. Boicu, M., Marcu, D., Tecuci, G., Kaiser, L., Uttamsingh, C., Kalale, N. (2018). Co-Arg: Cogent Argumentation with Crowd Elicitation, *Proceedings of the 2018 AAAI Fall Symposium "Artificial Intelligence in Government and Public Sector"*, Arlington, VA, October 18-20, Technical Report, AAAI Press: Palo Alto, CA, [arXiv:1810.01541](https://arxiv.org/abs/1810.01541) [cs.AI], <https://arxiv.org/abs/1810.01541v1>
8. Federal Cyber Workforce Management and Coordinating Working Group. (2022). State of the Federal Cyber Workforce: A Call for Collective Action. Cybersecurity and Infrastructure Security Agency.

Appendix A

Baseline Curriculum

ISM 650: Managing Network and Data Security

Target Audience: Graduate Students in the School of Business

Description: The course covers principles and practices of assessing network and data security needs, and implementation of the necessary security plan for communication networks and organizational data. The course addresses databases as well as unstructured data in files, including securing cloud infrastructure. More specifically, the course will focus on database security principles, database auditing, database reliability and implementation of database controls and security. In context of CIA triad of information security, threat classifications to the communication networks are discussed: eavesdropping (confidentiality), man-in-the-middle (integrity), and denial-of-service (availability). Real world examples of attack methods and cases of database breaches are discussed, as well as attacks on communication networks and networked applications are discussed to translate principles into reality. Security design and architecture consisting of authentication, authorization, access control, traffic monitoring, secure protocols are covered. Class project requires students to undertake security requirements assessment, conduct an audit and present a network and data security plan.

Course is set for Eight Weeks

Pre-requisite Courses

MIS 310: Database Management Systems (or equivalent)

AND

MIS 320: Networks & Security (or equivalent)

Corequisite Course

ISM 603: Fundamentals of Information Security Management

****Topics Taught in ISM 603:***

1. The Cyber Environment & Current / Emerging Threats
2. Information Security Fundamentals & Principles
 3. Identity & Access Management
4. Host Hardening & Vulnerability Assessments
5. Supply Chain Cyber Security Overview
6. Cloud Security & Managed Security Service Providers
7. Framework & Architecture: Cyber Security, Risk Management, & Zero Trust
 8. Incident Response & Business Continuity

Appendix B

Topics & Subtopics

1. Review of IT Networking & Database Management Fundamentals
 - IT Networking Technologies
 - IT Networking Protocols
 - OSI Reference & TCP/IP Network Models
 - Database Design Fundamentals
 - Database Technologies & Systems
 - Structured Query Language Basics
2. Examination of Past Network, Data, & Application Security-Related Attacks
 - Target Security Incident (2013)
 - U.S. Office of Personnel Management Security Incident (2015)
 - Equifax Security Incident (2017)
 - Marriott Security Incident (2018)
 - Colonial Pipeline Security Incident (2021)
 - SuperVPN, GeckoVPN, & ChatVPN Security Incident (2022)
3. Securing Applications & Their Data
 - Data Discovery, Classification, & Valuation
 - Securing Databases & User Devices
 - Access Control to the Data
 - Security of Data Backups & the Restoration Process
 - Data Retention & Disposal
 - Application Security: OWASP Top 10
4. Securing Communications & Networks
 - Network Security Basics: Firewalls; Intrusion Detection & Prevention Systems
 - Securing Network Hardware
 - Access Control to the Network
 - Securing Network Traffic
 - Virtual Private Networks
 - Securing & Auditing Network Operations
5. Cryptography & Encryption
 - Cryptography Basics: Hashes; Symmetric & Asymmetric Cryptography; Hybrid Cryptography
 - Encryption at Rest & in Transit; IPsec
 - Encryption Algorithms
 - Quantum Cryptography
6. Network & Data Security in the Cloud
 - Overview of Cloud Computing
 - Security Benefits in the Cloud Environment
 - Security Challenges in the Cloud Environment: Lack of Visibility, Multitenancy, & Misconfigurations
 - Implementing on-premises Security Solutions in the Cloud Environment

