Deloitte.

Can a Robot Do My Job?

A Study on the Potential of Artificial Intelligence to Take on Cybersecurity Tasks

Esther van Luit (S1789821)
7 December 2017

Thesis Supervisors:
Prof.dr.ir. Jan van den Berg, TU Delft & Universiteit Leiden
Dr. Jan Morsch, Nyenrode Business Universiteit

This thesis was written in fulfilment of the requirements of the Executive Master in Cybersecurity from the Cyber Security Academy.

Abstract

The cybersecurity industry is currently dealing with a shortage of qualified security workers, predicted to be 1.8 million workers in 2022. This thesis investigates the potential of artificial intelligence to take over cybersecurity work. Security tasks are derived from NIST SP800-181. Criteria are formulated on whether artificial intelligence can perform a certain task. It is assessed whether a task cannot, partially or fully be outsourced to artificial intelligence. 22.1% of security tasks could be fully outsourced and 37.1% of the tasks could be partially outsourced to artificial intelligence. Translating this to impact on specific work roles, 19 of the 52 work roles are mostly non-outsourceable, whereas only 4 roles are mostly outsourceable. The roles of which the majority could be outsourced to artificial intelligence are 'System Testing and Evaluation Specialist', 'Technical Support Specialist', 'Cyber Defense Analyst' and 'Cyber Defense Forensics Analyst'. The macro-economic impact of outsourcing work roles is assessed for the United States of America. Assuming a similar shortage for each role globally, an extrapolation is made to discern potential impact on the total cybersecurity skills gap in the context of various scenarios of artificial intelligence adoption over time. If 100% AI adoption were to occur by 2022 for all of the tasks named fully or partially outsourceable in this research, 45% of the global cybersecurity skills gap or over 800,000 jobs in 2022 could be outsourced to artificial intelligence. The results are put into context of more work being created because of the implementation and security needs of artificial intelligence and the results are validated by looking at the extent other new technologies have created technological unemployment.

Key words: *NIST SP800-181, NICE Cybersecurity Workforce Framework, artificial intelligence, computational intelligence, technological unemployment, cybersecurity skills gap.*

Preface

This thesis continues my previous thesis on skills gaps at Nyenrode Business Universiteit [1]. During my first job at a strategy consultancy in payments I let the topic of skills gaps rest. Next, working at Deloitte I found that the cybersecurity industry was in dire need of more and better professionals to guard cyberspace. So I started to research the topic once more, this time for the cybersecurity industry. During my career at Deloitte I was presented the opportunity to do a second Masters, in Cybersecurity. The matter of why we learn and what we should learn has always fascinated me. Therefore this thesis looks into technology's impact on the need to 'know' and 'be able' for one of our core activities - work. The question whether technology will replace some of us at our jobs is as old as technology itself. The answer has always been affirmative, but then again, new jobs have often emerged instead. Humankind faces a new frontier as the rise of general artificial intelligence nears. Sensorisation, deep learning and big data analysis make intelligent computerized labour attractive to employers. In cybersecurity as well, intelligent network security and advanced virus protection hold promise.

In February 2017 IBM released their 'Cognitive Security' product. It consists of IBM QRadar enhanced with IBM Watson. While not the first application to use machine learning, its uniqueness lay somewhere else. IBM marketed the product as the solution to the cybersecurity skills gap and lay out a roadmap for doing so. This triggered me to want to reflect on the potential of artificial intelligence in this industry. This thesis attempts to put that reflection into an academic, practical and societal perspective.

I would like to thank Deloitte Netherlands for offering me the chance to complete this second masters. The cybersecurity industry is broad and I feel this has allowed me to grow as a cybersecurity professional. I want to thank Deloitte UK for their support in me writing my thesis during my secondment and showing interest in the results. My partner Roland Schagen has been a great sounding board, despite having his hands full on his own thesis. He is probably the smartest person I know and an invaluable sidekick. I had two fantastic supervisors help me finish this thesis putting up with me being in the UK. Prof.dr.ir. Jan van den Berg has helped tremendously with his insights in the world of artificial- and computational intelligence. I am grateful he has supported me in what is in fact a technology-driven social study. Dr. Jan Morsch has stood by me during my thesis period at Nyenrode. He provided great feedback then and I was happy and honoured to receive his feedback again. Lastly, I would like to thank Will Markow from Burning Glass for providing me with the data required to perform this research and his professional advice.

Table of Contents

Abstract 2

Preface 3

1. Introduction 7

   1.1 Scope & Assumptions 9

   1.2 Relevance 10

   1.3 Outline 11

2. Cybersecurity Tasks 12

   2.1 Background 12

   2.2 Approach 13

   2.3 Results 15

   2.4 Summary 17

3. Artificial Intelligence Capabilities 19

   3.1 Background 19

   3.2 Approach 20

   3.3 Results 22

   3.4 Summary 28

4. Outsourcing Cybersecurity Tasks 29

   4.1 Background 29

   4.2 Approach 31

   4.3 Results 32

   4.4 Summary 38

5. Potential Impact on the Cybersecurity Skills Gap 39

   5.1 Background 39

5.2 Approach ... 40
5.3 Results ... 41
5.4 Summary ... 44
6. Conclusion & Discussion ... 46
6.1 Conclusion ... 46
6.2 Reflection ... 47
6.3 Counteracting effects of AI Adoption ... 49
6.3.1 Training, implementation and maintenance ... 49
6.3.2 Securing artificial intelligence ... 50
6.3.3 Using artificial intelligence against security ... 51
6.4 Comparison to Other New Technologies ... 51
6.5 Academic and Societal Relevance ... 54
6.6 Future Research ... 55
7. References ... 56
Appendix A – Specialties and Work Roles ... 61

Figures

Figure 1. Conceptual model 9
Figure 2. Security work roles demanded in 2015 GISWS [46] 13
Figure 3. Number of tasks per domain 15
Figure 4. Division of tasks over roles 18
Figure 5. Task analysis template (partial view) 33
Figure 6. Task analysis template upon completion (partial view) 33
Figure 7. Supply/demand ratio based on online job postings ([89], [90]) 42
Figure 8. Supply/demand ratio based on online job postings with 100% AI adoption ([89], [90]) .. 42
Figure 9. Status quo of the cybersecurity skills gap (no AI adoption for outsourceable tasks) 43
Figure 10. Various AI adoption scenarios and their impact on the cybersecurity skills gap 45
Figure 11. Percentage of USA workers in different industries (data from [32]) 52
Figure 12. Percentage of time spent in US occupations (adapted from [106]) 53

Tables

Table 1. Overview of the 33 NIST SP800-181 specialties [28] 14
Table 2. Average number of connections to role nodes 15
Table 3. Frequency table for number of connections to role nodes 16
Table 4. Human competencies for artificial intelligence (adapted from [54, p. 32]) 21
Table 5. CI competencies for artificial intelligence (based on [50]) 22
Table 6. Puigbo et al.'s tasks linked to competencies [58, p. 111] 24
Table 7. Poole et al.'s tasks for three artificial intelligence systems [59, pp. 13–17] 25
Table 8. Environment, task and agent characteristics (based on [61]) 27
Table 9. Current AI applications in security 30
Table 10. Division of tasks over 'Yes', 'Partial' and 'No' categories 32
Table 11. Frequency of reasons over 'Yes' category 35
Table 12. Frequency of reasons over 'Partial' category 35
Table 13. Frequency of reasons over 'No' category 35
Table 14. Division of tasks in categories per work role 37
Table 15. Division of tasks in categories per security domain 38
Table 16. Cybersecurity workforce, demand, supply and skills gap 41
Table 17. Outsourceable tasks per NIST SP-800-181 domain and impact per AI adoption scenario 43
Table 18. Global cybersecurity skill gap reduction 44
Table 19. Representative jobs to be created for AI development & maintenance (adapted from [30]) 50

1. Introduction

The website https://willrobotstakeovermyjob.com was launched on May 30th 2017. It uses research by Frey and Osborne from 2013 in which they calculated the risk of 702 job types being computerized [2], [3]. Visitors of the website can search their employment category, e.g. 'doctor'. They are answered in percentages on how likely their job is to be taken over by machines. Although the website in itself is amusing, the fact that it went viral may suggest a deeper-rooted interest in this topic [4]. 'A robot might take over my job' and 'artificial intelligence might do my work better' are thoughts people from some if not all work categories should entertain.

For the cybersecurity industry it might however be an opportunity to get work done that now cannot be done, or can be done faster. The cybersecurity industry has in its growth struggled to keep up with the right number of people to do the work. The Global Information Security Workforce Study is performed biannually by cybersecurity certification association (ISC)[2] and has each time around surveyed around 20,000 cybersecurity professionals globally on the state of the cybersecurity labour market since 2004. The 2017 edition concluded that 1.8 million more security professionals will be needed than are available in 2022 [5]. This differential has been termed the 'cybersecurity skills gap' (conceptualized in 2009 by the USA military, supposed first coinage in 2015 [6, p. 2], [7, p. 2]). Multiple governments have launched initiatives to remediate the cybersecurity skills gap. These consist mostly of improving the industry's image and embedding technology early on in education[8], [9], [10, p. 88]. Despite their positive influence on the gap, such strategies cannot be expected to wholly and swiftly resolve the cybersecurity skills gap [11].

In February 2017 IBM launched their 'Cognitive Security' platform. It uses IBM Watson to enhance IBM's QRadar SOC application. One of their claims to fame was that this was a new step forward in resolving the cybersecurity skills gap. Van Zadelhoff, IBM Security's general manager, was quoted to say the following at its inception:

> Even if the industry was able to fill the estimated 1.5 million open cyber security jobs by 2020, we'd still have a skills crisis in security. The volume and velocity of data in security is one of our greatest challenges in dealing with cybercrime. By leveraging Watson's ability to bring context to staggering amounts of unstructured data, impossible for people alone to process, we will bring new insights, recommendations, and knowledge to security professionals, bringing greater speed and precision to the most advanced cybersecurity analysts, and providing novice analysts with on-the-job training [12].

Regardless of whether IBM's product is capable of doing so, the question whether artificial intelligence can reduce the cybersecurity skills gap is an interesting one. This thesis explores different techniques, applications in general and plausible applications in the cybersecurity industry of artificial intelligence to answer that question.

To shed light on how artificial intelligence may impact the cybersecurity skills gap is not an easy feat for (at least) three reasons: a) there is much discussion about what artificial intelligence can and cannot do, now and in the future [13]–[16], b) there is an unclear taxonomy of artificial intelligence techniques and how they relate to each other [17], [18] and c) the applications of artificial intelligence techniques have not been studied for the cybersecurity industry as a whole, but only as point solutions [19]–[23].

Up until recently, what 'cybersecurity work' consisted of was a similarly difficult problem. Work roles, tasks, knowledge, skills and abilities were poorly defined [24, pp. 26–27], [25]. The field also is rapidly growing and adding new competencies to its Body of Knowledge [24, pp. 32–33], [26, p. 1], [27]. However, one of the aforementioned attempts of governments to address the cybersecurity skills gap was relatively fruitful. The USA government created the NICE Cybersecurity Workforce Development Framework. It was released in August 2017 as a NIST Standard (NIST SP800-181) and contains a comprehensive overview of cybersecurity work roles and their contents. This thesis makes use of the work roles and tasks outlined in NIST SP800-181 as it is the first document detailing the industry's workforce make-up in such detail and it has been developed with a variety of stakeholders from the government, the private sector and academia over the past six years [28, p. iv]. How competencies supplied by artificial intelligence and competencies demanded by the cybersecurity industry match up to each other has not been previously researched, but can now with the publication of NIST SP800-181. The main research question is thus as follows:

To what extent can artificial intelligence aid in reducing the cybersecurity skills gap?

Sub-questions that are of interest in answering this research question are as follows:

- What tasks does cybersecurity work consist of?
- What is artificial intelligence currently capable of?
- Which cybersecurity tasks can be performed by artificial intelligence?
- What could be the potential impact of outsourcing tasks to artificial intelligence in closing the cybersecurity skills gap?

The first sub-question must be asked as 'cybersecurity work' is not concrete enough to determine whether it can be outsourced to artificial intelligence. Tasks are a lower-level element for which it is possible to discern the actions, knowledge, skills and abilities required. They can therefore be used to determine whether tasks cannot, partly or wholly be performed by artificial intelligence.

The second sub-question must be asked to clarify what artificial intelligence is currently capable of. Artificial intelligence has branched out significantly since its origin and assessing its full breadth is necessary to understand which tasks can and cannot be performed by it.

The third sub-question must be asked to understand which and what proportion of tasks could be outsourced to artificial intelligence. From this it is possible to understand what part of cybersecurity work roles may be outsourced.

The fourth sub-question must be asked to understand the potential of artificial intelligence in closing the cybersecurity skills gap. Based on the part of roles that can be outsourced, one can determine the macro-economic implications in terms of FTE required. From these four sub-questions, the overall research question can be answered.

The conceptual model based on the research question is depicted in Figure 1, with sub-questions indicated in green bubbles. First the reader is provided with an understanding of what kind of work the cybersecurity industry contains as per sub-question 1. The independent variable in the research is artificial intelligence. By looking at various techniques based on human capabilities and those developed based on mathematics and natural patterns (CI), sub-question 2 is answered on what artificial intelligence is currently capable of, which is codified into a rule set. This rule set is based on what could feasibly be performed by artificial intelligence and whether it is deemed desirable to have a task performed by artificial intelligence (its 'business case'). Sub-question 3 uses this rule set to determine which tasks as indicated by sub-question 1 could be outsourced to artificial intelligence. This information is then used to answer sub-question 4 on our overarching dependent variable – the impact on the cybersecurity skills gap.

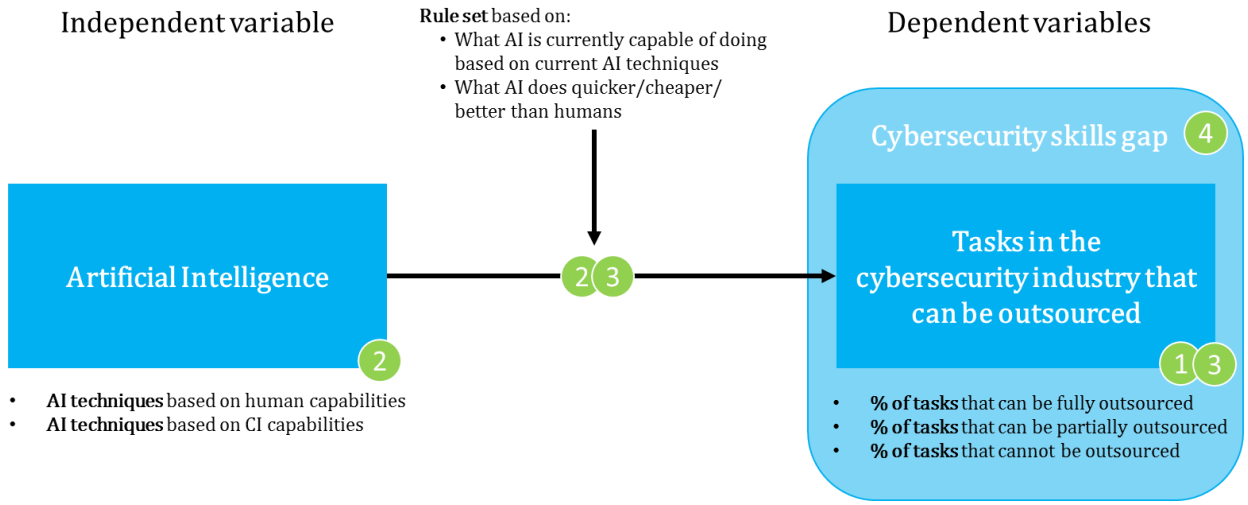

Figure 1. Conceptual model

1.1 Scope & Assumptions

This section outlines the specification and limitations to the scope for each variable. For artificial intelligence, the artificial intelligence techniques and applications under review are discussed. Regarding work to be done in the cybersecurity industry, the caveats of using NIST SP800-181 and the locality of the cybersecurity skills gap data available are discussed.

The field of artificial intelligence is continually expanding. Techniques are discovered, improved or revolutionized. New applications arise every day. This thesis looks at the set of currently accepted and tested techniques, not those with only an experimental or theoretical basis. As far as applications go, a limited set of applications in the security industry exists today. This research therefore does investigate new applications as long as they can make use of common techniques.

Most research into the make-up of the security industry has so far been conducted with the scope of the United States. Limited data is available for other geographical areas. Roles, tasks, knowledge, skills and abilities (KSAs) and the relative demand for them might not be the same in other regions of the world. Typical changes in terms of the framework are the adding/removing of specialisations, upskilling/downskilling, a shift in emphasis and scaling [29]. Predictions are made for other regions using extrapolation and assuming an American distribution of roles, tasks and KSAs required for the lack of more accurate figures.

This research only looks into the effect on the demand-side of the labour market; i.e. the number of fulfilled and unfulfilled jobs. It presumes the supply-side, i.e. the number of professionals fulfilling the jobs, to be consistent with the macro-economic predictions from the GISWS. The general demand for security services, whether executed by artificial intelligence or humans is assumed to not change from the GISWS predictions.

There may also be counteracting effects on the cybersecurity skills gap resulting from the proliferation of artificial intelligence applications. First of all, the training, implementation and maintenance of the artificial intelligence could create more work [30]. Secondly, the use of artificial intelligence for security may invite hackers to compromise their working [22]. This would call for extra professionals trying to secure the integrity of the artificial intelligence. Lastly, hackers might use these techniques and applications to their advantage and try to attack companies, governments and consumers using them [21]. This calls for extra professionals skilled in the defence against artificial intelligence attack vectors. There is limited data or theoretical framing available on these topics. Making assertions on how much the cybersecurity skills gap might widen based on these counteracting effects would be conjecture. Some qualitative insights are presented in the concluding chapter.

1.2 Relevance

The academic relevance of this thesis stems from two grounds. First, it assesses artificial intelligence capabilities for cybersecurity tasks in a systematic manner. Much has been said about what artificial intelligence can and cannot do, and when they can do more. A consistent evaluation based on rules supported by literature has not been attempted. Second, it ties the results with macro-economic data to estimate the potential impact on the cybersecurity labour market. This research may provide profound insights in how much we may expect from any technology to cause technological unemployment. Its rule-based approach can be used to assess the impacts of other technologies and on other industries. It might be worth replicating this research frequently to adapt for the evolving abilities of artificial intelligence and changing growth of the cybersecurity industry and labour market.

There is a continuing social unrest on what consequences new technologies will have on our employability [31]–[33]. Rather than investigating the impact of ‘automation’ on the cybersecurity skills gap, this thesis takes a single technology (artificial intelligence) as a starting point for showing what this impact might be to make a structural analysis of tasks concrete based on the specific capabilities of artificial intelligence. Likely a large part of the tasks remain infeasible for artificial intelligence for now and the near future. A new set of tasks probably has to be created to be able to implement artificial intelligence well [30]. A shift in the amount and the kinds of work is to be expected, not a definite loss. Looking back at research done on this topic and the impact of former ‘new’ technologies, this thesis seeks to substantiate that hunch.

1.3 Outline

This thesis continues by answering a sub-question per chapter. Each of the chapter discusses relevant background to the sub-question. The chapters then detail the approach followed to perform the analysis. The results section summarizes the results of the analysis and seeks to provide explanations. A summary recaps how the background, approach and analysis have answered the sub-question.

The second chapter answers the sub-question on what tasks the cybersecurity industry is currently comprised of. The NIST SP800-181 standard on cybersecurity work is analysed for patterns in the work. The third chapter answers the sub-question on what artificial intelligence is currently capable of. A literature review is performed to understand the competencies of artificial intelligence. A rule set is created to be able to discern what artificial intelligence can and cannot do. The fourth chapter answers the sub-question on what part of the cybersecurity tasks can be outsourced to artificial intelligence. Based on the rule set, reasons are formulated why artificial intelligence can, can partially or cannot perform a task. An analysis is performed on the NIST SP800-181 tasks to categorize tasks in these three categories. The fifth chapter answers the sub-question on how many jobs could potentially be outsourced to artificial intelligence. Based on their division across roles, an estimation is made of how much the cybersecurity skills gap could be closed. The sixth chapter concludes by evaluating and reflecting on the findings, and comparing them to previous technological revolutions. A short discussion is presented on counteracting effects that may enlarge the cybersecurity skills gap.

2. Cybersecurity Tasks

This chapter looks into the make-up of the cybersecurity industry when it comes to work. It starts off with a general exploration into the state of the current cybersecurity industry. It discusses what domains the cybersecurity industry consists, the growth of the industry and the parallel growth of the professionals required to perform the work. The standard NIST SP800-181 based on the NICE Cybersecurity Workforce Framework is used to obtain insight into various cybersecurity work roles and the tasks that they perform, their knowledge, skills and abilities. The results section presents conclusions on interesting patterns found in the dataset such as the number of tasks per role.

2.1 Background

Many put the birth of the cybersecurity industry around 1988, when Robert Morris created the first computer worm [34]–[37]. The conceptualization of the field dates back to as far as the 1960s, but it took until 21st century for cybersecurity to gain momentum [38]. The field has significantly branched out since then.

The academic realm has been divided about what cybersecurity exactly comprises [39]. It is considered to contain a variety of governance and technical topics in the security domain. This thesis makes use of the definition of 'cybersecurity' by Craigen et al.: "Cybersecurity is the organization and collection of resources, processes, and structures used to protect cyberspace and cyberspace-enabled systems from occurrences that misalign de jure from de facto property rights" [40, p. 17]. This definition combines elements of various well-known definitions and is intended to be inclusive of all cybersecurity specialisms.

A common overview of what elements make up cybersecurity work is provided by $(ISC)^2$, one of the dominant providers of professional security education. The Common Body of Knowledge associated with its CISSP certification contains eight domains: 1) security and risk management, 2) asset security, 3) security engineering, 4) communications and network security, 5) identity and access management, 6) security assessment and testing, 7) security operations and 8) software development security [41].

Up until recently there was no comprehensive overview of cybersecurity work in more detail. Then the NICE Cybersecurity Workforce Framework 1.0 (NCWF) was published in April 2013. It contained seven domains, 31 work roles and their tasks, knowledge, skills and abilities [42]. In August 2017, the latest version of the NCWF was published as NIST standard SP800-181. It now

consists of seven domains, 33 specialty areas, 52 work roles, 999 tasks, 587 knowledge elements, 365 skills and 175 abilities [28]. It can be seen as a testimony to how much the field has grown and has become more complex.

