Scoping the Cyber Security Body of Knowledge

Awais Rashid, George Danezis, Howard Chivers, Emil Lupu, Andrew Martin,
Makayla Lewis and Claudia Peersman

Cyber security is becoming an important element in curricula at all education levels. However, the foundational knowledge upon which the field of cyber security is being developed is fragmented and, as a result, it can be difficult for both students and educators to map coherent paths of progression through the subject. By comparison, mature scientific disciplines like mathematics, physics, chemistry and biology have established foundational knowledge and clear learning pathways. Within software engineering, SWEBOK, the IEEE Software Engineering Body of Knowledge [1], codifies key foundational knowledge upon which a range of educational programmes may be built. There are a number of previous and current efforts on establishing skills frameworks, key topic areas and curricular guidelines for cyber security (see sidebar). However, a consensus has not been reached on what the diverse community of researchers, educators and practitioners sees as *established foundational knowledge* in cyber security.

The Cyber Security Body of Knowledge (CyBOK) project [2] aims to codify the foundational and generally recognised knowledge on cyber security. In the same fashion as SWEBOK, CyBOK is meant to be a *guide* to the body of knowledge—the knowledge that it codifies already exists in literature such as textbooks, academic research articles, technical reports, white papers and standards. Our focus is, therefore, on mapping *established knowledge* and not fully replicating everything that has ever been written on the subject. Educational programmes ranging from secondary and undergraduate education through to post-graduate and continuing professional development programmes can then be developed on the basis of CyBOK.

Since the 1st of February 2017, we undertook a range of community consultations (cf. Tables 1 and 2), both within the UK and internationally, through a series of different activities designed to gain as much input as possible and from as wide an audience as possible. In addition, analysis of a number of relevant texts (44 in total), such as tables of contents of textbooks, calls for papers for conferences and symposia, standards, existing certification programmes, etc. was undertaken to complement the insights gained from the community consultations. The insights from these activities were synthesised to develop a Scope for CyBOK and 19 top-level Knowledge Areas (KAs) identified. We next discuss the Scoping Research before discussing the KAs that emerged.

Table 1: Scoping research activities and number of participants/responses

Online Survey	44 responses received
Analysis of relevant texts	44 separate texts analysed
In-depth interviews with key experts	10 interviews undertaken
Community workshops across the UK	11 workshops 106 attendees
Call for positions statements	13 statements received
Panel at Advances in Security Education Workshop at USENIX Security Symposium, Vancouver, Canada, October 2017	Paper-based exercise with 28 attendees

Table 2: Distribution of input from academia and practitioners

	Academic (%)	*Practitioner (%)*
Online Survey	51	49
In-depth interviews with key experts	50	50
Community workshops across the UK	55	45
Call for positions statements	62	38

Scoping Research

Consultation workshops

We took a participatory design approach to our consultation workshops that brought together 106 attendees from industry and academia in the UK – in a collaborative and creative environment – to discuss the KAs that ought to be included in CyBOK. Some workshops were dedicated to consultation with academia and others to consultation with practitioners. A subset also included representatives from both academia and practitioner communities.

The workshops were based on a *supermarket* metaphor (Figure 1) whereby participants were encouraged to think about what they considered to be the key KAs to be included in CyBOK. Participants discussed and identified a range of KAs collectively and put each KA into one of the four *supermarket* areas:

- *In the trolley* – KAs to be included;
- *On the shopper's heart* – KAs that are of interest to participants but not necessarily to be included;
- *On the shelf* – KAs to be discussed further;
- *In the bin* – KAs deemed out of Scope;

This sorting exercise was followed by a *15 items or less* task during which participants were asked to sort the 'in the trolley' KAs into groups of top-level and sub-level KAs.

