Automation of Cybersecurity Work

Stefan Varga, Teodor Sommestad, and Joel Brynielsson

1 Introduction

Automated solutions have shown great potential within the cybersecurity field, and much of current research is directed towards automated solutions. By automation, we here mean the execution of some task by a machine agent (usually a computer) that was previously carried out by a human, according to Parasuraman and Riley's definition [33, p. 231]. Examples of machine agents in cybersecurity, are intrusion detection systems and vulnerability assessment tools. In a future scenario where high-quality cybersecurity automated solutions are widespread, a plausible reality may be that:

- complex cyber operations, both offensive and defensive, are easy to execute, even for actors who lack resources,
- the ones who possess the most powerful automated solutions have dominant positions within the cyber domain, and
- some current (human) work roles are obsolete, significantly changed, or simplified.

S. Varga (✉)
KTH Royal Institute of Technology, Stockholm, Sweden
Swedish Armed Forces Headquarters, Stockholm, Sweden
e-mail: svarga@kth.se

T. Sommestad
FOI Swedish Defence Research Agency, Stockholm, Sweden
e-mail: teodor.sommestad@foi.se

J. Brynielsson
KTH Royal Institute of Technology, Stockholm, Sweden
FOI Swedish Defence Research Agency, Stockholm, Sweden
e-mail: joel@kth.se

T. Sipola et al. (eds.), *Artificial Intelligence and Cybersecurity*,
https://doi.org/10.1007/978-3-031-15030-2_4

This chapter seeks to examine the preconditions for further automation of cybersecurity tasks, especially in light of developments in the artificial intelligence, AI, field. The chapter mainly focuses on the third aspect listed above, namely: automation in relation to human work. Even if increased automation is desirable, there are several obstacles that have to be overcome first. For example, research on intrusion detection has not yet nullified the role of human system operators, and humans are still needed in the loop [22, 40]. Why have these efforts failed?

More precisely, we seek to shed light on the following research questions:

1. What variables affect how hard a cybersecurity role is to automate?
2. How likely is it that current cybersecurity roles will be automated?
3. What variables constrain the potential for automation of today's cybersecurity roles?

A keystone of the analysis is the *National Initiative for Cybersecurity Education*, NICE, framework for cybersecurity work roles, brought forward by the U.S. National Institute of Standards and Technology, NIST [31]. This document delineates cybersecurity roles by describing various tasks and the demands that the fulfillment of those tasks require in terms of knowledge, skills, and abilities.

2 Cybersecurity Automation Research

Progress is continuously being made within the fields of AI and automation. At the same time there is an urgency to cope with cybersecurity threats due to the increasingly severe implications that breaches may bring. There is, however, no comprehensive research that addresses the intersection of these perspectives: the full automation of cybersecurity work. This subject has, to our knowledge, only been treated as a cybersecurity issue synoptically. This chapter seeks to outline relevant theories and models that describe what is easy and what is hard to automate, both in general and specifically for cybersecurity. The characteristics of this topic is cross-disciplinary, and the research questions addressed in this chapter are neither fully answered, nor extensively treated in the cybersecurity literature. Relevant related work can instead be found within a variety of other (academic) fields. Work related to the two first research questions concerns economics, while references for the last research question, which mainly deals with technical issues, can be found in literature about technology development. The references in this chapter are therefore attributed mainly to fields other than cybersecurity.

2.1 Variables That Influence Automation in General

Even if automation largely can be seen as a cornerstone of human progress [26], empirical research that unequivocally states what circumstances that make a human

task easy or hard to automate, is hard to find. For instance, comprehensive studies of successful or failed automation attempts are difficult to find. The earliest identified attempt to create comprehensive assessment criteria to this respect was, to our knowledge, made by Frey and Osborne [21]. They used variables from the U.S. Department of Labor's Occupational Information Network, known as O*NET, which is used to describe requirement levels for various professions. They then let a panel of computer science machine learning researchers judge whether 70 randomly chosen professions were possible to automate or not. Panel members read the job descriptions and answered the question: "Can the tasks of this job be sufficiently specified, conditional on the availability of big data, to be performed by state of the art computer-controlled equipment" [21, p. 30].

The resulting judgments were then used as a baseline for a regression model that calculated the importance of the different O*NET variables. These steps resulted in the following list of variables deemed as relevant for whether an occupation is technically feasible to automate (e.g., whether they were to be seen as potential bottleneck indicators for automation, given data availability): perception and manipulation, creative intelligence, and social intelligence. It should be noted that some of the correlations, like the requirement for originality (as part of creative intelligence), were nonlinear.