Growth has not only occurred in quality, but also in quantity. In 2004, (ISC)² published their first annual Global Information Security Workforce Study (GISWS). In the same year the global cybersecurity market was evaluated at 3.5 billion US Dollars [43]. It is projected to grow to 120 – 175 billion US Dollars by the end of 2017 and market research suggests it may be valued at 233 billion US Dollars by 2022 [43], [44].

This growth is mirrored in the amount of cybersecurity professionals required. The supply of professionals has not kept up, and is not expected to do so in the future. The 2004 GISWS made no mention of a shortage of security professionals to do the work required [45]. The 2017 GISWS concluded that 1.8 million more security professionals will be needed than are available in 2022 [5]. The 2015 GISWS sheds some light on which security domains are most in demand according to survey respondents, as shown in Figure 2. It does not provide insight in how many are demanded nor how this relates to the total demand for professionals [46].

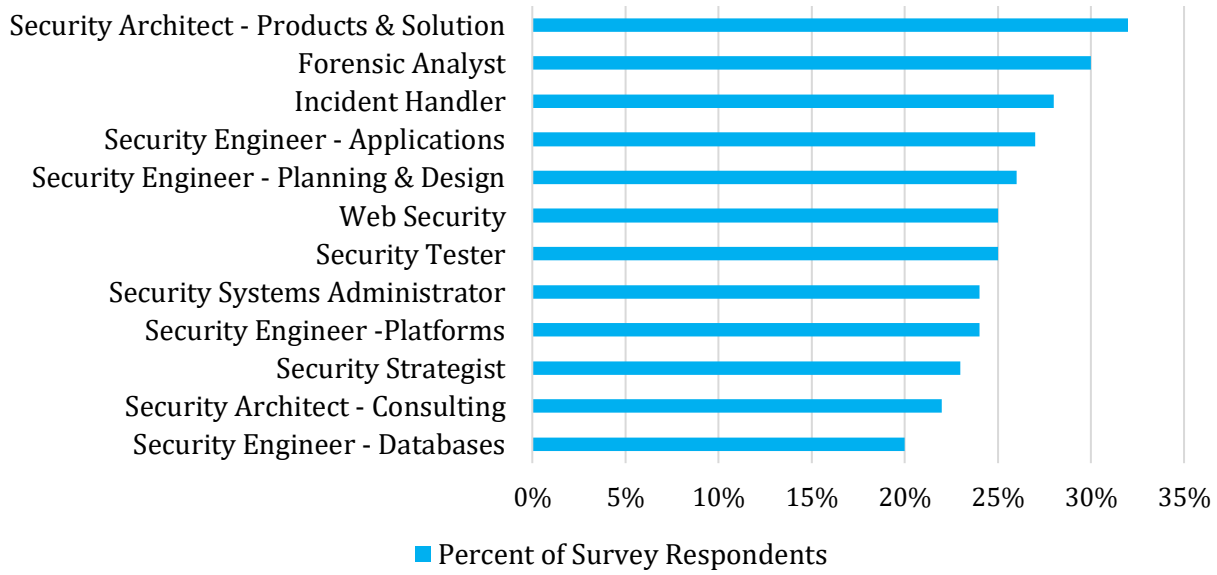

Figure 2. Security work roles demanded in 2015 GISWS [46]

To summarize, the cybersecurity industry is in its essence quite an old industry. Due to increased internet connectivity and an evolving threat landscape, this industry has seen rapid growth over the last years. More companies demand adequate protection of their IT infrastructure, people and information in the digital space. This has led to a vast expanse in work roles and the required knowledge, skills and abilities needed. The demand for the qualified cybersecurity professionals has significantly outpaced the supply.

2.2 Approach

The NIST SP800-181 standard is published in PDF and is accompanied by an Excel spreadsheet containing a Task Master list, a KSA Master list and Tasks and KSAs for each of the 52 roles. The structure of the framework in the standard is summarily explained here. The framework

contains seven top-level domains. These domains are 1) Securily Provision, 2) Operate and Maintain, 3) Oversee and Govern, 4) Protect and Defend, 5) Analyze, 6) Collect and Operate, and 7) Investigate. The layer below the domains are the specialties. There are 33 specialties in NIST SP800-181 with an n-to-1 relationship to domains, and Table 1 depicts them in relation to their parent domain nodes.

Table 1. Overview of the 33 NIST SP800-181 specialties [28]

Securily Provision
1) Risk Management, 2) Software Development, 3) Systems Architecture, 4)Technology R&D, 5) Systems Requirements Planning, 6) Test and Evaluation, 7) Systems Development
Operate and Maintain
1) Data Administration, 2) Knowledge Management, 3) Customer Service and Technical Support, 4) Network Services, 5) Systems Administration, 6) Systems Analysis
Oversee and Govern
1) Legal Advice and Advocacy, 2) Training, Education and Awareness, 3) Cybersecurity Management, 4) Strategic Planning and Policy, 5) Executive Cyber Leadership, 6) Program/Project Management and Acquisition
Protect and Defend
1) Cybersecurity Defense Analysis, 2) Cybersecurity Defense Infrastructure Support, 3) Incident Response, 4) Vulnerability Assessment and Management
Analyze
1) Threat Analysis, 2) Exploitation Analysis, 3) All-Source Analysis, 4) Targets, 5) Language Analysis
Collect and Operate
1) Collection Operations, 2) Cyber Operational Planning, 3) Cyber Operations
Investigate
1) Cyber Investigation, 2) Digital Forensics

The layer below that consists of work roles, which have an n-to-1 relationship to specialties. There are 52 work roles in NIST SP800-181. For the full list of roles and how they relate to specialties and domains, please refer to Appendix A – Specialties and Work Roles.

The lowest layer in the framework are tasks, knowledge, skills and abilities, which have an n-to-n relationship to work roles. There are 999 task nodes, 587 knowledge nodes, 365 skill nodes and 175 ability nodes. 8 Tasks, 43 knowledge elements, 9 skills and 1 ability were removed because they were present in an earlier standard version and were withdrawn or merged with another element in the NIST SP800-181 or duplicates. Knowledge, skills and tasks are required to perform tasks. This

entails that if artificial intelligence can perform a task, it implicitly also has the knowledge, skills and abilities that are needed for that task. They will therefore not be evaluated separately.

The suggested limitation for the locality and nature of the data as mentioned in 1.1 is supported by various tasks in the standard. Some of the tasks point to specific practices of the American Department of Defense (e.g. the Vulnerability Equity Process (VEP)) and reference American standards (e.g. Federal Information Processing Standards (FIPS)). In one case, a specific vendor supporting the DoD in vulnerability assessments (Blue Force) is referred to. Lastly, the specialty 'Targets' contains mostly offensive rather than defensive tasks. These tasks would be uncommon for most white hat cybersecurity companies. The fact that these specific instances are referred to is not likely to impact the assessment of whether artificial intelligence cannot or can do such a task.

2.3 Results

This section answers the sub-question on what tasks cybersecurity work consists of. Its focus is on tasks, but knowledge, skills and abilities are included in the analysis to provide context.

All the nodes were encoded in a graph database to better understand the relations between tasks, knowledge, skills and abilities, and work roles. A first query on the graph database reveals that not all nodes are connected. 140 nodes in total, of which 90 tasks out of 999 tasks, do not hold a relationship to a specific work role. It does not make sense to have tasks in the framework when they are not to be executed by a role. The researcher has reached out to NIST for an explanation. Their answer was that the inclusion of these elements was not a mistake. They were submitted to the NIST committee for inclusion in the standard, but the committee did not manage to build consensus to which work roles they belonged on time. It is likely that they will be matched up to work roles in a future edition of the NIST SP800-181 and new work roles may be created to accommodate this ambition [47]. The average number of connections to work role nodes are presented in Table 2. The number of tasks per security domain of NIST SP800-181 is depicted in Figure 3, based on the tasks allocated to the roles of each of the domains. Tasks may be double for roles within a domain, and may feature in various domains. Oversee and Govern is much larger than the other domains (374 tasks). Investigate (73 tasks) and Protect and Defend (68 tasks) are much smaller than the other domains.

Table 2. Average number of connections to role nodes

	Average (without 0s)	Average (with 0s)
Tasks	1.45	1.32
Knowledge	3.70	3.60
Skills	1.98	1.95
Abilities	3.03	2.51

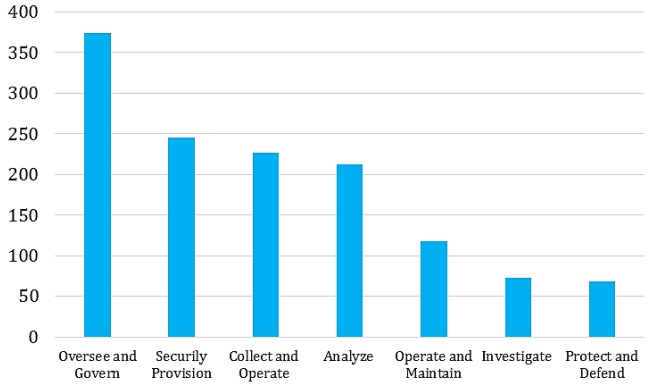

Figure 3. Number of tasks per domain

The full results of the frequency analysis are depicted in Table 3. A noteworthy feature is that tasks are only linked up to six work roles at maximum. Other elements can be linked up to 14 (skills), 15 (abilities) or 52 (knowledge) work roles. Six of the knowledge nodes are linked to all the work roles. This is of interest because it would entail that anyone working in the cybersecurity field ought to possess this knowledge.

Table 3. Frequency table for number of connections to role nodes

Number of connections to distinct roles	0	1	2	3	4	5	6	7	8	9	10
Tasks	90	624	204	52	19	7	3	0	0	0	0
Knowledge	15	194	163	63	29	18	25	14	18	7	9
Skills	5	174	110	41	11	12	5	2	3	1	0
Abilities	30	53	31	23	9	6	9	5	1	3	0
Number of connections to distinct roles (cont.)	**11**	**12**	**13**	**14**	**15**	**16**	**17**	**18**	**19**	**20**	**52**
Tasks	0	0	0	0	0	0	0	0	0	0	0
Knowledge	9	4	3	4	1	1	2	1	1	0	6
Skills	0	0	0	1	0	0	0	0	0	0	0
Abilities	1	1	1	1	1	0	0	0	0	0	0

These six knowledge elements are:

- "Knowledge of laws, regulations, policies, and ethics as they relate to cybersecurity and privacy.
- Knowledge of specific operational impacts of cybersecurity lapses.
- Knowledge of cyber threats and vulnerabilities.
- Knowledge of risk management processes (e.g., methods for assessing and mitigating risk).
- Knowledge of computer networking concepts and protocols, and network security methodologies.
- Knowledge of cybersecurity and privacy principles." [28]

All of these knowledge elements are quite generic tasks and seem to represent the background knowledge on risks, methods of being breached and impacts that one should have. Another outlier is the skill that has 14 connections to role nodes. This is the skill "…to apply cybersecurity and privacy principles to organizational requirements (relevant to confidentiality, integrity, availability, authentication, non-repudiation)"[28]. This indeed seems a skill that could apply to a broad range of work roles.

A query into the number of tasks per role reveals a large difference between roles, with a Privacy Officer having 73 tasks versus an Authorizing Official only having four tasks as depicted in

Figure 4 on page 18. On average work roles have 25 tasks. It is noteworthy that four out of the five work roles with the most tasks are 'managers'. Many of their tasks seem comprised of receiving and processing information from various categories of sub-ordinates. The reason the Privacy Officer has a large amount of tasks seems to be three-fold. First, some tasks are the same to security tasks, but contain the word 'privacy' rather than 'security'. An example would be "Conduct on-going privacy training and awareness activities" [28] whereas for security there is an entire role

(Cyber Instructor) dedicated to doing the same for cybersecurity. Second, many tasks evolve around collaboration with various security work roles. An example would be: "Collaborate with cybersecurity personnel on the security risk assessment process to address privacy compliance and risk mitigation" [28]. These tasks are replicated for various security specialties, albeit that the collaboration varies in nature. Third, it seems the Privacy Officer has strategic, tactical and operational tasks, whereas these are usually split out over different work roles for security (e.g. Executive Cyber Leadership and Cyber Operator). Examples of a strategic task for the Privacy Officer is "Serve in a leadership role for Privacy Oversight Committee activities" and an operational task for the Privacy Officer is "Interpret patterns of noncompliance to determine their impact on levels of risk and/or overall effectiveness of the enterprise's cybersecurity program" [28]. The Authorizing Official with only four tasks on the other hand has a very narrow scope of work in which he/she signs off on risk for systems.

2.4 Summary

This chapter has investigated the cybersecurity market and its labour requirements. The NIST SP800-181 standard has been used to analyse what kind of work roles, tasks and KSAs there are in the cybersecurity field. There is a large variety in how many tasks the various roles have. Privacy Officers have significantly more tasks than other roles according to NIST SP800-181. Authorizing officials have very few tasks. The next chapter investigates the current artificial intelligence ontology and applications. These can then be used to determine their usefulness for the cybersecurity market and its impact on the labour market.

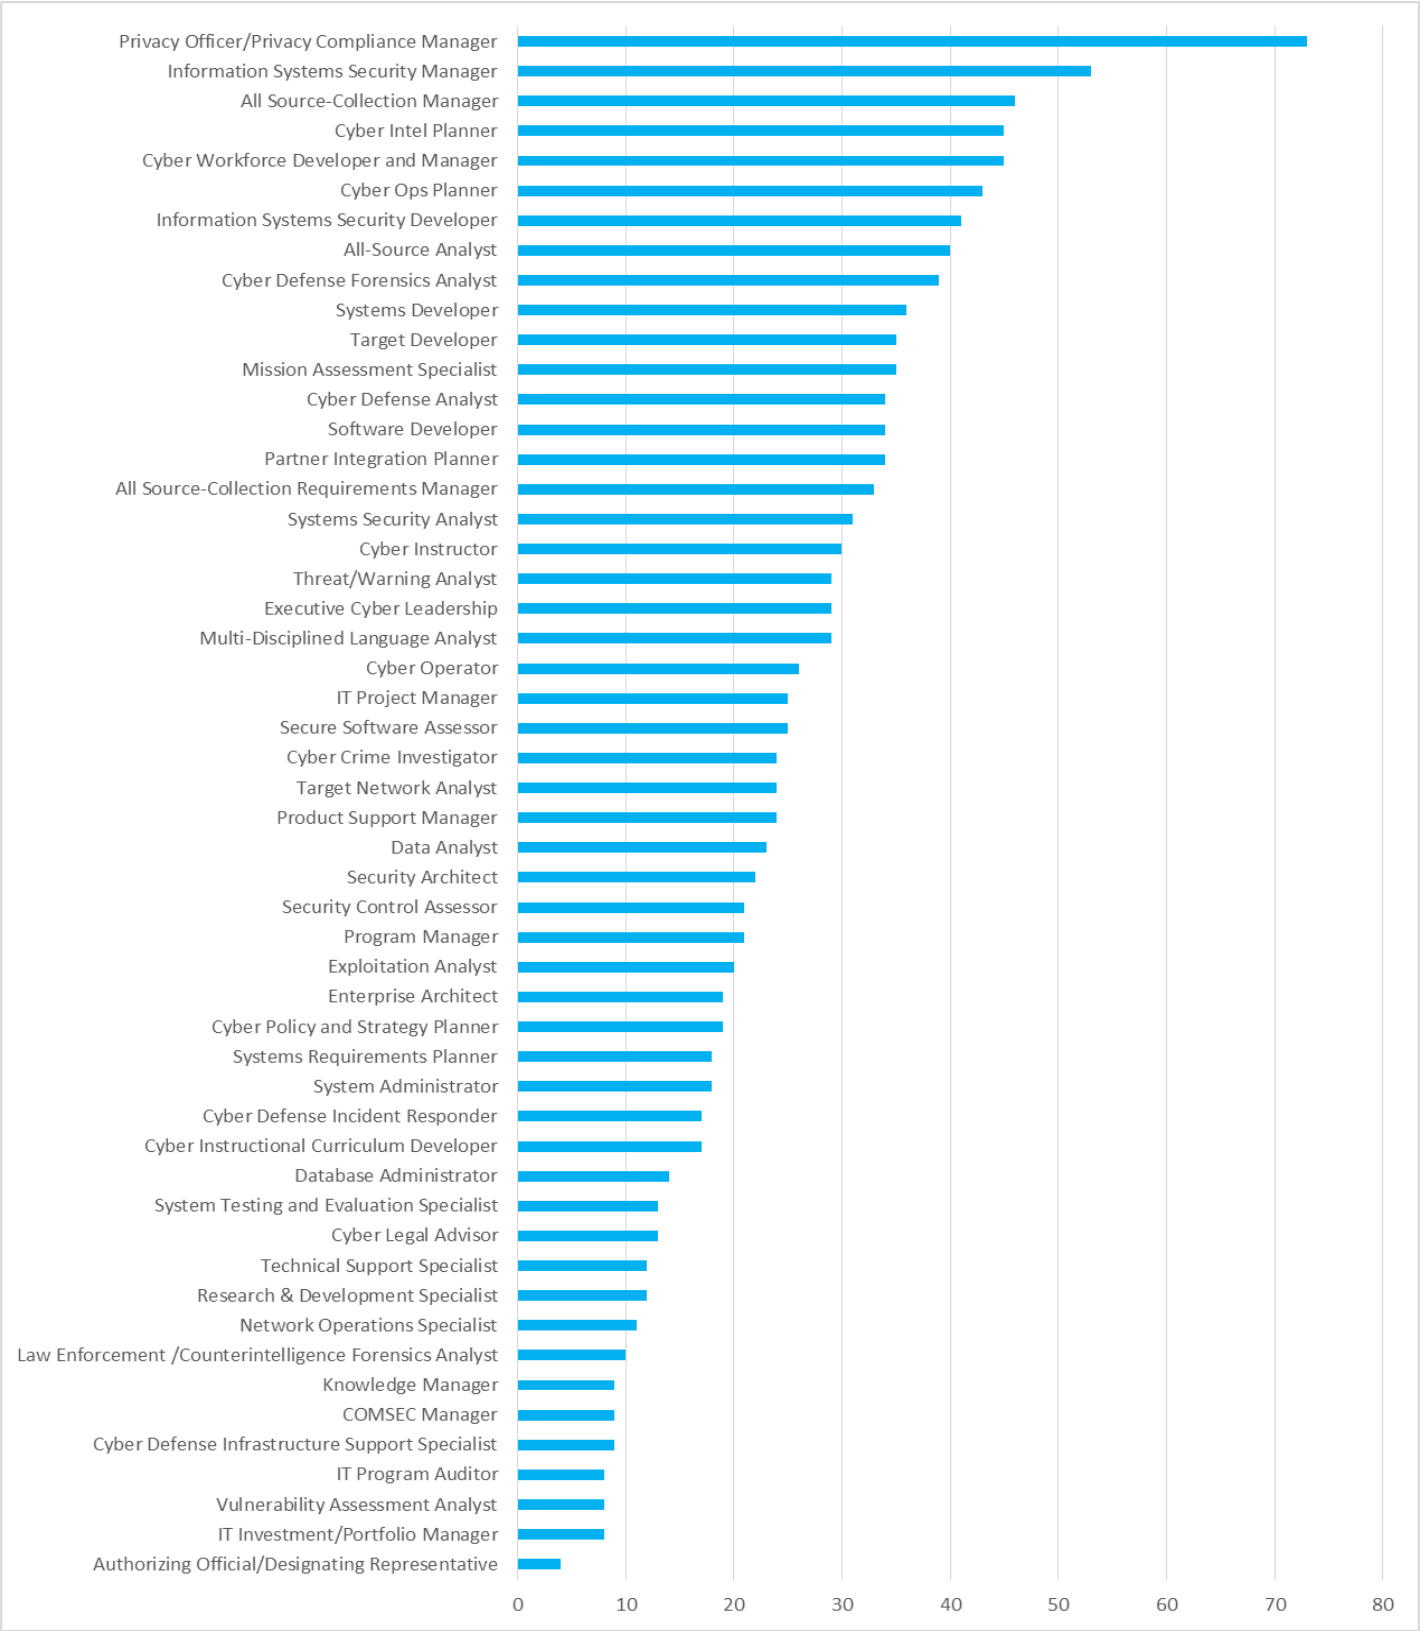

Figure 4. Division of tasks over roles

3. Artificial Intelligence Capabilities

This chapter documents the various current artificial intelligence competencies based on techniques and applications. The sub-question on what artificial intelligence can do based on current technology is answered through a literature study. The chapter starts with a definition of artificial intelligence. An introduction on artificial intelligence's origins and what it is capable of now are provided. The approach section takes this knowledge as input to determine a set of 'rules' on what artificial intelligence can and cannot do in terms of tasks. The results presented serve as input for sifting through the tasks of NIST SP800-181 to determine the proportion of outsourceable tasks in cybersecurity.

3.1 Background

This chapter looks into the competencies of artificial intelligence. A criticism of the term 'artificial' has been that it is inaccurate in the sense that just because it is not similar to human intelligence, it is therefore not 'real'. Nonetheless the term 'artificial intelligence' is used in this thesis as it is considered the industry standard. Two sub-fields within artificial intelligence are commonly discerned: Good old-fashioned artificial intelligence and computational intelligence.

Good old-fashioned artificial intelligence (GOFAI) (or 'Classical', 'Logical', 'Symbolic' or 'Conventional') refers to a specific sub-field of artificial intelligence. It uses intelligent agents with a knowledge base, rule sets and a learning capability to solve problems. It makes use of techniques as machine learning, deep learning, expert systems, search algorithms, (multi-)agent systems and natural language processing [48].

Computational intelligence (CI) is inspired by intelligent patterns in nature and mathematics and uses these patterns in multiple agents that solve problems through semi-stochastic behaviour [49]. It typically focusses on problems for which there is no effective computational algorithm, because problems are NP-hard or algorithms cannot be formulated, or because its learning, clustering, regression and classification techniques outperform GOFAI [50]. CI makes use of techniques as evolutionary computation algorithms and nature-inspired cooperative strategies such as genetic algorithms, artificial immune systems and particle swarm optimizations, artificial neural networks, statistic and probabilistic methods (e.g. Bayesian networks) and fuzzy logic [48].