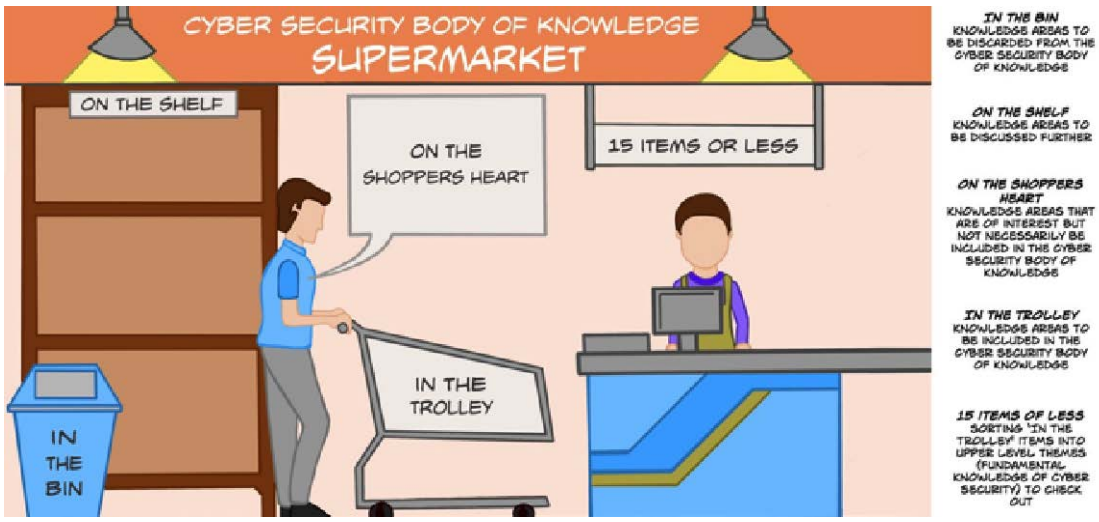

Figure 1: The *Supermarket* metaphor used for participatory workshops

This workshop design allowed for small group discussion on where KAs ought to be best placed and why. It also led to sub-topics within knowledge areas to be identified.

In addition to these workshops, consultations were also held at the Higher Education Academy Conference in Liverpool, UK in April 2017 and the Cyber Security Professionals Conference in York, UK in May 2017.

A panel discussion was also organised at the Advances in Security Education Workshop at the USENIX Security Symposium in Vancouver, BC, Canada in August 2017 and views on relative importance of particular topics emerging from the above workshops were sought via a paper-based exercise.

Complementing the Consultation Workshops

The workshop consultations were complemented by an online survey involving a series of open- and closed-ended questions on KAs that may form part of the CyBOK. The survey sought participants' views on topics such as: the KAs that had been most important background knowledge in their career; key KAs that ought to be covered in the CyBOK and those that should be out of scope; and topics that would be of most importance over the next 5 years.

Semi-structured interviews were conducted with 10 leading international experts in cyber security. The interviews included both technical experts in computer security and those studying topics such as human factors, governance, regulation, risk and law.

A small amount of input was also received through an open call for position papers.

Analysis of various texts listing key topics

We complemented the data arising from the above community consultations with analysis of a number of documents that typically list key topics relevant to security. Example documents included:

- Categorisations, such as the ACM CCS taxonomy;
- Certifications, such as CISSP and the IISP Skills Framework;
- Calls for Papers such as IEEE Symposium on Security and Privacy, USENIX Symposium on Usable Privacy and Security;
- Existing curricula, such as ACM Computer science curriculum, work of the ACM Joint Task Force on Cyber Security Education;
- Standards, such as BS ISO-IEC 27032 2021, NIST IR 7298;
- Tables of contents of various textbooks.

We used a variety of text mining techniques, such as Natural Language Processing (NLP) and automatic text clustering to cluster relevant topics and identify relationships between topics. Techniques utilised included semantic word cloud visualisations, Word Vectors, Ward clustering, K-means clustering and Latent Dirichlet Allocation (LDA).

Distilling the Knowledge Areas

Workshop participants identified key topic areas together with subsidiary topics that they considered should be included in each area. This provided the opportunity to visualise the workshop data as a graph in which nodes were highlighted according to the strength of recommendation as a key area, and edges weighted to show the strength of relationship between topics. Inevitably the workshops resulted in a large number of unique terms for topics: a total of 906 unique terms, with 660 occurring only once in the record. Some data cleaning was therefore necessary. Cleaning was carried out via an alias list that could be

inspected and reviewed, and this resulted in 483 unique terms, in which a core of 144 topics occurred more than once. The final graph was filtered by edge weight to allow review of the data at different levels of granularity. An example graph is shown in Figure 2. Note that the colours represent visual distinction into quartiles and that the size of each topic name is proportional to its frequency as a nominated *key area.*