Besides Frey and Osborne, others have performed similar analyses with other sets of variables. McKinsey [30], for example, used 18 variables split into five groups [30, p. 4], the groups being: sensory perception, cognitive capabilities, natural language processing, social and emotional capabilities, and physical capabilities. These types of requirements were not chosen for tasks that were particularly easy or hard to automate. They were rather used to reason about the requirement levels for various professions, and if automation to this level was feasible with today's technology or not. One example of a capability level, is the requirement for *natural language*, where the requirement level ranges from no requirements to "... nuanced human interaction" [30, p. 120]. To calibrate the meaning of these requirements, they set the highest level to correspond to the skill-level of the best quarter (25%) of the workforce. Van der Zande et al. [47] fused the McKinsey model with a simpler model conceived by Autor et al. [2]. They characterized the various professions into two new dimensions. First, on an axis ranging from whether they consisted mainly of routine tasks or not, to if they involved encounters of many novel situations or not. Second, on an axis ranging from if they primarily required physical capabilities, to whether they instead required cognitive capabilities. When van der Zande et al. [47] coupled the McKinsey analysis to these two new dimensions they, perhaps unsurprisingly, found that routine work is easier to automate than varied work.

Arntz et al. [1] used Frey and Osborne's analysis [21] as a blueprint to determine to what extent jobs in the United States are prone to automation. They suggested that this susceptibility was related to several variables divided into four categories: the characteristics of the workers (3 variables), the skill of the workers (3), the general characteristics of the work (11), and the characteristics of the tasks (25). Most of these variables were related to the automation assessments made by Frey and Osborne.

Suta et al. [43] aggregated variables used in studies such as the ones mentioned above, into factors that hinder or promote automation. The aggregation of the variables show only minor differences compared to the variables used by Frey and Osborne [21].

In a survey, Deloitte [16] also found that a major obstacle when new technology is introduced, is the lack of qualified and competent personnel who can use it. In other words, a certain kind of skill-set is required for the use of automated solutions. The following competencies were identified as desirable:

- specialist skills within the field of automation and AI to be able to create and deploy solutions,
- general abilities and subject matter knowledge possessed by users and various types of specialists (for example, economists), to be able to understand and make use of automated solutions in their respective fields,
- skills that can be leveraged to supplement the proposed automated solutions, when they are perceived to have deficiencies, e.g., to assess situations and to utilize interpersonal communication.

In addition to these competencies there are several other variables that can be assumed to greatly influence whether a profession or some part of it can or should be automated:

- Tasks that are performed with a high frequency by an expensive personnel category, e.g., personnel with high salaries, ought to be more prone for automation than tasks that are carried out seldom by low-cost personnel. In a development and innovation perspective, profitability and cost efficiency increase the market potential for automation [32].
- Resistance from personnel categories that may be negatively affected by automation is another variable that can affect whether a profession or task is automated or not. In some jobs, especially those that can be characterized by the so-called 3Ds: *Dirty*, *Dangerous*, and *Demanding* [10], such resistance can be believed to be low. For other jobs, especially ones that are perceived to be stimulating and creative, resistance can be expected to be greater.
- Resistance from potential *users* is another imaginable obstacle. According to the *Technology Acceptance Model* [27], information systems that are perceived to be useful and easy to use, are often readily accepted by users. For the opposite case, acceptance can be lower.
- Ethical and legal questions may affect whether a task is allowed to, or should, be automated. The question of liability for self-driving cars, or laws about personal privacy, serve as pertinent examples.

The motive for listing other potential factors that affect automation, is to highlight the existence of factors besides the ones that are brought up in this chapter.

2.2 Variables That Influence Automation of Cybersecurity Work

The analyses described this far are generic and pertain to the labor market in general. They encompass a range of jobs where only a fraction is connected to cybersecurity. For cyber jobs, Frey and Osborne [21] found, for example, that "[c]omputer systems analyst" jobs are difficult to automate, while "[c]omputer operators" are relatively easy to automate. Table 1 shows some cybersecurity jobs' proneness for automation according to Frey and Osborne. The previously mentioned McKinsey analysis [30] judged that 51% of the total hours worked in "[t]echnology, media and telecom", can be automated.