These lists of AI & CI techniques are not exhaustive. In the interest of this thesis, artificial intelligence will be considered any combination of GOFAI and CI techniques required and feasible order to execute a cybersecurity task from NIST SP800-181.

3.2 Approach

This section first looks at artificial intelligence with different capability levels. It then lists the competencies currently associated with human intelligence. It looks at competencies currently associated with computational intelligence.

DARPA has created a model dividing artificial intelligence into 'three waves' of advancedness. Typically, older artificial intelligence systems belong to the first wave, newer systems to the third wave. Each of the waves has its own inventory of four high-level competencies: perceiving, learning, abstracting and reasoning. A short summary of each of the four waves and examples of applications is provided here.

The first wave is 'handcrafted knowledge' and has no learning and abstracting, a low level of perceiving and a high level of reasoning. Artificial intelligence in this wave reason with regards to narrowly defined problems, for which its developers have created a set of rules to represent knowledge. There is no learning capability and uncertainty is handled poorly by the system. Examples of these systems are planning tools, tax return programs and early cybersecurity network monitoring tools [51]. The second wave is 'statistical learning' and has a high level of perceiving and learning, and a low level of abstracting and reasoning. Artificial intelligence in this wave have classification and prediction capabilities, enabled by their statistical models for a specific problem domain and big data training. They lack an understanding of their context and reasoning capabilities. Examples of these systems are virtual assistants (e.g. Siri, Cortana), text analysis, image recognition and AlphaGo [51]. The third wave is 'contextual adaptation' and has a high level of perceiving, learning and reasoning, and a medium level of abstracting. Artificial intelligence in this wave learn as they encounter new tasks and situations, enabled by their explanatory models for classes of real-world scenarios. A challenge in this wave is to achieve natural communication among machines and people. Examples of these systems are self-driving cars, autonomous delivery robots and medical diagnostic assistants [51], [52].

There has been a historic tendency to have artificial intelligence equal or outperform human intelligence. Animalistic lifeforms display intelligence humans may not understand, but is intelligence nonetheless. It is only natural that computational systems have some intelligent competencies that humans are not capable of [53]. Machine intelligence has gone beyond just human intelligence and now also takes inspiration from nature. From a human perspective, other non-human competencies that artificial intelligence may develop are difficult to design and predict [54, p. 31]. Entirely different forms of intelligence springing from artificial intelligence themselves may develop over time, especially if the state of art approaches artificial general intelligence (AGI).

The core competency categories of artificial intelligence are provided by Russell and Norvig as problem-solving, search, logic, planning, knowledge representation, probabilistic reasoning decision making, learning, communicating, perceiving and acting [48]. Artificial intelligence is capable of performing many individual competencies that were once unique to humans. It would however take artificial general intelligence to master all competencies. In the 2009 AGI Roadmap Workshop a group of academics drafted a list of core human competencies. On the road to creating AGI, it can be expected that from individual competencies they will evolve to perform all competencies from one competency domain to go on and perform competencies from all competency domains. The competency list the Workshop drafted contained the domains and competencies depicted in Table 4 [54], [55].

Table 4. Human competencies for artificial intelligence (adapted from [54, p. 32])

Perception	Communi-cation	Modelling Self/Other	Social Interaction	Emotion	Reasoning	Memory
C1 Vision	**C8** Gestural	**C15** Self-awareness	**C22** Communicatio n	**C28** Perceived	**C34** Induction	**C40** Working
C2 Smell	**C9** Verbal	**C16** Other-awareness	**C23** Appropriaten ess	**C29** Expressed	**C35** Deduction	**C41** Episodic
C3 Touch	**C10** Musical	**C17** Relationships	**C24** Social inference	**C30** Control	**C36** Abduction	**C42** Implicit
C4 Taste	**C11** Pictorial	**C18** Self-control	**C25** Cooperation	**C31** Under-standing	**C37** Physical	**C43** Semantic
C5 Audition	**C12** Diagram-matical	**C19** Theory of mind	**C26** Competition	**C32** Sympathy	**C38** Causal	**C44** Procedural
C6 Cross-modal	**C13** Language acquisition	**C20** Sympathy	**C27** Relationships	**C33** Empathy	**C39** Associational	
C7 Proprio-ceptive	**C14** Cross-modal	**C21** Empathy				
Learning	**Motivation**	**Planning**	**Actuation**	**Attention**	**Building/Crea tion**	**Quantitative**
C45 Imitation	**C50** Sub-goal creation	**C54** Tactical	**C58** Physical skills	**C62** Visual	**C66** Physical construction with objects	**C70** Count observed entities
C46 Reinfor-cement	**C51** Affect-based	**C55** Strategic	**C59** Tool use	**C63** Auditory	**C67** Formation of novel concepts	**C71** Grounded small number arithmetic
C47 Dialogical	**C52** Deferred gratification	**C56** Physical	**C60** Navigation	**C64** Social	**C68** Verbal invention	**C72** Compare quantitative properties of observed entities
C48 Media-oriented	**C53** Altruism	**C57** Social	**C61** Proprio-ceptive	**C65** Behavioural	**C69** Social organisation	**C73** Measure using simple tools
C49 Experimental						

The literature on the competencies of CI techniques is less well-defined than that of artificial intelligence [50]. It is perhaps because this domain is further away from human frame of reference and is therefore less intuitive to define. Looking at individual techniques and how they are used, the CI competencies that can be discerned are depicted in Table 5.

Table 5. CI competencies for artificial intelligence (based on [50])

Ref.	Competency	Ref.	Competency
C74	Optimisation	C81	Pattern recognition
C75	Exploration	C82	Perception
C76	Classification	C83	Object recognition
C77	Clustering	C84	Signal analysis
C78	Correlation	C85	Navigation
C79	Regression	C86	Proprioceptive actuation
C80	Prediction		

The next section reflects on the competencies listed in this section and distils a set of rules by which the total set of tasks in NIST SP800-181 can be analysed to determine which tasks can be performed by artificial intelligence and which cannot.

3.3 Results

This section reflects on the artificial intelligence capabilities listed in the section above. It discusses research that attempted combining capabilities to perform certain tasks by looking at characteristics of the latter. It continues with an overview of how sub-tasks may be strung together to perform more complex tasks. The section looks into the various criteria of deciding whether to have humans or artificial intelligence perform certain tasks. It then reflects on the influence of the structure of artificial intelligence, the nature of the task and the environment on the success of executing the task. It ends with a list of rules that determine whether tasks can be performed with the given capabilities. This list is used in the next chapter to determine which security tasks can be performed.

Literature covering the kinds of tasks that artificial intelligence can and cannot do was investigated. Tasks are divided into three categories increasing in level of complexity: mundane tasks, formal tasks and expert tasks. Each of the task categories relies on the capabilities displayed in the less complex categories [56, p. 2], [57, pp. 1620–1621].

Mundane tasks require abilities that most humans have naturally been born with. They often have a physical element or make use of sensors that simulate the human five senses. Examples of mundane tasks are natural language (understanding, generation and translation), perception (vision and speech), robot control and common-sense reasoning [56, p. 2]. Sensor-reasoning interaction for artificial intelligence have developed insofar that many mundane tasks can be performed by AI. Common-sense reasoning is something at which humans often outperform artificial intelligence (but progress has been made in some areas, such as translation). Our ability to generalize and learn from our experiences cannot always be replicated for artificial intelligence [57, pp. 1620–1621]. Mundane tasks make use of mostly the human competencies defined in Table 4.

Formal tasks are tasks performed in relation to a well-defined problem and often rely on mathematical and logic operations to understand and define the problem space. Computers typically excel in performing formal tasks beyond human capacity [57, p. 1620]. Examples of formal tasks are game playing (chess, go, checkers and backgammon) and mathematics (integral calculus, geometry, logic and proving properties of programs) [56, p. 2]. Formal tasks make use of mostly the computational intelligence competencies defined in Table 5.

Expert tasks require a knowledge set that provides context for mundane skills and functions as an application framework for formal methods for complex problem solving. Although expert systems are proliferating, having artificial intelligence perform these tasks remains complicated. To exceed humans for many tasks a very large knowledge database and complex set of rules is needed. It depends on the application and how good humans are in that application relative to artificial intelligence, whether it makes sense to have artificial intelligence perform these tasks. There is usually a significant cost involved in terms of development and training time [57, pp. 1620–1621]. Examples of expert tasks are engineering (design, fault finding and manufacturing planning), scientific analysis, medical diagnosis and financial analysis [56, p. 2]. Expert tasks combine human and computational intelligence competencies as depicted in both Table 4 and Table 5.

For artificial intelligence to perform tasks, it is necessary to frame them as problems to solve. Formal problems consist of four elements: an initial state, a state space, a goal test and a path cost [56]. It is first necessary to define a state space in which the artificial intelligence is to search. It is to contain all relevant configurations to solving the problem. Then an initial state is identified as a starting position for the search. For some problems this will be a logical state (i.e. the location of a distribution warehouse when calculating a route to deliver a package to a customer). For others it may be (semi-)randomly chosen position (e.g. evolutionary computations can start from any position). The goal state(s) that would represent acceptable solutions need(s) to be defined. The rule set specifying the operations and under which conditions these may be performed, is to be identified. A goal test confirms after each operation whether a goal state has been reached. The path cost needs to be minimized to find the 'best' solution to the problem [56, pp. 17–18]. It may sound as if this 'problemification' of tasks only applies to search challenges. However, effectively artificial intelligence always has to search for the optimal solution, whether it concerns the fastest delivery route or finding the textual equivalent of a spoken word by a human in its database.

In the quest for artificial general intelligence various studies look into whether artificial intelligence can perform certain roles by stringing together a number of tasks. Those studies are relevant to this research as the cybersecurity roles of NIST SP800-181 also contain multiple tasks (and sometimes these tasks contain multiple sub-tasks). Puigbo et al. investigate the cognitive architecture of a general purpose service robot [58]. They distinguish the following separate mundane tasks of which their robot is capable, of which some are chained into a more extensive task (e.g. 'look for the nearest exit and exit the area'). This research links these tasks to the competencies C1 – C86 listed in Table 6 as a demonstration of how such tasks may be deconstructed into sub-tasks that artificial intelligence is competent in performing. Multiple competencies may be involved depending on the exact function actualisation in artificial intelligence. The core competencies are highlighted in bold.

Another example of such a study was executed by Poole, Mackworth & Goebel [59]. They outlined the tasks an autonomous delivery robot, a diagnostic assistant and an 'infobot' were to master in order to be successful in their role. Their line-up of tasks is interesting as the level of complexity is quite close to that of the level of complexity in the NIST SP800-181 tasks, certainly when compared to Puigbo et al.'s tasks. Furthermore, when abstracting these roles they are close to some of the work roles in NIST SP800-181 (e.g. a diagnostic assistant does not differ that much from an information security systems assessor apart from their Body of Knowledge). Tasks for the three artificial intelligence systems are listed in Table 7.

Table 6. Puigbo et al.'s tasks linked to competencies [58, p. 111]

Puigbo et al. robot tasks	Link to competencies
Navigate to a location	**Navigation, physical planning** Proprioceptive actuation and perception, optimisation
Introduce himself	**Conversation, self-awareness** Verbal communication, relationships, self-awareness
Follow a specific person in front of him	**Object recognition, navigation** Other-awareness, physical planning, imitation, prediction, visual perception, proprioceptive perception
Look for objects in front of him	**Object recognition, visual perception** Proprioceptive perception
Look for someone in the area	**Object recognition, visual perception** Proprioceptive perception
Grasp a specific object	**Physical actuation, physical planning**
Deliver an object to the person	**Physical actuation, navigation**
Memorize a person's face and name	**Object recognition, visual perception, memory** Pattern recognition
Look for the nearest exit and exit the area	**Object recognition, visual perception, navigation** Proprioceptive perception, physical planning
Check the person in front as already known and retrieve the person's name	**Object recognition, visual perception** Pattern recognition, memory
Point to the location of a specific object	**Object recognition, physical reasoning, physical actuation** Visual perception

Table 7. Poole et al.'s tasks for three artificial intelligence systems [59, pp. 13–17]

Delivery robot	Diagnostic assistant	Infobot
1. Determine where individuals' offices are, where to get coffee, how to estimate the length of a trip, and so on.	1. Derive the effects of faults and interventions.	1. Derive information that is only implicit in the knowledge base(s), as well as interact in natural language.
2. Find a path between different locations and optimize this path for performance constraints.	2. Search through the space of possible faults or disease complexes.	2. Search through a variety of knowledge bases looking for relevant information.
3. Be able to represent knowledge about the domain so that inference can be quick, so that knowledge can be easily acquired, and so that the appropriate knowledge is represented.	3. Explain its reasoning to the human who is using it.	3. Find good representations of knowledge so that answers can be computed efficiently.
4. Plan how to carry out multiple goals, even when they use the same resources.	4. Derive possible causes for symptoms; rule out other causes based on the symptoms.	4. Explain how an answer was derived or why some information was unavailable.
5. Make default assumptions —for example, about where people will be.	5. Plan courses of tests and treatments to address the problems.	5. Make conclusions about lack of knowledge, determine conflicting knowledge, and be able to conclude disjunctive knowledge.
6. Make trade-offs about plans even though there may be uncertainty about what is in the world and about the outcome of its actions.	6. Hypothesize problems and use default knowledge that may not always be true.	6. Use default reasoning about where to obtain different information.
7. Learn about features of its domain, as well as learn about how its actions affect its position and its rewards.	7. Reason about the uncertainties about the artefact given only partial information about the state of the artefact, the uncertainty about the effects of the treatments, and the trade-offs between the alternate courses of action.	7. Make trade-offs between cheap but unreliable information sources and more expensive but more comprehensive information sources.
8. Sense the world, know where it is, steer around the corridors (avoiding people and other objects), and pick up and put down objects.	8. Learn about what symptoms are associated with the faults or diseases, the effects of treatments, and the accuracy of tests.	8. Learn about what knowledge is available where, and what information the user is interested in.

Adams et al. suggest the possibility of stringing together agents with specific competencies to complete a complex task ("big switch statement") [54]. They furthermore posit that matching artificial intelligence capabilities to tasks would benefit from grouping tasks into 'task families' by looking at the domains of tasks as outlined in Table 4 (e.g. Perception, Communication, Emotion etc.). A second layer of granularity could be to cluster tasks on the complexity of the tasks (i.e. mundane, formal, expert) [54, p. 38]. Adams et al. created scenarios (which for the purpose of this research could be seen as a 'work role') and identify various tasks in these scenarios. In these scenarios the agent is a child that grows up and moves through various stages of cognitive development to adulthood. Each of the scenarios comes with a set of tasks and human competencies to be developed. A concrete example task is "While Sam is in the room, Ben puts the red ball in the red box. Then Sam leaves and Ben moves the red ball to the blue box. Sam returns and Ben asks him where the red ball is. The agent is asked where Sam thinks the ball is" [54, p. 39]. This task belongs to the scenario 'Virtual Preschool' and has been mapped to the competency area of 'Modeling Self and Other' and the specific competency 'Theory of Mind'. Although the tasks outlined in the NIST SP800-181 are not as specific as these tasks, it does illustrate how they may be grouped into task families (e.g. Modeling Self and Other) that require similar competencies.

Shahaf and Horvitz investigate human-machine interaction in task markets [60]. Their goal is find the optimal distribution of tasks from a pool of tasks between the different actors in a generalized task market. Given the nature of a task, they determine whether the optimal set-up is skilled human only, semi-skilled human collaboration, human-machine collaboration or machine only. They take into account the availability of the actors, the competencies and preferences of the actors, and the price of actors to solve the problem (resource cost, time required, task performance quality). They state that every high-level tasks can be broken down into one or multiple low-level tasks. An interesting observation is that whether machines can execute tasks is not a matter of whether it can be done or not. Many low-level tasks can be performed by machines - their performance quality however may differ from the desired performance. When stringing these lower-level tasks together into its high-level task, the resulting quality may be so poor that we consider the machine to 'not be able' to perform the high-level task. The low-level tasks and high-level tasks can be equated in many cases to the mundane tasks and expert tasks discussed earlier.

A last consideration is that it depends on the structure of the artificial intelligence system, the task and the environment in which the task is to be executed whether a task can or cannot be done by artificial intelligence [61, p. 1]. Laird and Wray posited environment, task and agent characteristics and architecture requirements for artificial general intelligence to succeed, as depicted in Table 8 respectively.

Table 8. Environment, task and agent characteristics (based on [61])

Ref.	Characteristic	Explanation
CH1	The environment is complex, with diverse, interacting and richly structured objects.	The agent must be capable of recognising different objects in the environment and understanding how these may impact its performance of the task.
CH2	The environment is dynamic and open.	The environment can be in different states that impact the execution of the task, and the agent must be able to respond to these different states.
CH3	Task-relevant regularities exist at multiple time scales.	The environment is governed by systemic rules that make changes in state predictable.
CH4	Other agents impact performance.	Other agents may be present that aid or hinder the execution of the task. Additionally, the agent may learn from other agents performing the task.
CH5	Tasks can be complex, diverse and novel.	Tasks may come in various forms for which the agent needs to adapt its execution.
CH6	Interactions between agent, environment and tasks are complex and limited.	Agents can discern a limited range of changes in the environment and assess how it may impact the execution of the task.
CH7	Computational resources of the agent are limited.	The agent experiences bounded rationality and cannot search for a solution indefinitely.
CH8	Agent existence is long-term and continual.	The agent is to assume it will execute tasks for an extended period and prepare itself mentally and physically for continuous work.

Based on the literature studied on the types of tasks available in the sections above, the lessons learned were abstracted into general statements that should be true for artificial intelligence to perform a task. The various characteristics of the agent, the tasks and the environment, and the need to combine tasks to successfully execute an expert task and work role were deemed important enough to each get a separate rule. Initially rules contained specific verbs to search for in the NIST SP800-181 tasks. These were removed (now only listed as examples in some other rules) because the use of these verbs was too inconsistent to provide a reasonable means of selection. A manually executed test run for the task analysis revealed the initial rule list required adaptation in some of the wording. The following rules have been used to determine whether the task can be performed by artificial intelligence. The rules are numbered because they are referred to later during the task analysis; they do not indicate a ranking of importance.

Rule set:

1. A task must have a definable state space, an initial state, a goal state and a rule set for searching for a suitable solution. If it does not have one of these elements, a task cannot or only partially be outsourced to artificial intelligence. (based on [56])
2. If a task description contains multiple verbs, it requires the artificial intelligence to perform multiple competencies. (based on [58])
3. For mundane tasks, a task that requires common-sense reasoning is unlikely to be effectively outsourced to artificial intelligence. Verbs associated with common-sense reasoning appearing in the task description (e.g. 'comprehend', 'understand', 'determine') will thus be considered tasks not outsourceable to artificial intelligence, with exception of fields with sufficient progress (e.g. 'translate'). (based on [57])
4. Mundane tasks that require competencies C1- C73 and use verbs associated with those competencies in the task description, can be outsourced to artificial intelligence (e.g. 'select', 'assess', 'identify', 'navigate', 'translate', 'document'). (based on [57])
5. Formal tasks involving logic or mathematical operations can be outsourced to artificial intelligence. Verbs associated with formal tasks appearing in the task description (e.g. 'calculate', 'categorize', 'analyze', 'test', 'verify', 'reason', 'monitor', 'evaluate', 'search'). (based on [57])
6. Expert tasks consist of multiple sub-tasks and can be outsourced to artificial intelligence as long as it consists of stringing together various mundane and formal competencies. (based on [54], [57], [59])
7. When having the cost of having artificial intelligence perform the task to acceptable quality is higher than a human performing the task to acceptable quality, the task will be considered non-outsourceable (this may be different from the real-life industry, as there is such lack of human resource availability it might be necessary to have artificial intelligence execute tasks sub-par [62]). (based on [60])
8. Individual instances of a task follow a consistent pattern for the task to be outsourceable to artificial intelligence. (based on [61])
9. There is limited impact from variations in the environment and other agents on the performance of the task for the task to be outsourceable to artificial intelligence. (based on [61])
10. As artificial intelligence experience bounded rationality, the task must have a limited number of variables and a limited solution space. The larger the solution space and the more time is required to search this space, the higher the cost of the outsourcing to artificial intelligence (based on [60], [61])

3.4 Summary

This chapter was geared to providing examples of how tasks were framed in relation to artificial intelligence by other researchers. Their insights and formulations were used to draft a list of ten rules. In the next chapter these ten rules are used to determine which tasks can and which tasks cannot be outsourced to artificial intelligence from NIST SP800-181.

4. Outsourcing Cybersecurity Tasks

This chapter takes the rule set produced in the previous chapter and applies it to the tasks of NIST SP800-181. First, it looks into the current applications of artificial intelligence per NIST SP800-181 cybersecurity domain. Second, the rules of the previous chapter are adapted into criteria for whether a task can, can partially or cannot be outsourced to artificial intelligence. Third, it evaluates the results of the tasks being split into these categories. The evaluation consists of which reasons in which proportions determined the classification into the 'yes', 'no' or 'partial'-categories, the distribution of these categories over NIST SP800-181 work roles and the distribution over NIST SP800-181 cybersecurity specialty areas and domains.

4.1 Background

Before determining which tasks can be outsourced based on the rule set created in the previous chapter, this section presents an overview of successful current applications in cybersecurity. A literature review is performed by reviewing academic research involving various artificial intelligence techniques that have been operationalised for cybersecurity. The results are mapped back to the NIST SP800-181 domains and specialties and depicted in Table 9.

It is natural for those categories showing many current applications in artificial intelligence, to presume many tasks can be outsourced based on the rule set analysis as well. From Table 9 it becomes apparent that some domains have more extensive applications than others. A very large number of papers (not all included here) was found on various ways to use artificial intelligence techniques for Intrusion Detection Systems. This seems to be the major application of artificial intelligence in cybersecurity at this moment. Overall, the domains 'Protect & Defend' and 'Analyze' seem to have the most applications. Not many results were found for the category 'Operate and Maintain', perhaps because this category contains roles and tasks that can be seen as generic IT, rather than cybersecurity-specific. Lastly, it is interesting to note that the 'Oversee and Govern' domain only has supportive applications. This suggests that although artificial intelligence may aid roles in this category, they cannot perform them independently of humans.