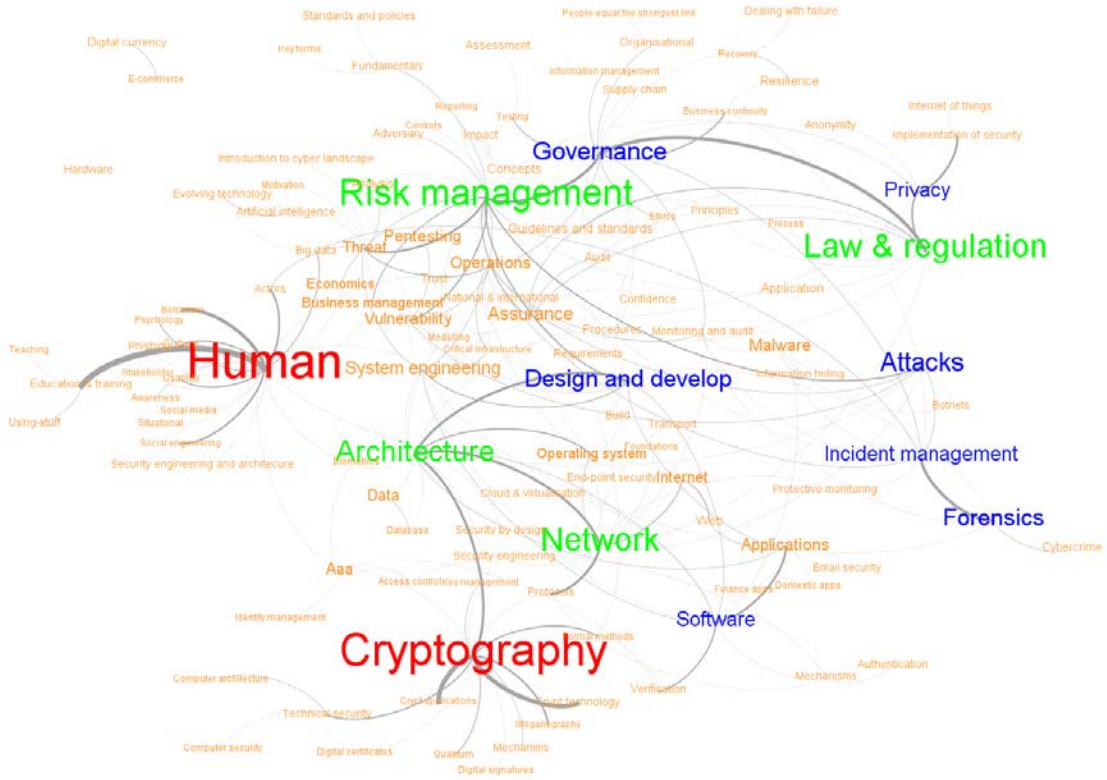

Figure 2: An example graph derived from the analysis of workshop data

We note that the graph in Figure 2 is one example of the types of graphs that served as the starting point for our distillation of KAs. The thematic clusters emerging from such graphs were cross-referenced against the data from the survey, the interviews and the position statements. The sub-topic lists in these thematic clusters were further compared with the clusters identified through the text mining analysis of documents listing key topics in cyber security. During this synthesis, we particularly attended to topics, for instance, Hardware Security and Cyber-Physical Systems Security, that appeared disconnected or did not formulate large clusters in the graphs but were highlighted by our survey or interview participants as key emerging topics of importance over the next five years.

This analysis and synthesis resulted in 19 KAs, grouped into five broad categories (see Figure 3 and Table 3). Figure 3 shows that system, infrastructure, software and platform security is shaped by human and organisational factors and vice versa. At the same time cyber security of technologies, people and organisation requires a deep understanding of attacker behaviours, attack technologies as well as effective responses for analysis of attacks, incident management and response. We note that other possible categorisations of these KAs may be equally valid. Also the categories are not necessarily orthogonal.