Current and emerging technologies hint about the limits of future automation. Technologies that can handle vulnerabilities and logs, prevent data loss, perform authentication control (including to networks), and advanced antivirus software, are common today. Technologies that ensure safety-critical development processes and services that extract intelligence from data, are less common. It has been argued, however, that any part of a job description that can be reduced to an algorithm or computational process can be automated [34], even if today's technology does not permit it. When tasks have been automated, the human part of the work could be simplified, reduced or transformed to some extent. Antivirus software, for example, checks software—a task that system administrators otherwise would have to do manually. The use of new technologies can also bring changes to current work practices. When, e.g., the signing of legally binding contracts is substituted for digital signatures, new additional cybersecurity work is probably needed to fully ensure the integrity of the involved IT systems. Another driver for cybersecurity, is the need to defeat tools and technologies used by attackers (and vice versa). This means that there is a cyber "arms race" going on. An example is when automatically generated attack code [3] forces defenders to rethink or completely redesign their intrusion detection processes. Another example is when the AI automation itself is the target for adversarial AI examples specifically designed to defeat cybersecurity

Table 1 The probability for cybersecurity-related jobs being automated according to Frey and Osborne [21]

Rank	Profession	Probability of automation
32	Computer systems analysts	0.0065
69	Computer and information research scientists	0.015
110	Database administrators	0.03
118	Computer and information systems managers	0.035
208	Information security analysts, web developers, and computer network architects	0.21
212	Computer occupations, all other	0.22
405	Computer, automated teller, and office machine repairers	0.74
428	Computer operators	0.78

efforts [14]. We assert that one of the goals of cybersecurity research is to automate to the widest extent possible, all in line with the goal of engineering at large. There are, however, oftentimes technical obstacles that prevent successful outcomes of automation efforts.

As an example case we highlight efforts to automate intrusion detection; a central problem in cybersecurity. This is an active research field that has gained much attention, but results this far have unfortunately not led to full automation of surveillance work in practice. Most approaches for automating surveillance tasks rely heavily on rule-based solutions [50, 51]. Existing rules are adapted and developed to suit organizational needs. This strand of research ultimately aims to develop models that are able to correlate events to automate threat classification and threat mitigation processes with different kinds of machine learning techniques. Possible reasons for the mediocre performance of event correlation are erroneous simplifications of the attack process, or reliance on incorrect information about the state of the IT system [39, 41]. Sommer and Paxson [37] propose the following explanations to why research based on machine learning approaches has failed:

1. the high cost of classification errors (e.g., sensitivity for false positives is too high due to a large amount of traffic),
2. the variety of formats in input data (e.g., numerous protocols with varying frequency of occurrence and content),
3. the lack of training data (e.g., realistic network traffic, where both attack-related and normal traffic flows have been appropriately labeled),
4. ambiguous output data (e.g., questions about the exact meaning of an anomaly and how it should be further investigated),
5. difficulties in determining the value or feasibility of proposed solutions (largely as an effect of items 3 and 4).

A second central problem in cybersecurity is to detect and identify software vulnerabilities [29]. The lack of automated *test oracles* that determine whether a function or calculation is adequate or not, is a problem [4]. Humans are as a rule still needed to conclusively judge whether a piece of software performs correctly or not. There are, however, many suggestions on how code reviews could be performed by computers. One proposal is to use formal methods to prove that software has certain secure properties. Formal methods, in short, use a specification of what the software is supposed to do, and then typically try to check that the software does precisely that and nothing else. This approach requires a human to set up an appropriate testing environment for the problem [17, p. 1176]. Another route to solve the problem is through more lightweight (semiformal) methods. Such methods do not require massive efforts during the problem formulation phase. However, although state of the art methods are fairly effective, they tend to result in additional work to ensure the safety of software [36]. The extra work is presumed to reduce the need for security-related work in the long run. A third research area targets vulnerability assessments of larger interconnected systems, computer networks and organizations. In this respect, many approaches model vulnerabilities

and conceivable attacks in the form of trees or graphs, in approximately the same way as fault trees are used to calculate risk or availability values.

The research community has produced several variants of modeling languages and associated software tools for the purpose of conducting vulnerability analyses using, e.g., graph-based models [28]. The actual use of such tools among practitioners is probably not widespread, neither in government, nor in the commercial sector. The sole test available that compares the data output produced by such a tool with the capabilities of actual attackers, show that the predictions provided are poor [41]. These kinds of disheartening results, however, are not only representative for cybersecurity research that models attacks in the form of graphs. In an extensive overview over research that seeks to quantify elements in security-related issues, Verendel [48] found that nearly no research had been validated with empirical data, and that the underlying research assumptions were not empirically validated either. In conclusion, it is difficult to assess whether any existing vulnerability assessment method actually works, if it has not been tested properly with realistic data. A commonly used excuse, or reason to waive this deficiency, is to refer to some need for secrecy. In general, one's cybersecurity stance appears to be a closely guarded and vigorous protection of the properties of so-called zero-day exploits [25], as suggested by the substantive underreporting of cybercrime [46].