Table 9. Current AI applications in security

NIST SP800 - 181 Domain	NIST SP800-181 Specialty	Usage found
Securily Provision	Risk Management	Decision support [62], [63], visual analytics [64], risk assessment support [65]
	Software Development	Secure software verification proof [66]
	Systems Architecture	Secure network architecture [67]
	Technology R&D	*None found*
	Systems Requirements Planning	*None found*
	Test and Evaluation	Secure system and software verification proof [66], recommend testing strategies [69]
	Systems Development	Secure system verification proof [66]
Operate & Maintain	Data Administration	*None found*
	Knowledge Management	*None found*
	Customer Service and Technical Support	*None found*
	Network Services	*None found*
	Systems Administration	Behavioural biometrics [68]
	Systems Analysis	Online system analysis [66]
Oversee & Govern	Legal Advice and Advocacy	*None found*
	Training, Education, and Awareness	Situational awareness [63]
	Cybersecurity Management	Decision support [62], [63], visual analytics [64]
	Strategic Planning and Policy	Security remediation planning [63], decision support [63]
	Executive Cyber Leadership	Decision support [62], [63]
	Program/Project Management and Acquisition	Decision support [62], [63], audit log and trail analysis [66]
Protect & Defend	Cybersecurity Defense Analysis	Intrusion prevention and detection systems [66], [69], [70], phishing & spam prevention [71], DDoS mitigation [63], anti-virus and anti-malware solution [67], botnet mitigation [71]
	Cybersecurity Defense Infrastructure Support	Intrusion prevention and detection systems [66], [69], [70], phishing & spam prevention [71], DDoS mitigation [63]
	Incident Response	DDoS mitigation [63]
	Vulnerability Assessment and Management	Attack planning [72], vulnerability analysis [73]
Analyze	Threat Analysis	Attack planning [72], situational awareness [63], cyber terrorism threat intelligence [71], visual analytics [64], anomalous behaviour detection [79]
	Exploitation Analysis	Attack planning [72], situational awareness [63]
	All-Source Analysis	Threat intelligence repository [69], situational awareness [63], cyber terrorism threat intelligence [71], visual analytics [64]
	Targets	Threat intelligence repository [69], situational awareness [63]
	Language Analysis	Situational awareness [63]
Collect & Operate	Collection Operations	Situational awareness [63]
	Cyber Operational Planning	Cybersecurity operations planning [63]
	Cyber Operations	Situational awareness [63]
	Cyber Investigation	Audit log analysis [66]
	Digital Forensics	Audit trail analysis [73], data carving [74]

4.2 Approach

As studied in the previous chapter, there are various ways of understanding tasks in relation to artificial intelligence. This section takes those ten rules and sorts them into various reasons why a task can (and in case of Rule 7, should), can partially or cannot be outsourced to artificial intelligence.

Reasons for a task being outsourceable are:

- Problem can be defined (Rule 1)
- Mundane human task (Rule 4)
- Uses CI competencies (Rule 4) / Formal task (Rule 5)
- All sub-tasks in the expert task are mundane or formal (Rule 6)
- Artificial intelligence can perform the task to an acceptable level quicker than a human (Rule 7)
- Artificial intelligence can perform the task to an acceptable level cheaper than a human (Rule 7)
- Artificial intelligence can perform the task better than a human (Rule 7)
- Task instances follow a highly consistent structure (Rule 8)
- Artificial intelligence can correct for the influence of the environment and other agents (Rule 9)
- The number of variables and the solution space of the task is limited (Rule 10)

Reasons for a task being partially outsourceable are:

- Problem can be defined for only part of the task (Rule 1)
- One of multiple verbs is not associated with mundane or formal tasks (Rule 2, 4, 5, 6)
- Common-sense reasoning with a poorly definable knowledge base or rule set is required for part of the task (Rule 3)
- Task instances have some consistent structural elements (Rule 8)
- Artificial intelligence needs input from another agent or the environment to perform the task (Rule 9)

Reasons for a task not being outsourceable are:

- No definable state space can be defined (Rule 1)
- No initial state can be defined (Rule 1)
- No goal state can be defined (Rule 1)
- No rule set can be defined (Rule 1)
- The expert task cannot be split into concrete mundane or formal sub-tasks (Rule 2, 3, 4, 5, 6)
- Humans can do the task better and/or cheaper and/or quicker (Rule 7)
- Task instances do not follow a consistent structure (Rule 8)
- The environment and other agents have an impact on the task or the artificial intelligence that cannot be corrected for by the artificial intelligence (Rule 9)
- The task has a very high number of variables or a large solution space, such that searching an appropriate solution to the task is unfeasible (Rule 10)

The analysis was conducted by reviewing each of the tasks manually and analysing the underlying sub-tasks. A template was created with the 999 NIST SP800-181 tasks, a 'verdict' column and three columns prepopulated with a list of potential reasons for a task being put into a category. Each category contains the specific reasons applying to that category. A screenshot of the template is provided for the reader's convenience in Figure 5 on page 33.

The tasks were each evaluated on whether they are fully, partially or not outsourceable to artificial intelligence by looking at the ten rules identified in the previous chapter. If a task was deemed to be feasible for artificial intelligence, the verdict 'Yes' was selected and the 'Yes' reasons that were not key to the task falling into this category were removed from the template. The same process was followed for the other categories. For example, T0001 was rated as a 'No'. Although multiple reasons were applicable for why the task could not be completed by artificial intelligence, a reason like 'No goal state can be defined (Rule 1)' was removed as it would be possible for us to specify and define the goal of the task ("necessary resources") to the AI, e.g. 'Obtain 1 million dollars'. The reason columns from other categories were emptied. An example of what the template looked like once verdicts were allocated and appropriate reasons provided, is depicted in Figure 6 on page 33 for the reader's convenience.

Each of the reasons falls into one of two more abstract categories. The first category consists of reasons why artificial intelligence *can or cannot* do a certain task, e.g. because we cannot define a formal problem to give to the artificial intelligence to solve, or because the techniques for solving the problem have not been invented yet. The second category of reasons consists of reasons related to whether we would *want* artificial intelligence to do a certain task, e.g. because it can do it cheaper (requires less training and incentives than its human counterpart, a reason most logical for requiring highly specialized and experienced human labour), quicker or simply better. Some tasks were in theory feasible for artificial intelligence, but did not make sense to outsource as our innate human abilities are so much better and cheaper (e.g. for the task "Consult with customers to evaluate functional requirements"[28]). The other way around occurred as well – e.g. for the task "Correlate incident data to identify specific vulnerabilities and make recommendations that enable expeditious remediation"[28] it is clear how if properly designed artificial intelligence could correlate data, identify vulnerabilities and provide recommendations better, quicker and cheaper than humans.

4.3 Results

Analysing the tasks on whether or not artificial intelligence would be able to fully, partially or not perform them delivered the results depicted in Table 10 - Table 13. Few tasks were simply mundane or formal (45 cases). Most tasks were expert tasks with many sub-tasks, and it was only feasible to outsource such a complex task (176 cases) to artificial intelligence in case all of its sub-tasks were deemed mundane or formal. In case an expert task was deemed to make particular use of a CI competency or human capability (eight cases) this was added as a separate reason. Reasons deemed key for the task being fully, partially or not outsourceable were listed for each of the tasks. The statistics are presented in Table 11, Table 12 and Table 13 with N=999. The more reasons were listed for a single task, the stronger the evidence was that the task belonged to that category.

Table 10. Division of tasks over 'Yes', 'Partial' and 'No' categories

Category	Frequency	Percentage of total number of tasks	Average number of reasons
Yes	221	**22.1%**	4.88 (out of 10)
Partial	370	**37.1%**	2.70 (out of 5)
No	407	**40.7%**	3.79 (out of 9)

Task ID	Task Description	Verdict	Reason Yes	Reason No	Reason Partial
T0001	Acquire and manage the necessary resources, including leadership support, financial resources, and key security personnel, to support information technology (IT) security goals and objectives and reduce overall organizational risk.	Yes/ Partial/ No	- Problem can be defined (Rule 1) - Mundane human task (Rule 4) - Uses CI competencies (Rule 4) / Formal task (Rule 5) - All sub-tasks in the expert task are mundane or formal (Rule 6) - Artificial intelligence can perform the task to an acceptable level quicker than a human (Rule 7) - Artificial intelligence can perform the task to an acceptable level cheaper than a human (Rule 7) - Artificial intelligence can perform the task better than a human (Rule 7) - Task instances follow a highly consistent structure (Rule 8) - Artificial intelligence can correct for the influence of the environment and other agents (Rule 9) - The number of variables and the solution space of the task is limited (Rule 10)	- No definable state space can be defined (Rule 1) - No initial state can be defined (Rule 1) - No goal state can be defined (Rule 1) - No rule set can be defined (Rule 1) - The expert task cannot be split into concrete mundane or formal sub-tasks (Rule 2, 3, 4, 5, 6) - Humans can do the task better and/or cheaper and/or quicker (Rule 7) - Task instances do not follow a consistent structure (Rule 8) - The environment and other agents have an impact on the task or the artificial intelligence that cannot be corrected for by the artificial intelligence (Rule 9) - The task has a very high number of variables or a large solution space, such that searching an appropriate solution to the task is unfeasible (Rule 10)	- Problem can be defined for only part of the task (Rule 1) - One of multiple verbs is not associated with mundane or formal tasks (Rule 2, 4, 5, 6) - Common-sense reasoning with a poorly definable knowledge base or rule set is required for part of the task (Rule 3) - Task instances have some consistent structural elements (Rule 8) - Artificial intelligence needs input from another agent or the environment to perform the task (Rule 9)
T0002	Acquire necessary resources, including financial resources, to conduct an effective enterprise continuity of operations program.	Yes/ Partial/ No	- Problem can be defined (Rule 1) - Mundane human task (Rule 4) - Uses CI competencies (Rule 4) / Formal task (Rule 5) - All sub-tasks in the expert task are mundane or formal (Rule 6) - Artificial intelligence can perform the task to an acceptable level quicker than a human (Rule 7) - Artificial intelligence can perform the task to an acceptable level cheaper than a human (Rule 7) - Artificial intelligence can perform the task better than a human (Rule 7) - Task instances follow a highly consistent structure (Rule 8) - Artificial intelligence can correct for the influence of the environment and other agents (Rule 9) - The number of variables and the solution space of the task is limited (Rule 10)	- No definable state space can be defined (Rule 1) - No initial state can be defined (Rule 1) - No goal state can be defined (Rule 1) - No rule set can be defined (Rule 1) - The expert task cannot be split into concrete mundane or formal sub-tasks (Rule 2, 3, 4, 5, 6) - Humans can do the task better and/or cheaper and/or quicker (Rule 7) - Task instances do not follow a consistent structure (Rule 8) - The environment and other agents have an impact on the task or the artificial intelligence that cannot be corrected for by the artificial intelligence (Rule 9) - The task has a very high number of variables or a large solution space, such that searching an appropriate solution to the task is unfeasible (Rule 10)	- Problem can be defined for only part of the task (Rule 1) - One of multiple verbs is not associated with mundane or formal tasks (Rule 2, 4, 5, 6) - Common-sense reasoning with a poorly definable knowledge base or rule set is required for part of the task (Rule 3) - Task instances have some consistent structural elements (Rule 8) - Artificial intelligence needs input from another agent or the environment to perform the task (Rule 9)

Figure 5. Task analysis template (partial view)

Task ID	Task Description	Verdict	Reason Yes	Reason No	Reason Partial
T0001	Acquire and manage the necessary resources, including leadership support, financial resources, and key security personnel, to support information technology (IT) security goals and objectives and reduce overall organizational risk.	No		- No definable state space can be defined (Rule 1) - No rule set can be defined (Rule 1) - Humans can do the task better and/or cheaper and/or quicker (Rule 7) The environment and other agents have an impact on the task or the artificial intelligence that cannot be corrected for by the artificial intelligence (Rule 9) - The task has a very high number of variables or a large solution space, such that searching an appropriate solution to the task is unfeasible (Rule 10)	
T0002	Acquire necessary resources, including financial resources, to conduct an effective enterprise continuity of operations program.	No		- No definable state space can be defined (Rule 1) - No rule set can be defined (Rule 1) - The expert task cannot be split into concrete mundane or formal sub-tasks (Rule 2, 3, 4, 5, 6) - Humans can do the task better and/or cheaper and/or quicker (Rule 7)	
T0003	Advise senior management (e.g., Chief Information Officer [CIO]) on risk levels and security posture.	Yes	- Problem can be defined (Rule 1) - All sub-tasks in the expert task are mundane or formal (Rule 6) - Artificial intelligence can perform the task better than a human (Rule 7) - Task instances follow a highly consistent structure (Rule 8) - The number of variables and the solution space of the task is limited (Rule 10)		

Figure 6. Task analysis template upon completion (partial view)

22.1% percent of the tasks seems to be outsourceable to artificial intelligence with the most dominant reasons being that the problem could be defined in terms an artificial intelligence system would understand (including an initial state, goal state, definable state space and rule set), the expert task (which most of them were) could be divided up into smaller mundane and formal tasks, and the reason that artificial intelligence with appropriate training would be able to execute the task faster than a human. Some examples of tasks that were deemed outsourceable are:

- "Correlate incident data to identify specific vulnerabilities and make recommendations that enable expeditious remediation;
- Provide technical summary of findings in accordance with established reporting procedures;
- Utilize models and simulations to analyze or predict system performance under different operating conditions;
- Optimize mix of collection assets and resources to increase effectiveness and efficiency against essential information associated with priority intelligence requirements;
- Evaluate and interpret metadata to look for patterns, anomalies, or events, thereby optimizing targeting, analysis and processing;
- Categorize the system and document the security categorization results as part of system requirements." [28]

37.1% of the tasks were deemed to be partially outsourceable. The most dominant reason for this categorization is that there is external input required. For instance, the current problem is quite 'soft' and would require a re-shaping of the task for artificial intelligence to be able to solve the problem, the artificial intelligence system depends on humans to relay non-digital and formal inputs to them or the environment variables need to be expressly conveyed to the artificial intelligence system to form an accurate internal model of knowledge and rule set. A human-on-the-loop or human-in-the-loop collaborative structure could be a feasible solution for these tasks [75, p. 4]. Another interesting reason for tasks being partially outsourceable applied in some cases where tasks contained more than one verb; in short, they were composed tasks of which one part could be outsourced to artificial intelligence and another part could not. Again, the solution for these tasks could be to construct a human-AI collaborative interface in case the sub-tasks are closely linked, or decouple them into multiple mundane/formal tasks when the sub-tasks are not that closely linked. In the next chapter, it will be presumed a 'partial' task can be outsourced to artificial intelligence for 50%; this would at least be an accurate assessment for those expert tasks that contain two verbs and could be split in half between artificial intelligence and humans.

40.7% of the tasks is deemed non-outsourceable. Typically these are tasks that would require artificial intelligence to 'ensure', 'review', 'approve', implement', 'apply', 'acquire', 'develop', 'build', 'manage', 'perform', 'define', 'design' and 'use'. Some of them require the artificial intelligence system to retrieve non-formal inputs from users and stakeholders, which would be quite inefficient (e.g. "Consult with customers about software system design and maintenance"[28]). Others require the artificial intelligence system to make decisions that companies are currently not likely to be willing to outsource to artificial intelligence (e.g. "Manage and approve Accreditation Packages (e.g., ISO/IEC 15026-2)"[28]).

Table 11. Frequency of reasons over 'Yes' category

Reason - Yes	Frequency (out of 221)	% Yes reasons
Problem can be defined	**211**	**95%**
All sub-tasks in the expert task are mundane or formal	**176**	**80%**
Artificial intelligence can perform the task to an acceptable level quicker than a human	**175**	**79%**
Artificial intelligence can perform the task to an acceptable level cheaper than a human	140	63%
Task instances follow a highly consistent structure	130	59%
Artificial intelligence can perform the task better than a human	76	34%
Artificial intelligence can correct for the influence of the environment and other agents	60	27%
The number of variables and the solution space of the task is limited	58	26%
Mundane human task	27	12%
Uses CI competencies / Formal task	26	12%

Table 12. Frequency of reasons over 'Partial' category

Reason - Partial	Frequency (out of 370)	% Partial reasons
Artificial intelligence needs input from another agent or the environment to perform the task	**300**	**81%**
Problem can be defined for only part of the task	222	60%
Common-sense reasoning with a poorly definable knowledge base or rule set is required for part of the task	177	48%
Task instances have some consistent structural elements	165	45%
One of multiple verbs is not associated with mundane or formal tasks	135	36%

Table 13. Frequency of reasons over 'No' category

Reason - No	Frequency (out of 407)	% No reasons
The expert task cannot be split into concrete mundane or formal sub-tasks	**362**	**89%**
Humans can do the task better and/or cheaper and/or quicker	**304**	**75%**
The task has a very high number of variables or a large solution space, such that searching an appropriate solution to the task is unfeasible	256	63%
The environment and other agents have an impact on the task or artificial intelligence that cannot be corrected for by the artificial intelligence	180	44%
Task instances do not follow a consistent structure	125	31%
No definable state space can be defined	124	30%
No rule set can be defined	102	25%
No goal state can be defined	55	14%
No initial state can be defined	33	8%

The reasons highlighted in bold featured most prominently in their category. As mentioned earlier in this chapter, there were two kinds of reasons – whether artificial intelligence could or could not do a task, and whether one would want it to do that task. Whether the problem could be defined was in almost all cases a key factor in deciding whether the artificial intelligence could do such a task. Whether or not the expert tasks (954 out of 999 tasks) could or could not be split up into mundane and formal tasks compatible with known artificial intelligence capabilities was a frequent indicator of a categorisation. Lastly, the 'business case' around artificial intelligence (whether it is better/quicker/cheaper) also played an important role in both the 'Yes' and 'No' categories. The reasons 'Mundane human task' and 'Uses CI competencies/ Formal task' occurred few times because the number of non-expert tasks was limited. In addition, initial states and goal states being non-definable did not occur that frequently, even if the process of getting from an initial state to a goal state was clearly not feasible. For initial states, a '0' or start at a random position would have worked. As for goal states, tasks were often formulated so broadly that a suitable formal goal could have been formulated for the artificial intelligence system.

The categorised tasks are mapped back to the NIST SP800-181 work roles. In Table 14 on page 37 the results are depicted, with cells highlighted in case it surpasses 50% (i.e. a majority of the task) and lightly highlighted in case it surpasses 33% (i.e. more than if equally divided amongst the three categories) in each column. 'Enterprise Architect' seems to be the most non-outsourceable task (79% No), whereas 'Cyber Defense Forensics Analyst' is the most outsourceable (72% Yes).

19 of the 52 work roles are mostly non-outsourceable (>50% No), whereas only 4 roles are mostly outsourceable (>50% Yes). The roles of which the majority could be outsourced to artificial intelligence are 'System Testing and Evaluation Specialist', 'Technical Support Specialist', 'Cyber Defense Analyst' and 'Cyber Defense Forensics Analyst'. These roles are all quite technical and require structured work and data analysis. 11 roles could be partially outsourced to artificial intelligence (>50% Partial). Especially for the 'No' category, results seem to be clustered in certain domains. The average percentage of tasks in each category per domain based on the underlying roles is depicted in Table 15 on page 38. 'Securily Provision' and 'Oversee and Govern' are both more than 50% non-outsourceable. Looking at the roles in these domains, they contain a lot of work and information that would be difficult to convey to artificial intelligence in a formalized way, such as management, leadership and design tasks (e.g. creating training programs and enterprise architectures). 'Protect and Defend' is a category that is 50% outsourceable – it contains tasks that are data-heavy and require monitoring and formal decision-making. The category 'Analyze' is mostly partially outsourceable, indicating there is some potential for artificial intelligence and humans to collaborate in this area.

Table 14. Division of tasks in categories per work role

Specialty	Role	# tasks	% Yes	% Partial	% No
Risk Management (RSK)	Authorizing Official/ Designating Representative	4	0%	25%	75%
	Security Control Assessor	21	0%	29%	71%
Software Development (DEV)	Software Developer	34	15%	41%	44%
	Secure Software Assessor	25	16%	40%	44%
Systems Architecture (ARC)	Enterprise Architect	19	5%	16%	79%
	Security Architect	22	5%	27%	68%
Technology R&D (TRD)	Research & Development Specialist	12	8%	17%	75%
Systems Requirements Planning (SRP)	Systems Requirements Planner	18	0%	22%	78%
Test and Evaluation (TST)	System Testing and Evaluation Specialist	13	62%	23%	15%
Systems Development (SYS)	Information Systems Security Developer	41	22%	34%	44%
	Systems Developer	36	17%	31%	53%
Data Administration (DTA)	Database Administrator	14	29%	50%	21%
	Data Analyst	23	30%	39%	30%
Knowledge Management (KMG)	Knowledge Manager	9	11%	33%	56%
Customer Service and Technical Support (STS)	Technical Support Specialist	12	50%	42%	8%
Network Services (NET)	Network Operations Specialist	11	36%	27%	36%
Systems Administration (ADM)	System Administrator	18	17%	67%	17%
Systems Analysis (ANA)	Systems Security Analyst	31	16%	35%	48%
Legal Advice and Advocacy (LGA)	Cyber Legal Advisor	13	38%	31%	31%
	Privacy Officer/Privacy Compliance Manager	73	14%	23%	63%
Training, Education, and Awareness (TEA)	Cyber Instructional Curriculum Developer	17	6%	29%	65%
	Cyber Instructor	30	7%	30%	63%
Cybersecurity Management (MGT)	Information Systems Security Manager	53	17%	30%	53%
	Communications Security (COMSEC) Manager	9	33%	22%	44%
Strategic Planning and Policy (SPP)	Cyber Workforce Developer and Manager	45	9%	27%	64%
	Cyber Policy and Strategy Planner	19	0%	42%	58%
Executive Cyber Leadership (EXL)	Executive Cyber Leadership	29	10%	28%	62%
Program/Project Management (PMA) and Acquisition	Program Manager	21	5%	33%	62%
	IT Project Manager	25	12%	32%	56%
	Product Support Manager	24	17%	29%	54%
	IT Investment/Portfolio Manager	8	13%	25%	63%
	IT Program Auditor	8	13%	50%	38%
Cybersecurity Defense Analysis (CDA)	Cyber Defense Analyst	34	68%	18%	15%
Cybersecurity Defense Infrastructure Support (INF)	Cyber Defense Infrastructure Support Specialist	9	0%	56%	44%
Incident Response (CIR)	Cyber Defense Incident Responder	17	47%	29%	24%
Vulnerability Assessment and Management (VAM)	Vulnerability Assessment Analyst	8	38%	38%	25%
Threat Analysis (TWA)	Threat/Warning Analyst	29	38%	48%	14%
Exploitation Analysis (EXP)	Exploitation Analyst	20	20%	55%	25%
All-Source Analysis (ASA)	All-Source Analyst	40	38%	53%	10%
	Mission Assessment Specialist	35	31%	51%	17%
Targets (TGT)	Target Developer	35	31%	54%	14%
	Target Network Analyst	24	46%	50%	4%
Language Analysis (LNG)	Multi-Disciplined Language Analyst	29	34%	45%	21%
Collection Operations (CLO)	All Source-Collection Manager	46	35%	30%	35%
	All Source-Collection Requirements Manager	33	33%	36%	30%
Cyber Operational Planning (OPL)	Cyber Intel Planner	45	18%	38%	44%
	Cyber Ops Planner	43	12%	40%	49%
	Partner Integration Planner	34	18%	38%	44%
Cyber Operations (OPS)	Cyber Operator	26	35%	62%	4%
Cyber Investigation (INV)	Cyber Crime Investigator	24	17%	58%	25%
Digital Forensics (FOR)	Law Enforcement/ Counter Intelligence Forensics Analyst	10	30%	30%	40%
	Cyber Defense Forensics Analyst	39	72%	21%	8%

Table 15. Division of tasks in categories per security domain

Domain	Number of tasks	% Yes	% Partial	% No
Securily Provision	245	14%	30%	56%
Operate and Maintain	118	25%	42%	32%
Oversee and Govern	374	13%	29%	58%
Protect and Defend	68	50%	28%	22%
Analyze	212	34%	51%	15%
Collect and Operate	227	24%	39%	37%
Investigate	73	48%	34%	18%

4.4 Summary

This chapter described previous academic practical research done on the applications of artificial intelligence in the cybersecurity industry to get a sense of which type of tasks are already outsourced or could be outsourced to artificial intelligence in this domain. The rules from the previous chapter were rewritten as reasons why a task should fall into a certain category. The statistics of the resulting task division were collected and analysed. The majority of tasks still require a human to perform them (40.7%) or a human component (37.1%). On the other hand, 22.1% of the tasks could be performed by artificial intelligence, and artificial intelligence involvement in cybersecurity tasks could go up to 59.2% if 'Partial' tasks were to be restructured to make them wholly feasible for artificial intelligence.