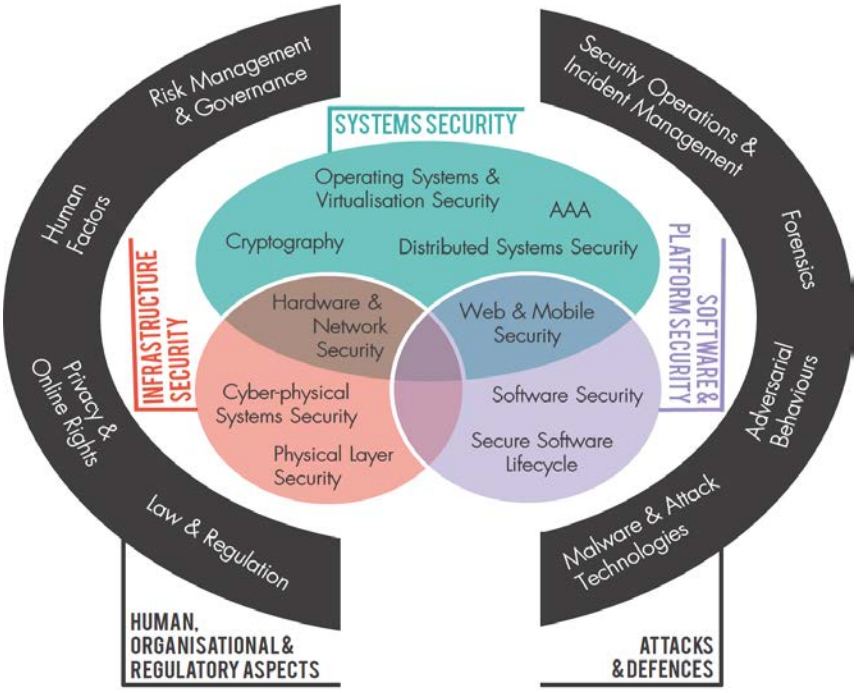

Figure 3: The 19 Knowledge Areas and their categorisation within CyBOK

Table 3: Overview of the 19 Knowledge Areas

Human, Organisational and Regulatory Aspects	
Risk Management and Governance	Security management systems and organisational security controls, including standards, best practices and approaches to risk assessment and mitigation.
Law and Regulation	International and national statutory and regulatory requirements, compliance obligations and security ethics, including data protection and developing doctrines on cyber warfare.
Human Factors	Usable security, social and behavioural factors impacting security, security culture and awareness as well as the impact of security controls on user behaviours.
Privacy and Online Rights	Techniques for protecting personal information, including communications, applications and inferences from databases and data processing. It also includes other systems supporting on-line rights touching upon censorship and circumvention, covertness, electronic elections and privacy in payment and identity systems.
Attacks and Defences	
Malware and Attack Technologies	Technical details of exploits and distributed malicious systems, together with associated discovery and analysis approaches.
Adversarial Behaviours	The motivations, behaviours and methods used by attackers, including malware supply chains, attack vectors and money transfers.
Security Operations and Incident Management	The configuration, operation and maintenance of secure systems including the detection of and response to security incidents and the collection and use of threat intelligence.
Forensics	The collection, analysis and reporting of digital evidence in support of incident or criminal events.
Systems Security	
Cryptography	Core primitives of cryptography as presently practised and emerging algorithms, techniques for analysis of these and the protocols which use them.
Operating Systems and Virtualisation Security	Operating systems protection mechanisms, implementing secure abstraction of hardware and sharing of resources, including isolation in multi-user systems, secure virtualisation and security in database systems.

Distributed Systems Security	Security mechanisms relating to larger scale coordinated distributed systems, including aspects of secure consensus, time, event systems, peer-to-peer systems, clouds, multi-tenant data centers, and distributed ledgers.
Authentication, Authorisation and Accountability	All aspects of identity management and authentication technologies, and architectures and tools to support authorisation and accountability in both isolated and distributed systems.
Software and Platform Security	
Software Security	Known categories of programming errors resulting in security bugs, and techniques for avoiding these errors - both through coding practice and improved language design, and tools, techniques and methods for detection of such errors in existing systems.
Web and Mobile Security	Issues related to web applications and services distributed across devices and frameworks, including the diverse programming paradigms and protection models.
Secure Software Design and Development	The application of security software engineering techniques in the whole systems development lifecycle resulting in software that is secure by default.
Infrastructure Security	
Network Security	Security aspects of networking and telecommunication protocols, including the security of routing, network security elements and specific cryptographic protocols used for network security.
Hardware Security	Security in the design, implementation, and deployment of general-purpose and specialist hardware, including trusted computing technologies and sources of randomness.
Cyber-Physical Systems Security	Security challenges in cyber-physical systems, such as IoT and industrial control systems, attacker models, safe-secure designs, security of large-scale infrastructures.
Physical Layer Security	Security concerns and limitations of the physical layer including aspects of radio frequency encodings and transmission techniques, unintended radiation, and interference.

Next Steps

The initial CyBOK Scope and KAs identified above were made publicly available for community comments in September 2017. While none of the 19 KAs needed to be removed or new ones added on the basis of the feedback, the topics to be covered under each KA have been refined. As a next step, authors will be invited to write detailed descriptions of KAs which will be reviewed by a small panel of peer-reviewers before being made available for public consultation. As each KA description is finalised, it will be made available on the CyBOK web site. We aim to complete all KA descriptions by the end of July 2019. Alongside, learning pathways through CyBOK and exemplar curricula at different education levels will be developed. We will undertake a series of consultations through workshops and interviews with stakeholders not only involved in university education but also those from primary and secondary education, as well as industrial training programmes. Combined with desk research on curricula, such consultations will form the basis to develop a set of exemplar learning pathways as a set of case studies for utilising CyBOK in educational programmes.