In relation to the factors related to automation identified in other domains, the examples given above illustrate requirements on creativity (e.g., when possible ways of attacking a system shall be identified in vulnerability analyses), requirements on social interaction (e.g., ambiguous data from detection systems), and the need for data (e.g., to train or test solutions).

3 Method

This section describes the method that was employed to answer our three research questions. Section 3.1 describes the NIST NICE framework. Section 3.2 outlines the evaluation criteria for automation, while Sect. 3.3 describes the assessment process. Finally, in Sect. 3.4 the process of aggregating the assessments is described.

3.1 The Content of Cybersecurity Tasks

This study uses the U.S. National Institute of Standards and Technology (NIST) *National Initiative for Cybersecurity Education (NICE) Cybersecurity Workforce Framework* [31], henceforth the NICE framework, as a reference for the work that needs to be carried out to increase the level of cybersecurity. It should be noted that there exist other similar frameworks. Examples include *The Cyber Security Body*

of Knowledge (CyBOK),[1] the British NCSC[2] *Certification for Cyber Security/IA Professionals* framework, and the U.S. military DoD Directive 8570, originally from 2005. This latter-mentioned directive was later replaced by DoD Directive 8140, which in turn drew heavily on NIST publications.

The NICE framework was chosen because it has several advantages compared to other alternatives. It was conceived and prepared by the U.S. government in cooperation with industrial partners, and is in our judgment very detailed and predominantly well-structured. It is also widely spread and used, not least as a basis for further work. Moreover, it is regularly updated.

A factor that speaks against the use of NICE is that it is adapted to U.S. circumstances that do not necessarily fully reflect realities in other countries. Chen et al. [9, p. 63] point out that the descriptions of knowledge, skills and abilities in the NICE framework omit mentioning many details about the work's cognitive and collaborative aspects. There are also, in our opinion, some less well-worked parts in the document (see Sect. 5.1.1).

The content of the NICE framework can be used in many different ways. The analyzed version contains a high-level classification that divides security work into seven categories; further, there are 33 specialty areas in which people can work, and there are 52 distinct roles. Every role is coupled to tasks, knowledge requirements, skills and abilities that a person who should master it, should have. In the framework version used for this chapter, a total of 1006 tasks, 589 knowledge requirements, 365 skills and 176 abilities were listed [31].

This chapter, like previous studies [21], focuses on skills and abilities for specific tasks. These skills and abilities are the factors that form the foundation from which assessments are made regarding how easy or hard it is to automate various roles and specialty areas. Some examples of how skills and abilities are listed in the NICE framework are given in Table 2.

Table 2 Examples of skills and abilities as listed in the NICE framework [31]

Skills	Abilities
Skill in reviewing and editing target materials	Ability to conduct vulnerability scans and recognize vulnerabilities in security systems
Skill in configuring and utilizing network protection components (e.g., Firewalls, VPNs, network intrusion detection systems)	Ability to operate common network tools (e.g., `ping`, `traceroute`, `nslookup`)
Skill to analyze strategic guidance for issues requiring clarification and/or additional guidance	Ability to monitor measures or indicators of system performance and availability
Skill in creating and utilizing mathematical or statistical models	Ability to identify the requirements of in-process accounting for communications security (COMSEC)

[1] https://www.cybok.org/.

[2] https://www.ncsc.gov.uk/.

The meaning of the term knowledge is ambiguous. What knowledge is, and how it should be defined, has been the subject for lively discussions through the years. The imprecise meaning of the term is duly noted in the NICE framework. For the sake of this chapter, neither further problematization, nor additional extensive discussions on this subject, is necessary. We simply assert that knowledge can be codified into the memory of computers with relative ease. This (our) standpoint diminishes or abolishes the knowledge requirement criterion in the assessment of whether a role is possible to automate or not. This assertion is consistent with how earlier similar analyses have handled this question.

3.2 Assessment Criteria

As described in Sect. 2.1, previous research largely agrees on what factors affect whether a task is easy or hard to automate. The four criteria used in this study (see Table 3) are in line with these. The first three mirror the barriers for automation used by Frey and Osborne, whilst the fourth is used to refine (quantify) their premise of data availability [21].