5. Potential Impact on the Cybersecurity Skills Gap

This chapter takes the division of cybersecurity tasks over the categories outsourceable, partially outsourceable and non-outsourceable to artificial intelligence, and translates it into the macro-economic impact on the cybersecurity skills gap. The background section presents previous statistics published by the collectors of the datasets being used (Burning Glass and GISWS) as to put the macro-economic analysis into historical context. Two datasets from Burning Glass from the website CyberSeek.org are used to determine the impact of AI outsourcing on demand for cybersecurity professionals in each domain in the United States of America. The results are extrapolated to the effect in other parts of the world to evaluate the total impact on the cybersecurity skills gap using the GISWS survey results.

5.1 Background

Burning Glass Technologies is a company that delivers labour market statistics on a commercial basis to employers, educators and the government to support policy, strategy and curriculum decision making. A Georgetown University study from 2014 claimed 60% - 70% of the job openings were then posted online. For jobs which required a Bachelor's degree or higher, this estimate is up to 80 – 90% [76]. Jobs that miss from their dataset mostly are low-skilled jobs and in small businesses, which are advertised informally [77]. Cybersecurity jobs are however relatively high-skilled (in the USA in 2014, only 16% of the cybersecurity jobs did not require a Bachelor's or Master's degree) [78, p. 7].

Burning Glass collects data on a daily basis by searching over 40.000 online job posting boards [79]. Individual job postings are collected and de-duplicated. Postings are then parsed and analysed for 70 different data and meta-data elements such as job title, occupation, employer, industry, credentials, required skills and salary. Cybersecurity jobs were flagged as such if they A) had a cybersecurity-related job title (e.g. network security engineer), B) required cybersecurity certification (e.g. CISSP), or C) required cybersecurity skills (e.g. malware analysis) [79].

They started analysing the cybersecurity labour market in 2007. The first indication of a cybersecurity skills gap can be found in their 2014 report. There were 209,749 cybersecurity online job posts in 2013, which was a growth of 74% since 2007. Cybersecurity job postings accounted for a significant 10% of total IT job openings in 2013, and that these posting stayed 24% longer

unfulfilled than general IT jobs despite an on average 20% salary premium on cybersecurity jobs versus IT jobs [80]. In 2014 they found 238.158 job cybersecurity online job postings with a growth of 91% since 2010 [78].

For global extrapolation, the 2015 GISWS and 2017 GISWS survey reports were used. The Global Information Security Workforce Study (GISWS) is conducted on a biannual basis by the Center for Cyber Safety and Education (Center) and $(ISC)^2$. 19,641 cybersecurity professionals from 170 countries were surveyed in it last edition between June and September 2016. The survey itself is conducted by Frost & Sullivan. There is always a general report published, and some special subject matter reports (e.g. women in cybersecurity, millennials in cybersecurity) are created with each edition. Unlike previous years, this year features five geographically-oriented specialisations [81].

Due to its long-running history and a transparent methodology, the GISWS is considered a reliable source for this type of data. It should be mentioned that there are other reports that present different cybersecurity skills gap sizes such as 3.5 million in 2021, and criticise the GISWS for having an approach to narrowly focussed on 'information security' rather than 'cybersecurity'. They include capabilities such as cyber warfare, Internet of Things and Industrial Internet of Things security, embedded security, automotive, maritime and aviation security, mobile security and medical device security [82]. Although these are undeniably important elements of modern-day security, there is no indication of the GISWS explicitly not including these elements and the methodology for reaching a 3.5 million worker shortage is not explained in more detail [83]. Other sources such as the Digital Skills Committee of the UK House of Lords stay much closer to the GISWS estimates (2 million people short in 2017) [84].

5.2 Approach

Two .csv datasets were obtained from Burning Glass, the data science company behind the production of the website CyberSeek.org in cooperation with NICE and CompTIA. The first dataset contains data from October 2015 – September 2016; the second dataset contains data from October 2016 – October 2017. This is the entire period Burning Glass collected data for Cyberseek.org [85].

Both datasets presents national statistics, statistics per state and statistics per metro areas of the United States of America and contain 15583 data points. The prime variables of interest to this research are the total number of job postings in an area, the total number of employed cybersecurity workers, and the supply-demand ratio. In addition, for each of the NIST SP800-181 domains, the number of online job postings in each geographical area is listed. Some of the variables contain administrative identifiers and population details of the geographical location the data has been recorded from. The top 10 demanded roles (using NIST SP800-181) are ranked for each area. The remainder of the variables presents the number of people that hold a certain security qualification versus the number of people with that security certification that are demanded. No missing values were encountered and the datasets contained no errors or irregularities. This research makes use of the national and state data presented in the data sets.

What demand, supply and the cybersecurity skills gap are in the Burning Glass data is not that straightforward. A discussion with Burning Glass revealed the following considerations [86]: the use of job postings as a comparable variable for labour demand is problematic as job postings are largely caused by churn from current employees [87]. The number of job postings is therefore not considered reliable enough to base the cybersecurity labour demand on. The data from Burning Glass that will be used is A) the number of current cybersecurity employees as 'labour supply', corrected

for the 2% unemployed cybersecurity workforce from friction unemployment (global average 2017 [88, p. 6]) and B) the ratio of job postings divided over each of the NIST SP800-181 domains. Burning Glass has allocated some job postings to multiple domains in case they felt jobs belonged to more than one category. Based on the nature of this research, it is presumed that these job postings in fact consist of sub-jobs in one FTE that each belong to a domain with their own task sets.

For extrapolation to a global impact on the cybersecurity skills gap, the 2015 GISWS and 2017 GISWS studies were used. The figures from 2016 to 2019 were taken from the 2015 report [46] as not all geographies have been published yet for the 2017 survey. The global cybersecurity skill gap of 1.8 million workers and the known overall growth rate of 20% from 2019 were used to calculate the local figures for 2022 [5], [88]. The numbers used as a basis for the extrapolation are shown in Table 16.

Table 16. Cybersecurity workforce, demand, supply and skills gap

Demand	2016	2017	2018	2019	2022
Global	4,416,000	4,908,000	5,424,000	5,963,000	7,155,000
EMEA	1,230,000	1,363,000	1,502,000	1,646,000	1,975,000
Americas	1,867,000	2,081,000	2,308,000	2,546,000	3,055,000
APAC	1,320,000	1,463,000	1,614,000	1,771,000	2,125,000
Supply	**2016**	**2017**	**2018**	**2019**	**2022**
Global	3,796,000	4,007,000	4,227,000	4,456,000	5,347,000
EMEA	1,072,000	1,134,000	1,200,000	1,267,000	1,520,000
Americas	1,596,000	1,692,000	1,792,000	1,897,000	2,276,000
APAC	1,127,000	1,180,000	1,235,000	1,292,000	1,550,000
Gap	**2016**	**2017**	**2018**	**2019**	**2022**
Global	621,000	901,000	1,172,000	1,536,000	1,808,000
EMEA	158,000	229,000	302,000	379,000	455,000
Americas	271,000	389,000	516,000	649,000	779,000
APAC	193,000	283,000	379,000	479,000	575,000

5.3 Results

An assessment is made of how much the impact is on the cybersecurity skills gap in the United States of America using Burning Glass data. Burning Glass recommends using the ratio between the current number of cybersecurity employees in their dataset and the number of unique online job postings as an indicator for the severity of the cybersecurity skills gap [86]. The 2017 ratio of 2.6 employees to 1 job posting versus the national USA average of 5.6 employees to 1 job posting entails that in order to align with regular market conditions, the cybersecurity workforce would need to more than double overnight in order to match demand [86]. Although job posting data may not be suitable for labour market predictions, it does allow one to spot the relative differences in geographies. An overview of the actual state ratios for 2016 and 2017 in the United States is shown in Figure 7, accounting for 2% unemployment. A ratio of less than 1 entails that the number of job postings for that year is larger than the total number of working security professionals in that area. South Dakota in 2016 was that kind of a stretched labour market. States with a ratio of less than 2.23 in 2016 and 2.61 in 2017 are below the national average for cybersecurity jobs. Only Wyoming in 2016 had a ratio of more than the national job average of 5.6. The national cybersecurity

demand/supply ratio slightly improved from 2016 to 2017 (also visible from the greener colours in the maps) but still lags far behind the ratio expected from an average labour market.

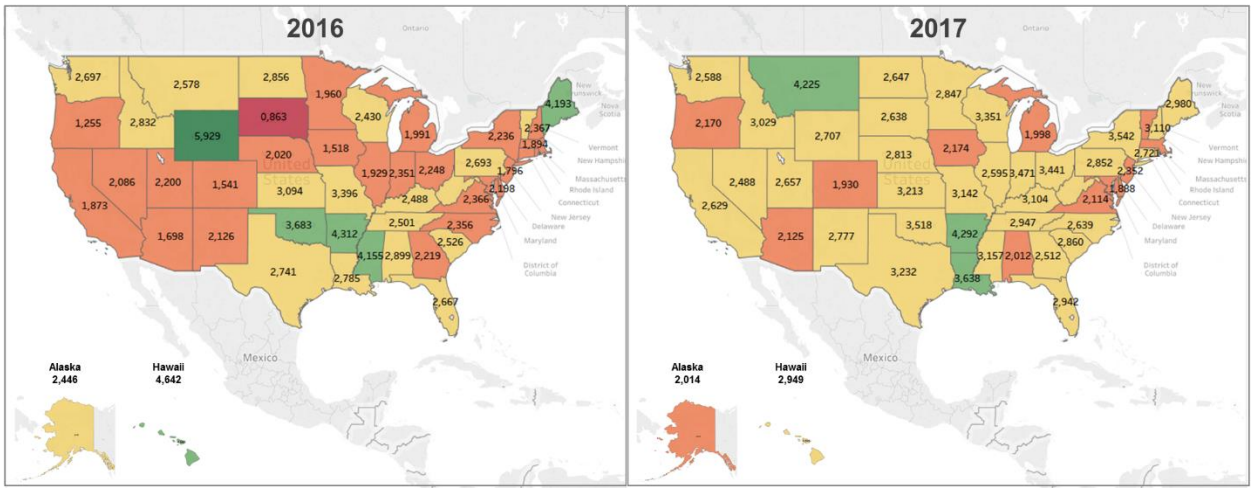

Figure 7. Supply/demand ratio based on online job postings ([89], [90])

It is presumed 100% of the work from tasks in the 'Yes' category and 50% of the work from the tasks in the 'Partial' category could be outsourced to artificial intelligence. It is presumed that the weight of tasks is equal over an FTE job and therefore a proportion of work in terms of tasks is equal to a proportion of FTE jobs. The proportional division of work over the NIST SP800-181 domains is taken into account, as well as a 2% unemployment rate in addition to the current supply of cybersecurity workers. Figure 8 has been generated assuming 100% adoption of artificial intelligence for the work that is outsourceable.

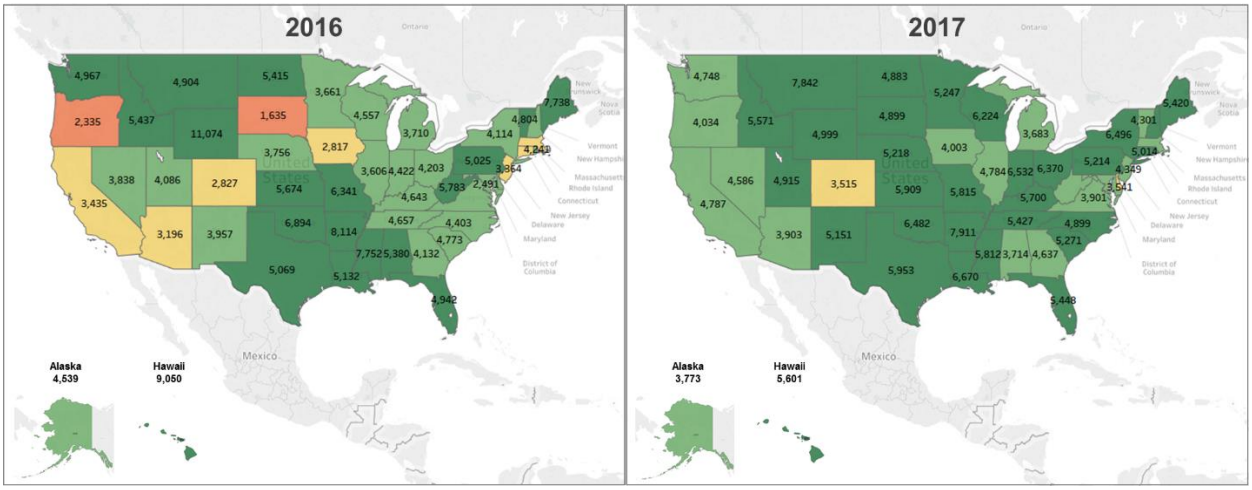

Figure 8. Supply/demand ratio based on online job postings with 100% AI adoption ([89], [90])

The maps in Figure 8 are quite greener than those in Figure 7 and indicate a less stretched cybersecurity labour market overall. Only South Dakota and Oregon in 2016 would still have had lower supply/demand ratios than the national cybersecurity average. A significant number of states in both 2016 and 2017 has now a supply/demand ratio higher than the current national average ratio for all jobs (5.6).

The next step is to extrapolate these results to the potential impact artificial intelligence could have on the global cybersecurity skills gap. Using the same conditions as in the previous calculations, the outsourceable percentage of work per NIST SP800-181 domain is calculated and percentage of remaining work is calculated (i.e. the cybersecurity skill gap), depicted in the top half of Table 17. The first row is the proportion of each of the domain to the total pool of cybersecurity work. Using the percentages of tasks that could be outsourced fully and partially (e.g. for the domain 'Securily Provision', 14%*100% + 30%*50% = 29%). These percentages are then put into context of various artificial intelligence adoption scenarios, as it is not realistic to expect artificial intelligence to be fully adopted for all outsourceable tasks at a given point in time. The lower part of Table 17 shows the remaining work percentage given various adoption rates for each of the NIST SP800-181 domains.

Table 17. Outsourceable tasks per NIST SP-800-181 domain and impact per AI adoption scenario

	Securily provision	Operate & Maintain	Oversee & Govern	Protect & Defend	Analyze	Collect & Operate	Investigate
Size (% Total Tasks)	22%	28%	12%	14%	15%	6%	3%
Impact Yes (100%)	14.0%	25.0%	13.0%	50.0%	34.0%	24.0%	48.0%
Impact Partial (50%)	15.0%	21.0%	14.5%	14.0%	25.5%	19.5%	17.0%
Impact No (0%)	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
Outsourceable tasks	29.0%	46.0%	27.5%	64.0%	59.5%	43.5%	65.0%
1 – (Outsourceable tasks*AI adoption rate)=							
100% adoption rem.	71.0%	54.0%	72.5%	36.0%	40.5%	56.5%	35.0%
75% adoption rem.	78.3%	65.5%	79.4%	52.0%	55.4%	67.4%	51.3%
50% adoption rem.	85.5%	77.0%	86.3%	68.0%	70.3%	78.3%	67.5%
25% adoption rem.	92.8%	88.5%	93.1%	84.0%	85.1%	89.1%	83.8%

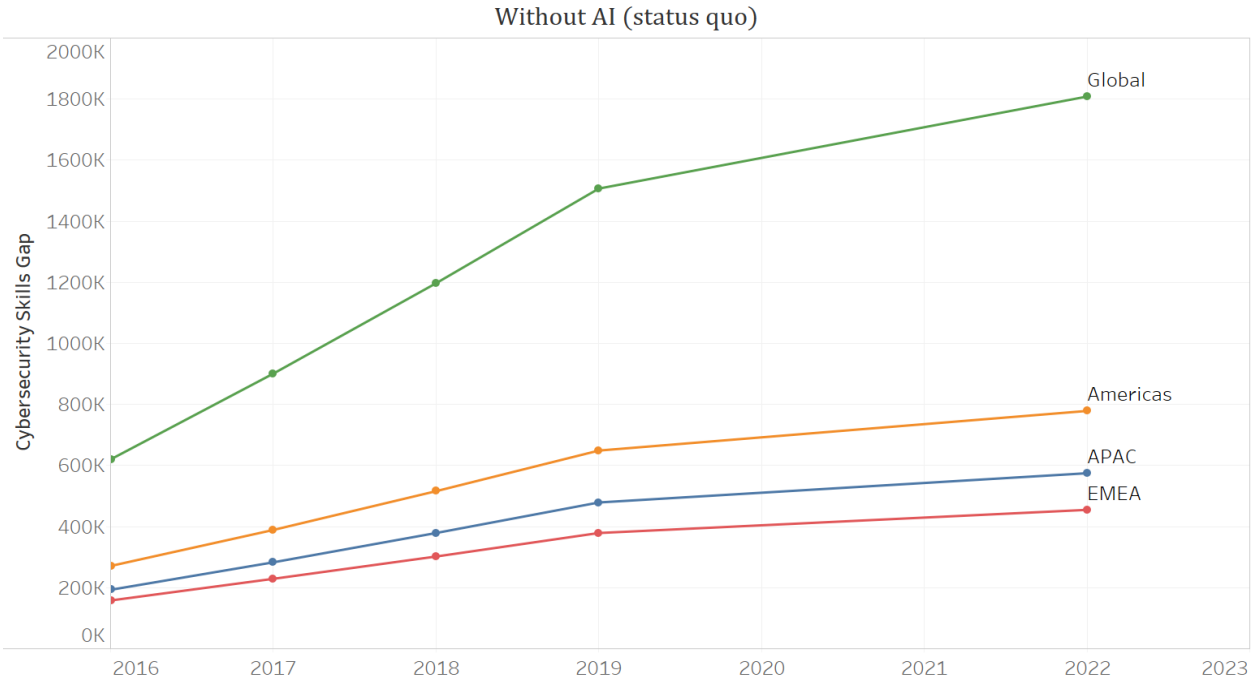

Figure 9. Status quo of the cybersecurity skills gap (no AI adoption for outsourceable tasks)

Figure 9 depicts the current cybersecurity skills gap presuming no adoption of artificial intelligence for the tasks this research deems outsourceable beyond what has already been outsourced to artificial intelligence and is accounted for in the labour market statistics of the GISWS. The Americas, APAC and EMEA combined equal the estimated global cybersecurity skills gap of 1.8 million cybersecurity workers in 2022. Earlier predictions of the cybersecurity skill gap up to 2019 follow a linear growth rate, but this growth rate drops off towards predicting the 2022 numbers. Figure 10 on the next page depicts the same graph for various rates of AI adoption (25%, 50%, 75% and 100%). AI adoption rates are not instantaneously at a certain level and then remain there. It is more likely that for a nearby point in time a low adoption rate applies (e.g. 25%) which grows over time (e.g. to 75%). Predictions on how fast artificial intelligence will be adopted vary and industries in which cybersecurity is practiced (e.g. finance, manufacturing) are likely to have different adoption rates for artificial intelligence [91], making it unfeasible to provide a single best digit as an estimation. These four adoption scenario graphs are therefore included so the reader can estimate the impact on the cybersecurity skills gap for different adoption rates at different points in time.

Table 18 summarizes the total impact of the different artificial intelligence adoption scenarios. These percentages were calculated by summing up ('Size of cybersecurity skills gap' * 'Proportion of the domain versus domain total') * ('Remainder of tasks under each AI adoption scenario') for all domains. The difference between the outcome for a specific AI adoption scenario and the status quo (0% adoption) was expressed in percentages. Percentages were equal across years; only the absolute number of jobs varied as the size of the cybersecurity skills gap varied. If 100% AI adoption were to be achieved for the tasks deemed outsourceable in this research, the global cybersecurity skills gap could be reduced with 45% (from 1.8 million to 1 million in 2022) or over 800,000 jobs. As artificial intelligence capabilities advance over time, more tasks might become executable for artificial intelligence and the cybersecurity skills gap can be decreased even further.