Cyber security is a rapidly changing and evolving field. As such the CyBOK will never be 'finished' per se. Future iterations will need to be undertaken to ensure that the coverage remains up-to-date and the KAs reflect both current state of knowledge in cyber security and emerging needs. The inclusion of KAs such as Hardware Security and Cyber-Physical Systems Security in the current Scope reflects such emerging needs. Any future maintenance of CyBOK will need to ensure that, whilst not ignoring the needs of contemporary and legacy systems, the CyBOK scope also reflects key challenges arising from the increasing integration of technology – and hence cyber security – into the very fabric of our society.

Sidebar: Related work on identifying core concepts in cyber security

The ACM, IEEE, AIS and IFIP Joint Task Force on Cyber Security Education (JTF) has developed guidelines for undergraduate curricula in cyber security [3]. Five principle knowledge areas are considered – based on the entities to be protected: Data Security, Software Security, Component Security, Connection Security, System Security, Human Security, Organisational Security and Societal Security. These are complemented by crosscutting concepts such as Confidentiality, Integrity and Availability. Undergraduate cyber security curricula can then be designed for particular disciplines, e.g., Computer Science, Software Engineering, etc. and/or linked to particular application areas. In contrast, CyBOK aims to codify foundational knowledge that can inform the design of cyber security education and training programmes at a range of levels: from secondary and undergraduate through to postgraduate and continuing professional development. It complements the work of the JTF by providing in-depth coverage of KAs and key resources that curriculum designers can utilise.

The National Initiative for Cybersecurity Education (NICE) Cybersecurity Workforce Framework (NICE Framework) [4] has established a set of speciality areas and mapped them to roles in the cyber security workforce. The focus is on skills and the tasks a particular role ought to be able to perform. CyBOK can form the basis of charting the learning pathways that such skilled roles may need to take across the 19 KAs (or a subset thereof) in order to be able to proficiently perform the required tasks.

The Cyber Security Assessment Tools (CATS) project has undertaken a Delphi study identifying the importance, difficulty and timelessness of particular cyber security topics [5]. Such understanding is essential to the design of cyber security education programmes. It would be interesting to explore where the topics of most difficulty and importance appear in the 19 CyBOK KAs and, combined with charting of learning pathways for the NICE framework, how this may inform pedagogical approaches to cyber security.

The security counterpart to the SWEBOK [1] is “Software Assurance: A Guide to the Common Body of Knowledge to Produce, Acquire and Sustain Secure Software” [6]. This has similar style and chapter headings to the SWEBOK and provides a summary of knowledge relating to software and the software lifecycle. Software development is within the scope of the CyBOK, which in contrast has the wider scope of fundamental and applied knowledge in all aspects of cyber security.

Knowledge required for the CISSP examination has also been codified in a body of knowledge [7]. The CISSP CBK documents the knowledge required for a specific examination in a summary textbook; this is in contrast with other bodies of knowledge and the CyBOK, the contents of which guide readers to knowledge contained in authoritative references.

Acknowledgements
The CyBOK project is sponsored by the National Cyber Security Programme in the UK. The authors also thank Yvonne Rigby, Project Manager for CyBOK, for her excellent work on coordinating the various strands of research across the project.

References

[1] Guide to the Software Engineering Body of Knowledge (SWEBOK Guide), https://www.computer.org/web/swebok
[2] Cyber Security Body of Knowledge (CyBOK), https://www.cybok.org/
[3] Joint Task Force on Cyber Security Education (CSEC), https://www.csec2017.org

[4] NICE Framework, http://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-181.pdf
[5] G. Parekh, D. DeLatte, G. L. Herman, L. Oliva, D. Phatak, T. Scheponik, A. T. Sherman, "Identifying Core Concepts of Cybersecurity: Results of Two Delphi Processes", IEEE Transactions on Education, 2017.
[6] S. T. Redwine, Ed., "Software Assurance: A Guide to the Common Body of Knowledge to Produce, Acquire and Sustain Secure Software, Version 1.1," U.S. Department of Homeland Security, Washington, 2006.
[7] Gordon, Adam, ed. "Official (isc) 2 Guide to the Cissp CBK." CRC Press, 2015.