The motive for why the criterion "existing statistical data" has been included in the assessment, is the assumption that the availability of data is important for cybersecurity. It has been shown that data availability can be a problem. The reasons for the unavailability of data can be that it often is sensitive and difficult to use due to security (confidentiality) concerns, and because many tasks are performed in environments where the meaning of specific data points is unclear, e.g., if

Table 3 Criteria used to grade abilities and skills, and cases that correspond to simple and difficult automation prerequisites, respectively. These represent the extremes in our grading scheme 1–5

	The nature of the skill or ability	Easy to automate (1)	Difficult to automate (5)
1	Requirements for creativity	There is only one (natural) way to perform the task that does not vary over time	The task may require that action alternatives no one has previously thought of be identified and applied
2	Requirements for social interaction	Does not require any interaction with humans	Requires situation-adapted and/or dynamic interpersonal communication with nuances
3	Requirements for physical work	Can be solved completely within a computer. Requires no physical work	Requires varied fine motor work with little tolerance for errors
4	Existing statistical data	High-quality basic data in sufficient quantity is available. Data produced to describe or document the work is available	Data does not exist, or exists to a very limited extent. There are very few cases to learn from

certain traffic patterns are indeed part of an attack or not (see Sect. 2.2). These two factors contribute to the difficulties in extracting and making high-quality conclusive data available for training of machine learning models. Furthermore, these are also prohibiting factors for the automation of cybersecurity work, and therefore motivates the introduction of this assessment variable.

The characteristics of the skills and abilities were judged according to the four criteria (see Table 3) on a scale ranging from one (1) to five (5), where one (1) corresponds to that a task is easy to automate, and five (5) corresponds to that it is hard. The assessors also judged whether the NICE description of the skill or ability was at all understandable.

3.3 Assessment Process

The 541 descriptions (176 skills and 365 abilities) were extracted and assessed by four researchers. They were:

- a Ph.D. (computer science) and deputy research director within the cybersecurity field,
- an associate professor (computer science) and research director working with decision support systems,
- a Ph.D. student (computer science) and military officer with a background from intelligence and cybersecurity,
- a master of laws (LL.M.) graduate specializing in cyber operations.

The work started with a calibration round where 30 randomly chosen task descriptions were assessed independently. This initial work was carried out at different geographical locations, without any communication between the assessors. It was later found that the researchers in some cases had to acquire additional background material from the NICE framework to make fair assessments. The need for this arouse from the desire to get an understanding for the context of the various descriptions of skills and abilities. The correlation coefficients between the assessments in the calibration round were overall moderately positive (e.g., ranging between 0.3–0.6 for the "requirements for social interaction" criterion), but also highlighted some differences in how the four researchers interpreted and implemented the criteria. The largest difference was detected for the "requirements for creativity" criterion.

After the initial calibration round, a seminar was held. Detected differences as well as other results were discussed. The seminar resulted in a refined scale. The seminar also improved consensus as of how to interpret the NICE texts, and thus probably contributed to a higher quality in the assessments for the main body of the data. After the seminar, the assessment work was continued with the assessments of all 541 descriptions. It could be noted that some differences between the views of the assessors remained even after the seminar. The mean deviations

were 0.7 for "existing statistical data", 0.6 for "requirements for creativity", 0.5 for "requirements for social interaction", and 0.1 for "requirements for physical work".

In Table 4, the correlation coefficients between the four researchers' assessments are listed. They varied between 0.24 and 0.76 with a mean value of 0.45. The differences were probably not a result of a faulty research design, but rather the fully natural variance of the competencies of the researchers due to their backgrounds. The mean value can thus be seen as an approximation of how a skill and an ability can be viewed from various perspectives. Overall mean values of the different criteria can also be seen in Table 4. A deeper data analysis (not shown here) indicates that the requirements for creativity and statistical data resemble normal distributions, with their mean values in the middle of the scale. The frequency of requirements for social interaction is linearly decreasing, which means that there are significantly more tasks with low requirements than high. Requirements for physical work are almost nonexistent.

3.4 Aggregation and Analysis

The assessments resulted in judgments about each and every skill and ability with regard to what requirements they have on (1) creativity, (2) social interaction, (3) physical work, and (4) how difficult it is to produce statistical data to be used for machine learning (or similar), in order to train a computer to perform a task. Mean values were used to obtain an aggregated answer from the panel. The examples displayed in Table 2 were used to obtain the assessments in Table 4.