Table 18. Global cybersecurity skill gap reduction

AI adoption rate	0%	25%	50%	75%	100%
Global cybersecurity skill gap reduction	0%	11%	23%	34%	45%

Some of these tasks may already be fully or partially performed by artificial intelligence, as indicated in the list of security applications in section 4.1. This could mean that the industry is e.g. already at 5% adoption, and that a 100% AI adoption scenario would lead to a reduction of 760,000 jobs in 2022 rather than 800,000. In addition, adopting artificial intelligence will create a plethora of AI related jobs and although these may or may not be seen as 'cybersecurity jobs' they will limit the reduction of workers needed. This 'gap-expanding' considerations are addressed in section 6.3.

5.4 Summary

This chapter has shown how much impact fully outsourcing all feasible tasks to artificial intelligence could have on the employee shortage in the cybersecurity industry. A detailed use case on the United States of America was presented using online job posting data. This helped define the distribution of the labour market over the NIST SP800-181 domains. This division was then used in conjunction with the GISWS data to calculate the impact on outsourcing tasks to artificial intelligence under various scenarios of AI adoption over time. If 100% AI adoption were to occur by 2022 for all of the tasks named fully or partially outsourceable in this research, 45% of the global cybersecurity skills gap or over 800,000 jobs in 2022 could be outsourced to artificial intelligence.

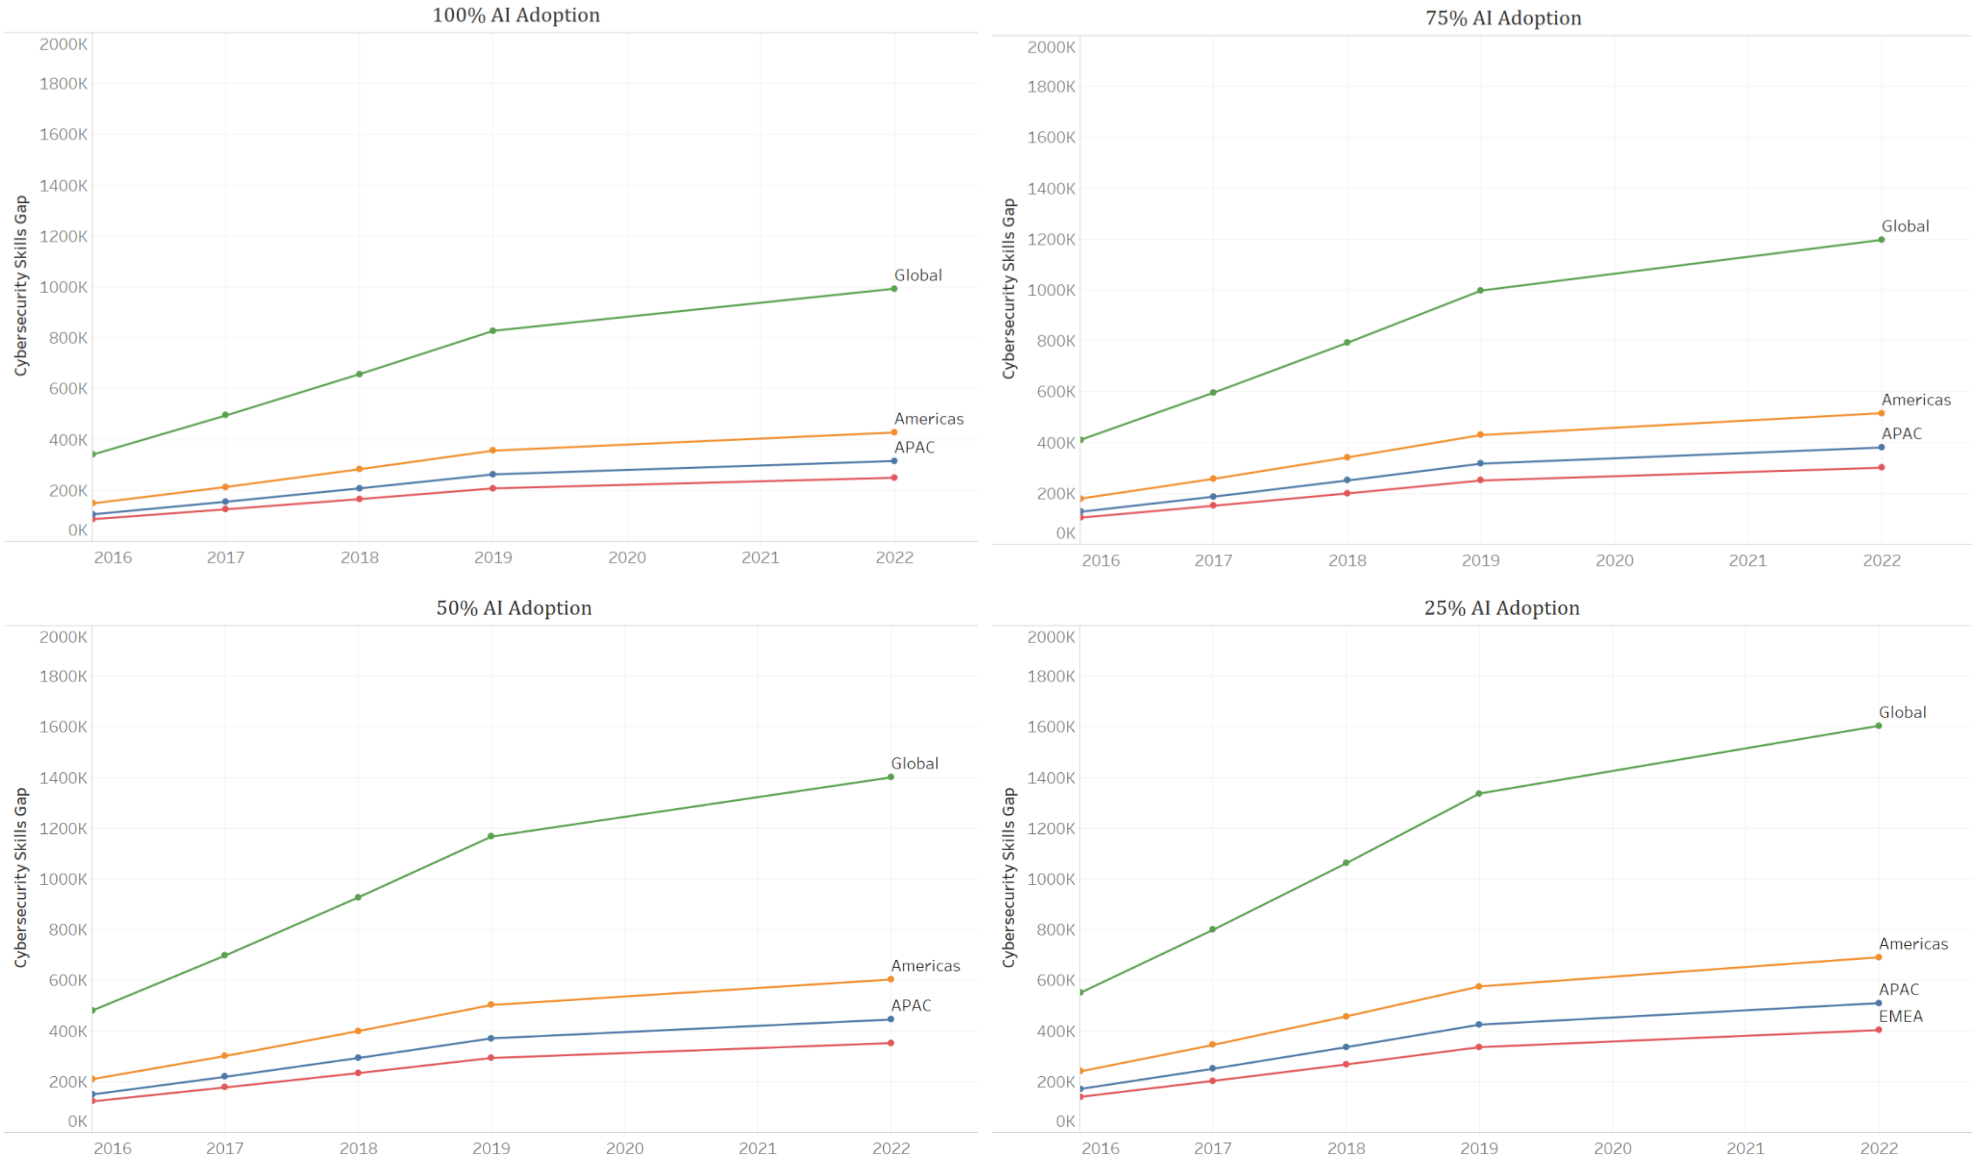

Figure 10. Various AI adoption scenarios and their impact on the cybersecurity skills gap

6. Conclusion & Discussion

This chapter looks at each of the elements core to answering the research question, i.e. cybersecurity tasks, artificial intelligence capabilities, the possibility to outsource tasks and the impact on the cybersecurity skills gap. The work done is shortly summarized and caveats and considerations are addressed. A discussion is presented on the counteracting effects of introducing artificial intelligence for these tasks, as they might reduce the reduction of the cybersecurity skills gap (or even enlarge it). Lastly, the results are validated and put into context of previous 'new technologies' and to what extent they caused technological unemployment. Potential future research is indicated

6.1 Conclusion

This thesis had as its goal to turn a potential risk into an opportunity and answer the question whether artificial intelligence can take over work in the cybersecurity industry. The cybersecurity industry has been struggling to find sufficient employees – both in quantity and in quality. Whereas for other industries artificial intelligence could entail people losing their jobs, in this industry artificial intelligence could help perform work that now cannot be done due to the cybersecurity skills gap. This research has provided an answer to that question by taking a structured approach to determining which tasks could be performed by artificial intelligence and deducing the macro-economic effects of such an outsourcing operation.

The NIST SP800-181 standard was used to index and understand the labour market of the cybersecurity industry. The standard consists of domains, specialties, work roles, tasks, knowledge, skills and abilities. Tasks were chosen as an appropriate level of analysis for determining whether work could be outsourced to artificial intelligence. There was a large variation in the number of tasks per role (from 4 up to 73), the number of tasks per domain (from 68 up to 374) and, since tasks could be n-to-n for work roles across all security domains, the number of connections between work roles and tasks (from 0 to 6).

This thesis looked at artificial intelligence as a subject of interest, and includes both Good Old Fashioned Artificial Intelligence and Computational Intelligence in its analysis. The three DARPA waves of artificial competency were used to put the artificial intelligence techniques and applications currently available into context. Based on a literature review an overview of human and mundane competencies, and computational intelligence competencies have been provided. Examples from literature on artificial intelligence applications and what tasks and work roles artificial intelligence

systems performed, were provided and used to build a rule set to help determine which tasks can and cannot be outsourced.

Based on literature research an inventory of how artificial intelligence is currently used in the cybersecurity industry was put forward to show what applications are already considered feasible and in existence. The rule set from the previous chapter was rewritten into reasons split over the 'Yes', 'Partial' and 'No' categories and a structured task analysis was performed using these reasons for determining to what extent tasks could potentially be outsourced to artificial intelligence. 22.1% of the tasks was deemed outsourceable, whereas 37.1% of the tasks was partially outsourceable and 40.7% not outsourceable. The most frequent reasons given to put a task into a certain category were related to whether or not a formal problem could be defined, whether and expert task could be split up into mundane and formal sub-tasks and whether there was a business case for letting a task be done by artificial intelligence). 19 out of 52 work roles were mostly not outsourceable. For four roles ('System Testing and Evaluation Specialist', 'Technical Support Specialist', 'Cyber Defense Analyst' and 'Cyber Defense Forensics Analyst') the majority of the work could be outsourced to artificial intelligence. Converging these roles into their respective security domains showed that 'Securily Provision' and 'Oversee and Govern' could for the most part not be outsourced to artificial intelligence. 'Protect and Defend' was the only domain from which a majority of the tasks could be performed by artificial intelligence.

The datasets of Burning Glass Technologies and the Global Information Security Workforce Study (GISWS) were used to determine the potential impact of the outsourceability of tasks on the cybersecurity skills gap. Since online job posting data is not by itself suited for labour market demand analysis, the supply/demand ratio as status quo and with 100% AI adoption were analysed for the United States of America, presuming 100% of the tasks in the 'Yes' category and 50% of the tasks in the 'Partial' category could be outsourced. Combining the proportional division of Burning Glass data on the seven NIST SP800-181 domains and the GISWS FTE data on the cybersecurity skills gap, a scenario analysis on cybersecurity skills gap reduction was performed under various rates of AI adoption. The potential impact on the global cybersecurity skills gap, given the scope being limited to the tasks and artificial intelligence techniques analysed in this research and an AI adoption of 100%, is that 800,000 jobs could be performed by artificial intelligence in 2022.

6.2 Reflection

Since the first edition of the NIST SP800-181 standard was only published in August this year (2017), this is bound to be one of the first academic works to inspect the standard up close. It being the first version, it is not surprising some irregularities were encountered. An example of this were the 90 tasks that were not linked to any role, specialty or domain. They were assessed like any other task on whether or not they could be fully, partially or not outsourced to artificial intelligence. Since they could not be related to a role or a domain, they have not been included in the macro-economic analysis and this might be an improvement for future research. Besides the withdrawn and integrated tasks that were clearly indicated as such, four more tasks were removed for being a duplicate value. Some of the task, knowledge, skill and ability descriptions were occasionally vague and required research and assumptions from the researcher to gain a proper understanding. More accurate descriptions would undoubtedly aid in an even more precise classification of tasks on outsourceability. It would be best if the bias towards American governmental cybersecurity tasks could be adjusted for in a new version of the standard. Regardless, the NIST SP800-181 is a large step

forward in understanding work in the cybersecurity industry and the knowledge, skills and abilities required to perform that work.

Drawing conclusions on the current competencies of artificial intelligence was difficult as there was no consensus in academic literature on the relations between and terminology of certain elements (e.g. artificial intelligence versus computational intelligence, and to which of the two some artificial intelligence techniques belong). This has hindered an initial plan to provide a reader with a full ontological overview of artificial intelligence techniques. Fortunately there was a good spread of academic literature from which the competencies, tasks and work roles that artificial intelligence could perform, could be drawn and on which the rule set was based. Since nothing similar to this rule set in literature has been found, this is perhaps the first attempt ever to do so and other researchers might seek to add to it based on future research. The initial rule set was refined during the initial stages of the task analysis, and has been made more robust before commencing the full task analysis.

The task analysis was performed manually by the researcher. The rule set could have been formalized even further and given to artificial intelligence to classify the tasks accordingly. Initial attempts proved to be difficult due to the different meanings of words in a specific context and an inconsistent use of verbs that could indicate whether or not an artificial intelligence would be able to do a task. The manual analysis however poses the drawback that bias may have been introduced by the researcher. This has been limited as much as possible by having to list the appropriate reasons for selecting a category for the task. The researcher therefore has confidence that the results are reasonably accurate. It is imaginable that one could debate the choice of the 'key' reasons and perspectives on how much effort it would cost to have artificial intelligence perform a task to acceptable standard may vary depending on the skill of the artificial intelligence developer, time and resources available and what is considered 'acceptable'. For instance, self-driving cars are already approximating human-level driving capacities and are expected to become a safer driving alternative. However, regulators have already indicated they expect self-driving cars to be safer than human drivers before allowing them on the road [92]. Acceptability may highly vary per industry, organisation and even the specific use case.

The calculation of the impact on the cybersecurity skills gap for the United States of America and the world has been done to the best extent possible given the data. The Burning Glass data on job postings was divided into the seven NIST SP800-181 domains. The researcher however learned that some tasks has been placed into multiple categories, but that data on how many tasks, the nature of these tasks and in what categories was not available. It was therefore presumed that they were equally divided over all categories and the proportions as-is were used to perform calculations on the GISWS data. Since the 2017 GISWS publication only mentioned the expected shortage of cybersecurity professionals in 2022 and the expected growth rate, the local 'gaps' had to be deduced from those numbers. The 2016 – 2019 figures have been taken from the 2015 GISWS and these predictions may therefore be slightly outdated if the growth rate of the industry has altered. AI adoption was assumed to take place in equal pace over all NIST SP800-181 domains. Tasks were furthermore assumed to take up an equal time expenditure for each of the work, which is unlikely to be wholly true (tasks = work = FTE), but a better assumption was lacking. It is likely that the outsourcing of 800,000 jobs can only occur if the part of the work role that could be outsourced, was cut loose from the FTE work role so that the FTE could assume other non-outsourceable tasks. Despite the assumptions made for the analysis, the extrapolation gives a good indication of the significant impact that artificial intelligence could make in reducing the cybersecurity skills gap.

6.3 Counteracting effects of AI Adoption

As mentioned in the introduction of this thesis, this research looks at how much artificial intelligence could reduce the cybersecurity skills gap. However it is likely that the adoption of artificial intelligence will have other effects on the cybersecurity industry and its labour market. Three of those potential effects have been identified in literature and previous research of the researcher [93], [94] and will be elaborated on in this section. Quantitative data on these effects is limited and hinders any assumption one might make on how jobs they require and will thus reduce the reduction on or even expand the cybersecurity skills gap.

6.3.1 Training, implementation and maintenance

In order to have artificial intelligence do all the tasks as stipulated in this research, people will be needed to shape the environment and task appropriately for artificial intelligence, design, develop, train and maintain the artificial intelligence system over time. These new jobs are not per se cybersecurity jobs and it is not clear whether they should be counted as expanding the cybersecurity skills gap. It is however likely that this category of jobs is to quickly grow in the future and that society will struggle fulfilling the demand for artificial intelligence designers and trainers. This may prove to be a bottleneck for the adoption of artificial intelligence systems in the cybersecurity industry.

Research performed by job posting website Indeed in 2017 revealed that there were already 2.3 vacancies for artificial intelligence experts and data scientists per qualified candidate on their platform [95]. A survey of 1000 companies performed by Capgemini suggested four out of five companies saw new jobs created in relation to the introduction of artificial intelligence in their organisation [96]. Other sources suggest that although it may be true that new jobs will be created because of artificial intelligence, these will be high-skilled jobs asking for a complex KSA set whereas the jobs lost will be mostly low-skilled jobs. This is supported by a report from Forrester Research that suggests that Cognitive Technology (which also includes robotics and automation in addition to artificial intelligence) will see 16% of USA jobs lost in 2025, but 9% new jobs created – jobs lost will be low-skilled, jobs gained will be high-skilled [31]. This would suggest the creation of an Artificial Intelligence Skills Gap and the occurrence of technological unemployment nonetheless [97], [98].

One particular publication of interest suggests there will be three new categories of jobs in relation to artificial intelligence based on early developments they have discerned in the artificial intelligence labour market [30]. The job categories are teachers, explainers and sustainers and respective example job titles are listed in Table 19. Teachers will train artificial intelligence systems, reduce errors in their performance and optimise their search strategy. For some applications (e.g. customer service chat bots), mimicking human behaviour will be part of that performance and requires an 'empathy trainer'. Explainers will bridge the gap between technologists, the business and the customer. They are to provide clarity on the possibilities and the limitations of artificial intelligence, especially if the artificial intelligence system is considered to be opaque. A case of particular interest is the General Data Protection Regulation's (GDPR) 'right to explanation' that effectively allows the consumer to question and fight any decision made by artificial intelligence that affects them [99]. The sustainer category ensures the artificial intelligence system operates as intended and addresses unintended consequences. An important role in the initial stages of AI adoption in this category will likely be the ethics compliance manager that acts like a watchdog to address any biases in the artificial intelligence system's code [30].

Table 19. Representative jobs to be created for AI development & maintenance (adapted from [30])

Teachers	
Customer-language tone and meaning trainer	Teaches AI systems to look beyond the literal meaning of a communication by, for example, detecting sarcasm.
Smart-machine interaction modeller	Models machine behaviour after employee behaviour so that, for example, an AI system can learn from an accountant's actions how to automatically match payments to invoices.
Worldview trainer	Trains AI systems to develop a global perspective so that various cultural perspectives are considered when determining, for example, whether an algorithm is 'fair'.
Empathy trainer	Trains AI systems to assimilate human behaviour and emotions.
Explainers	
Context designer	Designs smart decisions based on business context, process task, and individual, professional, and cultural factors.
Transparency analyst	Classifies the different types of opacity (and corresponding effects on the business) of the AI algorithms used and maintains an inventory of that information.
AI usefulness strategist	Determines whether to deploy AI (versus traditional rules engines and scripts) for specific applications.
AI forensics analyst	Identifies what variables and steps led to a certain outcome.
Sustainers	
Automation ethicist	Evaluates the noneconomic impact of smart machines, both the upside and downside.
Automation economist	Evaluates the cost of poor machine performance.
Machine relations manager	'Promotes' algorithms that perform well to greater scale in the business and 'demotes' algorithms with poor performance.

Not all of the above jobs are necessarily high-skilled; e.g. an empathy trainer makes use of innate human and mundane competencies to teach the artificial intelligence system. This would thus require organisations looking to build an artificial intelligence implementation workforce to train both low-skilled and high-skilled people.

6.3.2 Securing artificial intelligence

As artificial intelligence adoption grows in the cybersecurity industry and increasingly becomes part of the fabric of many organisations and governments, it will by itself become an attractive target for black-hat hackers. Implementing artificial intelligence will therefore require new work roles in cybersecurity that are specialised in protecting the artificial intelligence system.

In cybersecurity three security aspects of a system are of interest: its confidentiality, its integrity and its availability. Breaching the confidentiality of an artificial intelligence system could lead to hacker becoming aware of the initial state and definable state space of the artificial intelligence system's problem solving, but more importantly the goal state and the rule set for achieving that goal state could become known. Knowing about a rule set could enable an attacker to

adapt inputs fed to the artificial intelligence system in such a way that the machine learning patterns could become compromised (its integrity) as it learns from inaccurate data, that the artificial intelligence system would not know how to handle (e.g. a non-solvable version of the problem impacting its availability) or that it will make undesirable decisions based on the faulty inputs. The integrity of the artificial intelligence system could furthermore be compromised by altering the components of the artificial intelligence system themselves (the initial state, goal state, definable state space based on its worldview and the rule set itself). This would entail that despite receiving correct data, the artificial intelligence system would start to make incorrect decisions. Availability of the system could be attacked (similar to other systems) by e.g. compromising its physical environment, isolating it from required inputs and performing a (Distributed) Denial of Service attack ((D)Dos).

To achieve these goals there are three types of attacks that can be performed: evasion attacks, algorithm attacks and poison attacks. Evasion attacks seek to explicitly bypass the artificial intelligence system's controlling function. Algorithm attacks seek to compromise the foundations on which the artificial intelligence system makes its decisions and thereby impact its decision making capabilities. Poison attacks seek to compromise the data fed to the artificial intelligence system and thereby impact its decision making capabilities. Research into how these attacks would work and how they can be protected against is still in an early phase, although the some academic research groups are currently laying the ground work (e.g. AdverseriaLib from PRA Lab [100]).