7. Architecting Network & Data Security Design

- Requirements Determination with Business Stakeholders
- Unified Threat Management (UTM)
- Endpoint Protection & BYOD
- Data Loss/Leak Prevention (DLP)
- Zero-Trust Security Design
- Physical Security of Networks, Data Centers, & Devices

8. Managing Network & Data Security Operations

- Configuration Management: Change & Patch Management
- Managing Exceptions in Network & Data Security
- Security Information & Event Management (SIEM)
- Managed Security Service Providers (MSSPs)
- Automated Continuous Auditing, Compliance Checking, & Vulnerability Assessment
- Continuous Improvement of Network & Data Security

Appendix C

Topics & Subtopics

1. Review of IT Networking & Database Management Fundamentals
 - a. IT Networking Technologies
 - b. IT Networking Protocols
 - c. OSI Reference & TCP/IP Network Models
 - d. Database Design Fundamentals
 - e. Database Technologies & Systems
 - f. Structured Query Language Basics
2. Examination of Past Network, Data, & Application Security-Related Attacks
 - a. Target Security Incident (2013)
 - b. U.S. Office of Personnel Management Security Incident (2015)
 - c. Equifax Security Incident (2017)
 - d. Marriott Security Incident (2018)
 - e. SolarWinds Security Incident (2019)
 - f. Colonial Pipeline Security Incident (2021)
 - g. SuperVPN, GeckoVPN, & ChatVPN Security Incident (2022)
3. Securing Databases
 - a. Data Discovery, Classification, & Valuation
 - b. Database Security Threats: (i.e. SQL Injection)
 - c. Securing Databases & SQL Security Commands: GRANT, REVOKE, VIEW, etc.
 - d. Database Reliability: Security of Data Backups & the Restoration Process
 - e. Data Retention & Disposal
 - f. Database Auditing
 - g. Application Security's Impact on Database Security: OWASP Top 10
 - h. Securing Unstructured Data
4. Securing Communications & Networks
 - a. Network Security Basics: Firewalls; Intrusion Detection & Prevention Systems
 - b. Common Network Threats: Eavesdropping, man-in-the-middle, & denial-of-service
 - c. Securing Network Hardware
 - d. Securing Network Traffic
 - e. Virtual Private Networks
 - f. Securing & Auditing Network Operations
5. Cryptography & Encryption
 - a. Cryptography Basics: Hashes; Symmetric & Asymmetric Cryptography; Hybrid Cryptography
 - b. Encryption at Rest & in Transit; IPsec
 - c. Encryption Algorithms
 - d. Overview & Impacts of Quantum Cryptography
6. Network & Data Security in the Cloud
 - a. Security Benefits in the Cloud Environment
 - b. Security Challenges in the Cloud Environment: Lack of Visibility, Multitenancy, & Misconfigurations

- c. Implementing on-premises Security Solutions in the Cloud Environment
- 7. Designing & Architecting Network & Data Security
 - a. Requirements Determination with Business Stakeholders
 - b. Unified Threat Management (UTM)
 - c. Endpoint Protection & BYOD
 - d. Data Loss/Leak Prevention (DLP)
 - e. Zero-Trust Security Design
 - f. Physical Security of Networks, Data Centers, & Devices
 - g. Planning Security Implementations & Upgrades
- 8. Managing Network & Data Security Operations
 - a. Configuration Management: Change & Patch Management
 - b. Managing Exceptions in Network & Data Security
 - c. Security Information & Event Management (SIEM)
 - d. Managed Security Service Providers (MSSPs)
 - e. Automated Continuous Auditing, Compliance Checking, & Vulnerability Assessment
 - f. Continuous Improvement of Network & Data Security