6.3.3 Using artificial intelligence against security

Although artificial intelligence could do much good in the cybersecurity industry, it could also be of interest to black-hat hackers as a means to attack organisations, governments and consumers. The 2016 DARPA Grand Cyber Challenge demonstrated the current potential of machine learning systems to attack other machine learning systems [94]. Although the systems were built for a specific DARPA environment and general applicability was limited, the artificial intelligence systems demonstrated their ability to identify vulnerabilities, exploit the vulnerability on a target system, create and deploy a patch whilst maintaining system availability and share vulnerability- and patch information with other systems. When pitted against human hackers in a Capture the Flag contest the artificial intelligence systems still lagged behind their human counterparts but this can be expected to change as artificial intelligence systems become more advanced.

With artificial intelligence systems being weaponised, the parties under attack will have to adapt their defences to compensate for their increased attack speed, increased attack volume and large diversity of attacks. The results of the DARPA CGC however suggested that the artificial intelligence systems performed better at defence than attack [94, p. 27], and so training artificial intelligence systems to defend against attacking artificial intelligence systems might be a feasible solution – although it could result in an artificial intelligence loss of arms.

6.4 Comparison to Other New Technologies

This section attempts to validate the findings regarding artificial intelligence's impact on the cybersecurity skills gap by looking at previous 'new technologies' and see what their net impact on employment was. Technological unemployment has been known to vary depending on whether it concerned process or product innovation, with process innovation being likely to create more new jobs than product innovation as they reduce costs and fuel product demand. In the case of artificial intelligence, both product and process innovation are likely at play [33, p. 3]. Feldman surveys academic research done on this topic, some of which find a positive net creation of jobs whereas

others find a negative net creation of jobs. Results have not been clear-cut as there are potential confounding factors such as trade union density, collective bargaining coverage and wage bargaining coordination and measured outputs such as R&D expenditure and number of patents requested are poor proxies for innovation [33, p. 13]. One of the trends that has been noticed is that an initial effect of technological change is a transitory loss of jobs, which dissipates after three years when new jobs start to compensate for the loss of jobs [101]. Another empirical finding is that the faster the technological change occurs, the higher the initial rate of unemployment will be [33, p. 25].

The Industrial Revolution (1760 – 1840) is one of the most-well known movements in history when it comes to replacing human labour with machines. A frequently cited example of a technology that was feared to cause technological unemployment was the weaving machine. Elisabeth I even refused to grant a patent for the weaving machines because of this fear [102]. The machines came nonetheless and were met with hostility by a group of weavers known as the Luddites, who destroyed a large number of machines in protest against the feared job loss. This was later named the Luddite Fallacy as products became cheaper and increased demand to such extent many of the workers kept their jobs and transitioned to work in which they maintained the machines [103]. After an initial period of poverty in the Industrial Revolution, productivity climbed and workers transitioned to the modernized parts of the industries with higher wages [102]. This trend is depicted in Figure 11 and shows how the percentage of workers in the United States in America declined over time in agriculture and initially rose in manufacturing. As our economies experienced another shift to service economies, the percentage of people working in manufacturing declined again, demonstrating the transitive nature of work. The advent of the steam machine, the copy machine, computers and the Internet have created similar 'shifting work' movements.

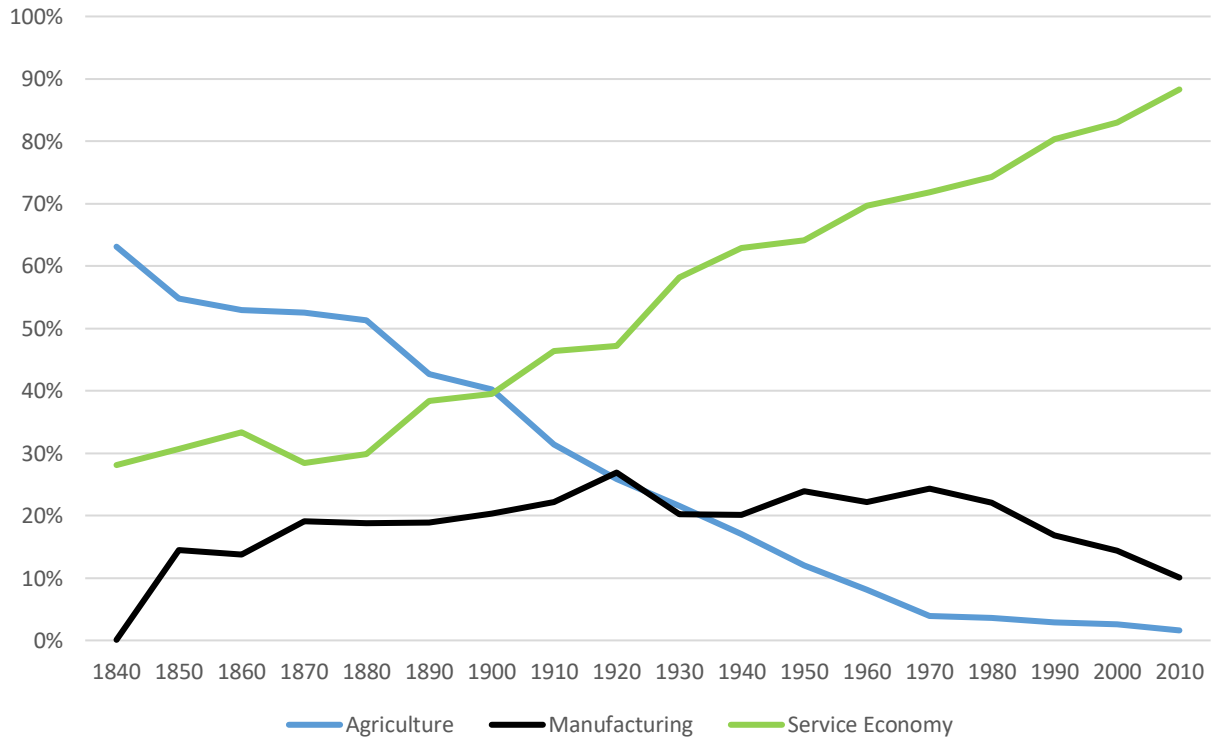

Figure 11. Percentage of USA workers in different industries (data from [32])

Unfortunately there is little research available on the impact of specific technologies creating technological unemployment. Many modern publications summarize all technologies under 'automation'. This entails that it is possible to draw conclusions on the general patterns of technological unemployment that may arise, but not determine how technologies differentiated from each other and what factors might contribute to this differentiation. Two studies by McKinsey indicate the automation potential for 54 countries, 12 and 19 sectors and 800 occupations [104], [105]. 44.67% of the work (12 sectors) and 46.84% of the work (19 sectors) were deemed to automatable. This is very similar to the 45% outsourceability predicted by this research for artificial intelligence. There were no large differences between the proportions of work that could be automated for 54 countries studied, indicating that the assumptions of equal outsourceability over all countries for the global cybersecurity skills gap extrapolation in this research might hold. McKinsey has taken a similar approach to this problem as this research, i.e. using work activities rather than occupations to estimate the feasibility of outsourcing work. The time spent per activity for all US occupations and an overview of how much time in that task could be automated is depicted in Figure 12. There were few tasks involving physical work in NIST SP800-181, so these task categories are not of much relevance. 'Data collection' and 'data processing' are however noted as the two other task categories highly susceptible to automation. For artificial intelligence outsourceability there was also an inclination towards tasks that were data-heavy. Our artificial intelligence research also perceived that 'managing others' and 'stakeholder interactions' was difficult to outsource to non-human actors. A major difference between the McKinsey research and this research is that the 'applying expertise' category is listed as least susceptible by McKinsey. In this research on artificial intelligence planning and decision-making were (in case appropriate data was available) capabilities that artificial intelligence could possibly take over quite well. Creative tasks from the 'applying expertise' category were noted to be more difficult to outsource to artificial intelligence, e.g. the creation of security awareness training programs.

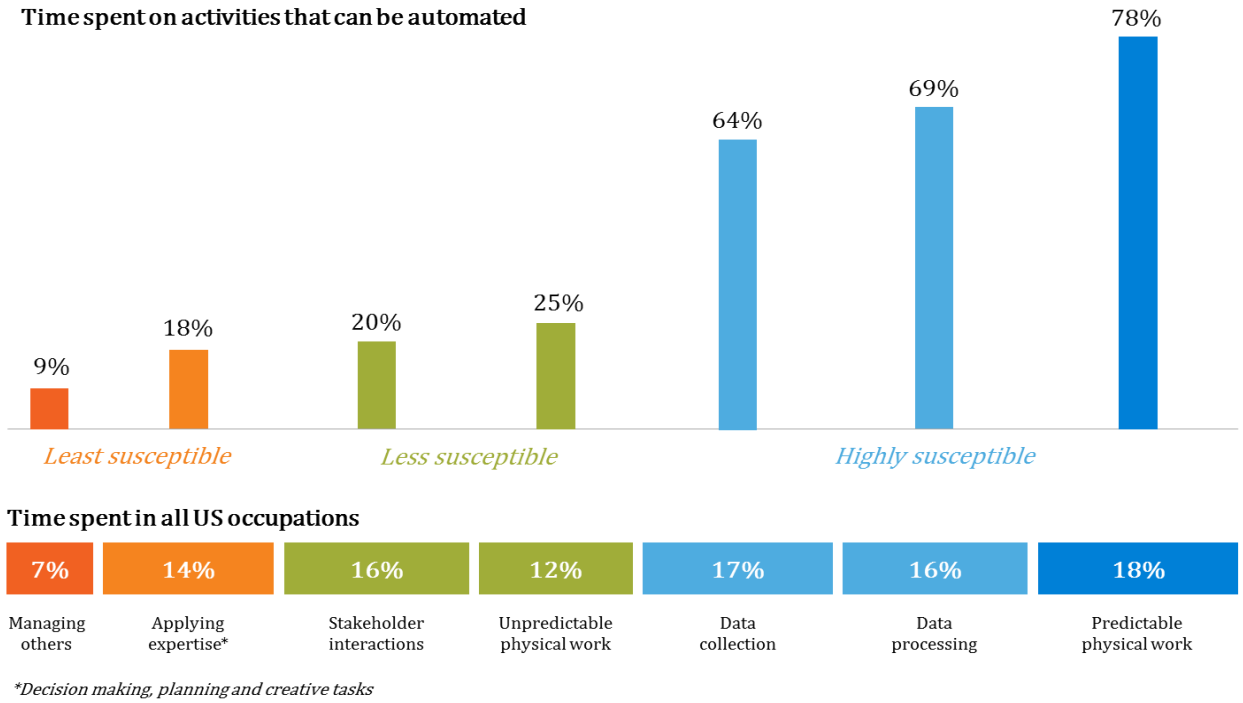

Figure 12. Percentage of time spent in US occupations (adapted from [106])

The overall patterns of the McKinsey automation research seem to be fairly closely aligned with the findings of this artificial intelligence outsourceability research. In addition, we can conclude based on the literature review performed it is likely that initially artificial intelligence will help in reducing the cybersecurity skills gap, that this effect will be more extreme if we decide to adopt artificial intelligence rapidly, but in the long term other jobs will be added to the cybersecurity labour market to take care of artificial intelligence implementation and maintenance, artificial intelligence security and security defence against artificial intelligence attacks.

A game changers versus previous technological innovations is that as artificial intelligence matures enough, it might become capable of writing programming code independently and expanding upon its own task set beyond what is currently possible. If this occurs, it is imaginable that artificial intelligence will take over those jobs that previously came into existence as side-effects ('managing the machines'). Since previous innovations have led to an increase in product demand and this would be an undesirable situation with regards to the cybersecurity skills gap, it might not be necessarily bad if artificial intelligence were to reach this increased level of independence.

6.5 Academic and Societal Relevance

Previous academic research on artificial intelligence has looked at the potential and impact of specific artificial intelligence techniques. Societal discussion on the impact of artificial intelligence on the world of work is proliferating, yet the impact is rarely qualified beyond 'a significant proportion of the work' or a single percentage. This research attempted to offer the best of both worlds by creating a structured approach based on ten rules to determine the outsourceability of tasks to artificial intelligence, and providing a transparent extrapolated calculation from the proportions of outsourceable work to the impact on the cybersecurity skills gap for various years, geographies and AI adoption rates.

The first specific academic contribution is thus a set of ten rules by which tasks can be assessed for their outsourceability to artificial intelligence. Hopefully it is a product that other researchers will build upon. This research applied those rules to cybersecurity tasks, but the rules have been formulated in a generic way so that they are applicable to all kinds of work domains.

The second academic contribution is that the research provides a cybersecurity case study on how to take these outsourceability ratios and translate them to the impact on work roles. Doing the same for other industries could shed some light on how the overall labour market across industries is likely to change in the future given various AI adoption scenarios.

As far as societal impact goes, providing insight into the outsourceability of cybersecurity tasks (and similarly for tasks from other industries) can aid organisations with deciding in which knowledge, skills and abilities to invest given the likelihood of outsourcing for each of the tasks and work roles. It can also help organisations determine in which artificial intelligence techniques and applications to invest for adoption, given a confirmation of their partial or full outsourceability in this cybersecurity research and those for other industries. Lastly, these insights will also allow the general public to steer their careers towards those areas less likely to be outsourced to artificial intelligence. To this end, the researcher is looking towards publishing a Point of View on this thesis with Deloitte in order to convey its messages to a wider audience than the academic community.

6.6 Future Research

As mentioned in section 6.2 there are some things that could still be improved for this research. These elements could be fruitful directions for future research. First of all, the nature of the data available led to some generalisations and assumptions for the extrapolation. As awareness on the cybersecurity skills gap grows and so does the need to track its development, available statistics will hopefully proliferate. In particular, data on the demand and supply of professionals for the NIST SP800-181 instead of the security domains would be of great value in understanding how recruitment strategies, educational curricula and training programs should be adapted to better match demand and supply. Also, insight in the time equivalent of each of the tasks for a certain work role (all tasks were assumed to have an equal time expenditure in this research) would enhance one's understanding of what the proportion of time is that can be outsourced for each FTE. Lastly, the estimated impact on the cybersecurity gap depends on the GISWS predictions on for the size of the cybersecurity skills gap. As these predictions change and as artificial intelligence will be able to take on new tasks, it would be of interest to see whether the proportion of work that could be outsourced to artificial intelligence were to change.

In addition to repeating the research with improved data, further research and quantification of the counteracting effects of artificial intelligence implementation that could widen the cybersecurity skills gap is recommended. Ideally it would be possible to quantify it in a similar fashion as the reduction of the cybersecurity skills gap was calculated, allowing for a 'net impact' calculation. A large part of bridging this gap would be an understanding of the full set of new artificial intelligence implementation and security roles that would be created, and how the demand and supply are likely to develop given the various AI adoption scenarios. In order to do so it is key to understand what likely paths are for artificial intelligence system attack and defence. Many publications focus on the safety implications of artificial intelligence, and research on security implications should be expanded upon for the industry to be adequately informed and protected as AI adoption grows. Lastly, it would be interesting to research how other industries compare in their AI outsourceability given the rule set provided in this thesis. The AI outsourceability rule set has been formulated in a generic way to enable such research.

7. References

[1] E. van Luit, 'Valuing Skills: Towards a Better Methodology for Matching Competencies and Jobs', Nyenrode Business Universiteit, Breukelen, 2014.

[2] D. Raykov and M. Iqbal, 'Will Robots Take My Job?', *Will Robots Take My Job?*, 30-May-2017. [Online]. Available: https://willrobotstakemyjob.com/. [Accessed: 22-Aug-2017].

[3] C. B. Frey and M. Osborne, 'The Future of Employment', *Susceptible Are Jobs Comput.*, 2013.

[4] M. Iqbal, 'From Idea to 4M Page Views in 4 Weeks', *From Idea to 4M Page Views in 4 Weeks*, 05-Jun-2017.

[5] (ISC)², 'Cybersecurity Workforce Shortage Projected at 1.8 Million by 2022', *(ISC)² Blog*, 15 February. [Online]. Available: http://blog.isc2.org/isc2_blog/2017/02/cybersecurity-workforce-gap.html. [Accessed: 15-Oct-2017].

[6] S. Cobb, 'Mind this Gap: Criminal Hacking and the Global Cybersecurity Skills Shortage, a Critical Analysis', presented at the Virus Bulletin Conference, 2016, pp. 1–8.

[7] CISCO, 'Mitigating the Cybersecurity Skills Shortage: Top Insights and Actions from Cisco Security Advisory Services', CISCO, 2015.

[8] M. Matsubara, 'How Japan Is Aiming to Close the Cybersecurity Skills Gap Before Tokyo 2020 - Palo Alto Networks Blog', *CSO Perspective*, 15-May-2017.

[9] The Wharton School, University of Pennsylvania, 'How America Is Closing the Cybersecurity Skills Gap', *Knowledge@Wharton*, 16-Aug-2017.

[10] OECD, *OECD Digital Economy Outlook 2017*. OECD Publishing, 2017.

[11] General Auditor, 'The UK Cyber Security Strategy: Landscape Review', House of Commons, London, 2013.

[12] C. Lalan, 'IBM Watson to Tackle Cybercrime', *IBM News Room*, 10-May-2016.

[13] K. Kelly, 'The Myth of a Superhuman AI', *WIRED*, 25-Apr-2017.

[14] A. Ng, 'What AI Can and Can't Do', *Harvard Business Review*, 09-Nov-2016.

[15] R. Atkinson, 'Artificial Intelligence, Robotics, and the Future of Work: Myths and Facts', *Information Technology & Innovation Foundation*, 19-Sep-2017.

[16] R. Goodwins, 'Debunking the Biggest Myths about Artificial Intelligence', *Ars Technica*, 25-Dec-2015.

[17] J. McCarthy, 'What is Artificial Intelligence?' Stanford University, 2007.

[18] J. Fetzer, 'What is Artificial Intelligence?', in *Artificial Intelligence: Its Scope and Limits*, vol. 4, Dordrecht: Springer, 1990, pp. 3–27.

[19] S. Greengard, 'Cybersecurity Gets Smart', *Commun. ACM*, vol. 59, no. 5, pp. 29–31, Apr. 2016.

[20] C.-V. Oancea, 'Artificial Intelligence Role in Cybersecurity Infrastructures', *Int. J. Inf. Secur. Cybercrime*, vol. 4, no. 1, pp. 59–62, 2015.

[21] C. E. Landwehr, 'Cybersecurity and Artificial Intelligence: From Fixing the Plumbing to Smart Water', *IEEE Secur. Priv.*, vol. 6, no. 5, pp. 3–4, 2008.
[22] R. V. Yampolskiy, 'Artificial Intelligence Safety and Cybersecurity: a Timeline of AI Failures', *ArXiv Prepr. ArXiv161007997*, pp. 1–12, 2016.
[23] A. Randrianasolo, 'Artificial Intelligence in Computer Security: Detection, Temporary Repair and Defense', Texas Tech University, 2012.
[24] L. Tsado, 'Analysis of Cybersecurity Threats and Vulnerabilities: Skills Gap Challenges and Professional Development', PhD, Texas Southern University, 2016.
[25] C. Alexander, 'The Cybersecurity Skills Gap', no. December 2014.
[26] J. Hammerstein and C. May, 'The CERT Approach to Cybersecurity Workforce Development', Carnegie-Mellon, Hanscom, USA, 2010.
[27] A. McGerrick, L. Cassel, M. Dark, E. Hawthorne, and J. Impagliazzo, 'Toward Curricular Guidelines for Cybersecurity', presented at the SIGCSE, Atlanta, Georgia, 2014, pp. 81–82.
[28] W. Newhouse, S. Keith, B. Scribner, and G. Witte, 'National Initiative for Cybersecurity Education (NICE) Cybersecurity Workforce Framework', National Institute of Standards and Technology, Gaithersburg, MD, NIST SP 800-181, Aug. 2017.
[29] R. Slane, 'How O*NET Classification Helps Us Match Jobs and Skills', *Emsi*, 11-Jun-2013. [Online]. Available: https://www.economicmodelling.co.uk/2013/06/11/how-onet-classification-helps-us-match-jobs-and-skills/. [Accessed: 01-Nov-2017].
[30] J. Wilson, P. Daugherty, and N. Morini-Bianzino, 'The Jobs That Artificial Intelligence Will Create', no. Summer 2017, 23-Mar-2017.
[31] C. Le Clair and J. Gownder, 'The Future of White-Collar Work: Sharing Your Cubicle with Robots', Forrester Research, Jun. 2016.
[32] J. Manyika *et al.*, 'Harnessing Automation for a Future That Works', *McKinsey & Company*, Jan-2017. [Online]. Available: https://www.mckinsey.com/global-themes/digital-disruption/harnessing-automation-for-a-future-that-works. [Accessed: 30-Nov-2017].
[33] H. Feldmann, 'Technological Unemployment in Industrial Countries', *J. Evol. Econ.*, vol. 23, no. 5, pp. 1099–1126, 2013.
[34] R. Bellamy, 'The History of Cyber Security: Everything You Ever Wanted to Know', *SentinelOne*, 10-Mar-2017.
[35] T. Julian, 'Defining Moments in the History of Cyber-Security', *Infosecurity Magazine*, 04-Dec-2014.
[36] G. Press, 'This Week in Tech History: The Birth of the Cybersecurity and Computer Industries', *Forbes*, 01-Nov-2015.
[37] NATO Review, 'The History of Cyber Attacks: A Timeline', *Cyber - the good, the bad and the bug-free*, 17-Jun-2013.
[38] M. Warner, 'Cybersecurity: A Pre-history', *Intell. Natl. Secur.*, vol. 27, no. 5, pp. 781–799, Oct. 2012.
[39] D. Klaper and E. Hovy, 'A Taxonomy and a Knowledge Portal for Cybersecurity', in *Proceedings of the 15th Annual International Conference on Digital Government Research*, 2014, pp. 79–85.
[40] D. Craigen, N. Diakun-Thibault, and R. Purse, 'Defining cybersecurity', *Technol. Innov. Manag. Rev.*, vol. 4, no. 10, pp. 13–21, 2014.
[41] (ISC)², 'CISSP Exam Outline'. (ISC)², Apr-2015.
[42] National Initiative for Cybersecurity Education, 'The National Cybersecurity Workforce Framework'. NICE, 2013.
[43] A. Ross, 'Want Job Security? Try Online Security', *WIRED UK*, 25-Apr-2016.
[44] MarketsandMarkets, 'Cybersecurity Market by Solution & Service - 2022'. [Online]. Available: http://www.marketsandmarkets.com/Market-Reports/cyber-security-market-505.html. [Accessed: 15-Oct-2017].

[45] A. Carey, '(ISC)2 Information Security Global Workforce Study', (ISC)2, 2004.
[46] Frost & Sullivan, 'The 2015 (ISC)2 Global Information Security Workforce Study', (ISC)2, 2015.
[47] NICE, 'Unconnected Tasks & KSAs in NIST SP800-181', 22-Nov-2017.
[48] S. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach*, Third. Pearson, 2017.
[49] M. S. Bittermann, 'Artificial Intelligence (AI) Versus Computational Intelligence (CI) for Treatment of Complexity in Design', *Des. Comput. Cogn. DCC*, vol. 10, p. 8, 2010.
[50] W. Duch, 'What is Computational Intelligence and where is it going?', *Chall. Comput. Intell.*, pp. 1–13, 2007.
[51] A. Prabnhakar, 'Powerful But Limited: A DARPA Perspective on AI', 31-Jan-2017.
[52] N. Nilsson, *The Quest for Artificial Intelligence*. Stanford University, 2009.
[53] K. Warwick, *Artificial Intelligence*. Routledge, 2012.
[54] S. Adams *et al.*, 'Mapping the Landscape of Human-level Artificial General Intelligence', *AI Mag.*, vol. 33, no. 1, pp. 25–42, 2012.
[55] R. Sanz, 'Requirements for AGI: Using Requirements-Driven Processes to Build AGIs', presented at the AGI Summer School 2012, Reykjavik, Aug-2012.
[56] R. Chopra, *Artificial Intelligence: A Practical Approach*. S Chand & Co Ltd, 2012.
[57] H. Tipton and M. Krause, *Information Security Management Handbook*, 5th ed. CRC Press, 2003.
[58] J.-Y. Puigbo, A. Pumarola, C. Angulo, and R. Tellez, 'Using a Cognitive Architecture for General Purpose Service Robot Control', *Connect. Sci.*, vol. 27, no. 2, pp. 105–117, Apr. 2015.
[59] D. Poole, A. Mackworth, and R. Goebel, *Computational Intelligence: A Logical Approach*. New York/Oxford: Oxford University Press, 1998.
[60] D. Shahaf and E. Horvitz, 'Generalized Task Markets for Human and Machine Computation', in *AAAI*, 2010.
[61] J. E. Laird and R. E. Wray III, 'Cognitive Architecture Requirements for Achieving AGI', 2010, p. 6.
[62] S. Miller, C. Wagner, U. Aickelin, and J. M. Garibaldi, 'Modelling Cyber-security Experts' Decision Making Processes Using Aggregation Operators', *Comput. Secur.*, vol. 62, pp. 229–245, 2016.
[63] E. Tyugu, 'Artificial Intelligence in Cyber Defense', in *2011 3rd International Conference on Cyber Conflict (ICCC)*, 2011, pp. 1–11.
[64] V. Lavigne and D. Gouin, 'Visual Analytics for Cyber Security and Intelligence', *J. Def. Model. Simul. Appl. Methodol. Technol.*, vol. 11, no. 2, pp. 175–199, Apr. 2014.
[65] P. Xie, J. H. Li, X. Ou, P. Liu, and R. Levy, 'Using Bayesian networks for cyber security analysis', in *Dependable Systems and Networks (DSN), 2010 IEEE/IFIP International Conference on*, 2010, pp. 211–220.
[66] A. C. Schultz, 'Artificial Intelligence and Security: An Overview', in *afips*, 1899, p. 73.
[67] S. Dilek, H. Cakır, and M. Aydın, 'Applications of Artificial Intelligence Techniques to Combating Cyber Crimes: A Review', *Int. J. Artif. Intell. Appl.*, vol. 6, no. 1, pp. 21–39, Jan. 2015.
[68] B. Purgason and D. Hibler, 'Security Through Behavioral Biometrics and Artificial Intelligence', *Procedia Comput. Sci.*, vol. 12, pp. 398–403, 2012.
[69] D. Dasgupta, 'Computational Intelligence in Cyber Security', in *Computational Intelligence for Homeland Security and Personal Safety, Proceedings of the 2006 IEEE International Conference on*, 2006, pp. 2–3.
[70] A. Rehman and T. Saba, 'Evaluation of Artificial Intelligent Techniques to Secure Information in Enterprises', *Artif. Intell. Rev.*, vol. 42, no. 4, pp. 1029–1044, Dec. 2014.

[71] S. Abu-Nimeh, D. Nappa, X. Wang, and S. Nair, 'A Comparison of Machine Learning Techniques for Phishing Detection', in *Proceedings of the Anti-phishing Working Groups 2nd Annual eCrime Researchers Summit*, 2007, pp. 60–69.
[72] M. S. Boddy, J. Gohde, T. Haigh, and S. A. Harp, 'Course of Action Generation for Cyber Security Using Classical Planning', in *ICAPS*, 2005, pp. 12–21.
[73] G. Grieco, G. L. Grinblat, L. Uzal, S. Rawat, J. Feist, and L. Mounier, 'Toward Large-scale Vulnerability Discovery Using Machine Learning', in *Proceedings of the Sixth ACM Conference on Data and Application Security and Privacy*, 2016, pp. 85–96.
[74] U. Karabiyik and S. Aggarwal, 'Audit: Automated Disk Investigation Toolkit', *J. Digit. Forensics Secur. Law JDFSL*, vol. 9, no. 2, p. 129, 2014.
[75] C. R. Tschan, E. H. Kivelevitch, and K. J. Melcher, 'Advancing Intelligent Systems in the Aerospace Domain', presented at the AIAA Infotech@Aerospace (I@A) Conference, Boston, 2013, pp. 1–10.
[76] A. Carnevale, T. Jayasundera, and D. Repnikov, 'Understanding Online Job Ads Data', Center of Education and the Workforce of Georgetown University, Georgetown, 2014.
[77] Burning Glass Technologies, 'FAQ on Real-Time Jobs Data', *Burning Glass Technologies*, 09-Jan-2017. [Online]. Available: http://burning-glass.com/faq-on-real-time-jobs-data/. [Accessed: 27-Nov-2017].
[78] D. Restuccia, 'Job Market Intelligence: Cybersecurity Jobs, 2015', Burning Glass Technologies, 2015.
[79] T. Herbert and W. Markow, 'Illuminating the Cybersecurity Workforce', presented at the NICE Cybersecurity Working Group, Teleconference, 07-Nov-2016.
[80] D. Restuccia, 'Job Market Intelligence: Report on the Growth of Cybersecurity Jobs', Burning Glass Technologies, Mar. 2014.
[81] Center for Cyber Safety and Education, 'Global Information Security Workforce Study', *Center for Cyber Safety and Education*, 2017.
[82] S. Morgan, 'Cybersecurity Jobs Analysis', *Cybersecurity Ventures*, 07-Jun-2017.
[83] Cybersecurity Ventures, 'Cybersecurity Jobs Report', Herjavec Group, 2017.
[84] Digital Skills Committee, 'Parliamentlive.tv', 28-Oct-2014. [Online]. Available: http://www.parliamentlive.tv/Event/Index/a2f9595f-1699-446c-bf51-3df1084191f6. [Accessed: 27-Nov-2017].
[85] CyberSeek, 'Cybersecurity Supply and Demand Heat Map', *CyberSeek*, 2017. [Online]. Available: http://cyberseek.org/heatmap.html. [Accessed: 27-Nov-2017].
[86] W. Markow, 'Request CyberSeek Inquiry.msg', 27-Nov-2017.
[87] R. Sentz, 'How Should We Look At Jobs? A Discussion of Labor Market Data and Job Postings', *Emsi*, 09-Apr-2013.
[88] (ISC)², '2017 Global Information Security Workforce Study Benchmarking Workforce Capacity and Response to Cyber Risk', (ISC)², 2017.
[89] Burning Glass Technologies, 'CyberSeek.org Data 2016'. Oct-2015.
[90] Burning Glass Technologies, 'CyberSeek.org Data 2017'. Oct-2016.
[91] M. Chui, 'Artificial Intelligence the Next Digital Frontier?', *McKinsey Co. Glob. Inst.*, p. 47, 2017.
[92] A. Wilkins, 'Elon Musk Says Tesla's Self-Driving Tech Already as Good as a Human Driver', *Inverse*, 03-Nov-2017. [Online]. Available: https://www.inverse.com/article/38049-elon-musk-self-driving-autopilot-tesla. [Accessed: 30-Nov-2017].
[93] E. van Luit, '"Smarter Toys": How Artificial Intelligence-enabled Toys Are to Become Next-level Social Engineering Malicious Insiders', Cyber Security Academy, The Hague, 2017.
[94] M. Dirksen, E. van Luit, and C. Rugers, 'Lessons from the DARPA Cyber Grand Challenge: What Security Artificial Intelligence Do Well?', Cyber Security Academy, The Hague, 2017.

[95] B. Marr, 'Instead Of Destroying Jobs Artificial Intelligence (AI) Is Creating New Jobs In 4 Out Of 5 Companies', *Forbes*, 12-Oct-2017. [Online]. Available: https://www.forbes.com/sites/bernardmarr/2017/10/12/instead-of-destroying-jobs-artificial-intelligence-ai-is-creating-new-jobs-in-4-out-of-5-companies/. [Accessed: 29-Nov-2017].
[96] C. Stancombe *et al.*, 'Turning AI into Concrete Value: The Successful Implementers' Toolkit', Capgemini, 2017.
[97] N. Heath, 'Why AI could destroy more jobs than it creates, and how to save them', *TechRepublic*. [Online]. Available: https://www.techrepublic.com/article/ai-is-destroying-more-jobs-than-it-creates-what-it-means-and-how-we-can-stop-it/. [Accessed: 29-Nov-2017].
[98] F. Levy and R. Murnane, *The New Division of Labor: How Computers Are Creating the Next Job Market*. Russell Sage Foundation, 2005.
[99] A. Burt, 'Is there a "right to explanation" for machine learning in the GDPR?', *IAPP*, 01-Jun-2017. .
[100] I. Corona, B. Biggio, and D. Maiorca, 'AdversariaLib', *Pattern Recognition and Applications Lab*. [Online]. Available: http://pralab.diee.unica.it/en/AdversariaLib. [Accessed: 30-Nov-2017].
[101] J. Mokyr, C. Vickers, and N. L. Ziebarth, 'The History of Technological Anxiety and the Future of Economic Growth: Is This Time Different?', *J. Econ. Perspect.*, vol. 29, no. 3, pp. 31–50, Aug. 2015.
[102] S. Kessler, 'The Optimist's Guide to the Robot Apocalypse', *Quartz*, 09-Mar-2017.
[103] T. Pettinger, 'The Luddite Fallacy', *Economicshelp*, 15-Jan-2016.
[104] McKinsey Global Institute, 'Where Machines Could Replace Humans - and Where They Can't (Yet): International Automation', *Tableau*, 23-Jan-2017. [Online]. Available: https://public.tableau.com/views/InternationalAutomation/WhereMachinesCanReplaceHumans?%3Aembed=y&%3AshowVizHome=no&%3Adisplay_count=y&%3Adisplay_static_image=y&%3AbootstrapWhenNotified=true. [Accessed: 30-Nov-2017].
[105] McKinsey Global Institute, 'Where Machines Could Replace Humans - and Where They Can't (Yet): Automation by Sector', *Tableau*, 15-Jan-2017. [Online]. Available: https://public.tableau.com/views/AutomationBySector/WhereMachinesCanReplaceHumans?%3Aembed=y&%3AshowVizHome=no&%3Adisplay_count=y&%3Adisplay_static_image=y&%3AbootstrapWhenNotified=true. [Accessed: 30-Nov-2017].
[106] M. Chui, J. Manyika, and M. Miremadi, 'Where Machines Could Replace Humans - and Where They Can't (Yet)', *McKinsey & Company*, Jul-2016. [Online]. Available: https://www.mckinsey.com/business-functions/digital-mckinsey/our-insights/where-machines-could-replace-humans-and-where-they-cant-yet. [Accessed: 30-Nov-2017].

Appendix A – Specialties and Work Roles

Specialty	Role	Role content
Risk Management (RSK)	Authorizing Official/Designating Representative	Senior official or executive with the authority to formally assume responsibility for operating an information system at an acceptable level of risk to organizational operations (including mission, functions, image, or reputation), organizational assets, individuals, other organizations, and the Nation (CNSSI 4009).
	Security Control Assessor	Conducts independent comprehensive assessments of the management, operational, and technical security controls and control enhancements employed within or inherited by an information technology (IT) system to determine the overall effectiveness of the controls (as defined in NIST SP 800-37).
Software Development (DEV)	Software Developer	Develops, creates, maintains, and writes/codes new (or modifies existing) computer applications, software, or specialized utility programs.
	Secure Software Assessor	Analyzes the security of new or existing computer applications, software, or specialized utility programs and provides actionable results.
Systems Architecture (ARC)	Enterprise Architect	Develops and maintains business, systems, and information processes to support enterprise mission needs; develops information technology (IT) rules and requirements that describe baseline and target architectures.
	Security Architect	Ensures that the stakeholder security requirements necessary to protect the organization's mission and business processes are adequately addressed in all aspects of enterprise architecture including reference models, segment and solution architectures, and the resulting systems supporting those missions and business processes.
Technology R&D (TRD)	Research & Development Specialist	Conducts software and systems engineering and software systems research to develop new capabilities, ensuring cybersecurity is fully integrated. Conducts comprehensive technology research to evaluate potential vulnerabilities in cyberspace systems.
Systems Requirements Planning (SRP)	Systems Requirements Planner	Consults with customers to evaluate functional requirements and translate functional requirements into technical solutions.
Test and Evaluation (TST)	System Testing and Evaluation Specialist	Plans, prepares, and executes tests of systems to evaluate results against specifications and requirements as well as analyze/report test results.
Systems Development (SYS)	Information Systems Security Developer	Designs, develops, tests, and evaluates information system security throughout the systems development life cycle.
	Systems Developer	Designs, develops, tests, and evaluates information systems throughout the systems development life cycle.
Data Administration (DTA)	Database Administrator	Administers databases and/or data management systems that allow for the secure storage, query, protection, and utilization of data.

	Data Analyst	Examines data from multiple disparate sources with the goal of providing security and privacy insight. Designs and implements custom algorithms, workflow processes, and layouts for complex, enterprise-scale data sets used for modeling, data mining, and research purposes.
Knowledge Management (KMG)	Knowledge Manager	Responsible for the management and administration of processes and tools that enable the organization to identify, document, and access intellectual capital and information content.
Customer Service and Technical Support (STS)	Technical Support Specialist	Provides technical support to customers who need assistance utilizing client-level hardware and software in accordance with established or approved organizational process components (i.e., Master Incident Management Plan, when applicable).
Network Services (NET)	Network Operations Specialist	Plans, implements, and operates network services/systems, to include hardware and virtual environments.
Systems Administration (ADM)	System Administrator	Responsible for setting up and maintaining a system or specific components of a system (e.g. for example, installing, configuring, and updating hardware and software; establishing and managing user accounts; overseeing or conducting backup and recovery tasks; implementing operational and technical security controls; and adhering to organizational security policies and procedures).
Systems Analysis (ANA)	Systems Security Analyst	Responsible for the analysis and development of the integration, testing, operations, and maintenance of systems security.
Legal Advice and Advocacy (LGA)	Cyber Legal Advisor	Provides legal advice and recommendations on relevant topics related to cyber law.
	Privacy Officer/Privacy Compliance Manager	Develops and oversees privacy compliance program and privacy program staff, supporting privacy compliance, governance/policy, and incident response needs of privacy and security executives and their teams.
Training, Education, and Awareness (TEA)	Cyber Instructional Curriculum Developer	Develops, plans, coordinates, and evaluates cyber training/education courses, methods, and techniques based on instructional needs.
	Cyber Instructor	Develops and conducts training or education of personnel within cyber domain.
Cybersecurity Management (MGT)	Information Systems Security Manager	Responsible for the cybersecurity of a program, organization, system, or enclave.
	Communications Security (COMSEC) Manager	Individual who manages the Communications Security (COMSEC) resources of an organization (CNSSI 4009) or key custodian for a Crypto Key Management System (CKMS).
Strategic Planning and Policy (SPP)	Cyber Workforce Developer and Manager	Develops cyberspace workforce plans, strategies, and guidance to support cyberspace workforce manpower, personnel, training and education requirements and to address changes to cyberspace policy, doctrine, materiel, force structure, and education and training requirements.
	Cyber Policy and Strategy Planner	Develops and maintains cybersecurity plans, strategy, and policy to support and align with organizational cybersecurity initiatives and regulatory compliance.
Executive Cyber Leadership (EXL)	Executive Cyber Leadership	Executes decision-making authorities and establishes vision and direction for an organization's cyber and cyber-related resources and/or operations.
Program/Project Management (PMA) and Acquisition	Program Manager	Leads, coordinates, communicates, integrates, and is accountable for the overall success of the program, ensuring alignment with agency or enterprise priorities.
	IT Project Manager	Directly manages information technology projects.
	Product Support Manager	Manages the package of support functions required to field and maintain the readiness and operational capability of systems and components.
	IT Investment/Portfolio Manager	Manages a portfolio of IT investments that align with the overall needs of mission and enterprise priorities.
	IT Program Auditor	Conducts evaluations of an IT program or its individual components to determine compliance with published standards.

Cybersecurity Defense Analysis (CDA)	Cyber Defense Analyst	Uses data collected from a variety of cyber defense tools (e.g., IDS alerts, firewalls, network traffic logs) to analyze events that occur within their environments for the purposes of mitigating threats.
Cybersecurity Defense Infrastructure Support (INF)	Cyber Defense Infrastructure Support Specialist	Tests, implements, deploys, maintains, and administers the infrastructure hardware and software.
Incident Response (CIR)	Cyber Defense Incident Responder	Investigates, analyzes, and responds to cyber incidents within the network environment or enclave.
Vulnerability Assessment and Management (VAM)	Vulnerability Assessment Analyst	Performs assessments of systems and networks within the network environment or enclave and identifies where those systems/networks deviate from acceptable configurations, enclave policy, or local policy. Measures effectiveness of defense-in-depth architecture against known vulnerabilities.
Threat Analysis (TWA)	Threat/Warning Analyst	Develops cyber indicators to maintain awareness of the status of the highly dynamic operating environment. Collects, processes, analyzes, and disseminates cyber threat/warning assessments.
Exploitation Analysis (EXP)	Exploitation Analyst	Collaborates to identify access and collection gaps that can be satisfied through cyber collection and/or preparation activities. Leverages all authorized resources and analytic techniques to penetrate targeted networks.
All-Source Analysis (ASA)	All-Source Analyst	Analyzes data/information from one or multiple sources to conduct preparation of the environment, respond to requests for information, and submit intelligence collection and production requirements in support of planning and operations.
	Mission Assessment Specialist	Develops assessment plans and measures of performance/effectiveness. Conducts strategic and operational effectiveness assessments as required for cyber events. Determines whether systems performed as expected and provides input to the determination of operational effectiveness.
Targets (TGT)	Target Developer	Performs target system analysis, builds and/or maintains electronic target folders to include inputs from environment preparation, and/or internal or external intelligence sources. Coordinates with partner target activities and intelligence organizations, and presents candidate targets for vetting and validation.
	Target Network Analyst	Conducts advanced analysis of collection and open-source data to ensure target continuity; to profile targets and their activities; and develop techniques to gain more target information. Determines how targets communicate, move, operate and live based on knowledge of target technologies, digital networks, and the applications on them.
Language Analysis (LNG)	Multi-Disciplined Language Analyst	Applies language and culture expertise with target/threat and technical knowledge to process, analyze, and/or disseminate intelligence information derived from language, voice and/or graphic material. Creates and maintains language-specific databases and working aids to support cyber action execution and ensure critical knowledge sharing. Provides subject matter expertise in foreign language-intensive or interdisciplinary projects.
Collection Operations (CLO)	All Source-Collection Manager	Identifies collection authorities and environment; incorporates priority information requirements into collection management; develops concepts to meet leadership's intent. Determines capabilities of available collection assets, identifies new collection capabilities; and constructs and disseminates collection plans. Monitors execution of tasked collection to ensure effective execution of the collection plan.
	All Source-Collection Requirements Manager	Evaluates collection operations and develops effects-based collection requirements strategies using available sources and methods to improve collection. Develops, processes, validates, and coordinates submission of collection requirements. Evaluates performance of collection assets and collection operations.

Cyber Operational Planning (OPL)	Cyber Intel Planner	Develops detailed intelligence plans to satisfy cyber operations requirements. Collaborates with cyber operations planners to identify, validate, and levy requirements for collection and analysis. Participates in targeting selection, validation, synchronization, and execution of cyber actions. Synchronizes intelligence activities to support organization objectives in cyberspace.
	Cyber Ops Planner	Develops detailed plans for the conduct or support of the applicable range of cyber operations through collaboration with other planners, operators and/or analysts. Participates in targeting selection, validation, synchronization, and enables integration during the execution of cyber actions.
	Partner Integration Planner	Works to advance cooperation across organizational or national borders between cyber operations partners. Aids the integration of partner cyber teams by providing guidance, resources, and collaboration to develop best practices and facilitate organizational support for achieving objectives in integrated cyber actions.
Cyber Operations (OPS)	Cyber Operator	Conducts collection, processing, and/or geolocation of systems to exploit, locate, and/or track targets of interest. Performs network navigation, tactical forensic analysis, and, when directed, executes on-net operations.
Cyber Investigation (INV)	Cyber Crime Investigator	Identifies, collects, examines, and preserves evidence using controlled and documented analytical and investigative techniques.
Digital Forensics (FOR)	Law Enforcement /CounterIntelligence Forensics Analyst	Conducts detailed investigations on computer-based crimes establishing documentary or physical evidence, to include digital media and logs associated with cyber intrusion incidents.
	Cyber Defense Forensics Analyst	Analyzes digital evidence and investigates computer security incidents to derive useful information in support of system/network vulnerability mitigation.