The values indicate what is easy and hard to automate, given our method. The ability "to operate common network tools (e.g., `ping`, `traceroute`, `nslookup`)" receives lower values than the skill "to analyze strategic guidance for issues requiring clarification and/or additional guidance" on all four criteria, i.e., has fewer obstacles for being automated. The mean value of the different judgments can be seen as an indication on whether a skill or ability can be performed by a computer. A word of caution is in place here, because several additional factors not included in our model, can also come into play:

- As has been described in, e.g., Sect. 2.1, previous studies found nonlinear relations between the criteria and proneness for automation. It is imaginable that the highest level of creativity is next to impossible to achieve, when at the same time the preceding level is quite easily reached, e.g., the "distance" between the requirement levels vary in a nonlinear fashion. Similarly, it is also possible that the requirement for physical work at the highest level is unattainable. The obstacles for automation can in other words be an exponentially increasing function of the criteria. A way to remedy this problem, related to our method, can be to further "calibrate" the different assessors, that is, to make sure that they think about the problem in the same way through additional seminars or rules.

Table 4 Examples of skills and abilities together with the panel's assessments. Mean deviation between assessors can be seen within brackets

Skill or ability	Requirements for creativity	Requirements for social interaction	Requirements for physical work	Existing statistical data
Skill in reviewing and editing target materials	2.0 (0.5)	2.5 (1.5)	1.3 (0.4)	2.8 (0.4)
Skill in configuring and utilizing network protection components (e.g., Firewalls, VPNs, network intrusion detection systems)	2.5 (1.0)	1.0 (0.0)	1.0 (0.0)	2.0 (1.0)
Skill to analyze strategic guidance for issues requiring clarification and/or additional guidance	3.3 (0.4)	3.3 (0.4)	1.3 (0.4)	4.0 (0.5)
Skill in creating and utilizing mathematical or statistical models	3.0 (1.5)	1.0 (0.0)	1.0 (0.0)	2.8 (0.9)
Ability to conduct vulnerability scans and recognize vulnerabilities in security systems	2.0 (0.5)	1.0 (0.0)	1.0 (0.0)	2.5 (0.8)
Ability to operate common network tools (e.g., `ping`, `traceroute`, `nslookup`)	1.8 (0.4)	1.0 (0.0)	1.0 (0.0)	1.8 (0.8)
Ability to monitor measures or indicators of system performance and availability	1.5 (0.5)	1.0 (0.0)	1.0 (0.0)	2.3 (1.3)
Ability to identify the requirements of in-process accounting for communications security (COMSEC)	2.3 (0.9)	1.3 (0.4)	1.3 (0.4)	2.7 (0.9)

- It may be difficult to overcome several obstacles by combining different techniques. It could, for example, be more difficult to conceive a creative solution that is able to act in a social context, than it is to create separate solutions; where one is creative and the other one social, for the same problem. The difficulty to automate can in other words consist of the product of the criteria.
- There may exist interactions between the criteria that make it either easier or harder to automate them. A requirement for creativity may be easy to fulfil if there is an abundance of available statistical data to utilise. This problem would perhaps be very hard, or close to impossible, to solve without *any* statistical data. This serves to illustrate that the availability of statistical data affects/constrains the values of the *other* criteria.
- A single criterion can sometimes be decisive for whether a skill or an ability can be automated at all (even if all the other criteria suggest that it can). It may be the case that *any* task that requires a high or very high social competency, in principle is impossible to automate. This could be seen as a definite obstacle for automation, rather than the maximum value for the four criteria combined.

Due to the reasons given above, calculations were made with other models besides the mean value analysis. In total, we used five different models to compute the numerical values of the criteria (k_1–k_4). All models yield a value for each skill or ability, based on the scores from all four assessors. The five models are described in Table 5. Model one is the easiest and most straightforward. When nothing else is stated, the values from this model are used in discussions. The motive for our preference for model one, is that simple predictive models often outperform more complex ditto [20].

All the computed values (according to the mentioned models) provide information about whether a skill or ability is susceptible to automation in the form of a scalar (e.g., 3.3 or 47.2), where a low value relative to the other computed values according to the same model, suggest that automation is easy. The maximum values that can be obtained vary with the different models—from the value five (5) in model five, to 625 in model three. To achieve commensurability the values were normalized

Table 5 The five different models that were used to assess automability

Model	Explanation	Formal description
1	The effect is simple, additive, and linear. The sum of the four criteria is used	$\sum_{i=1}^{4} k_i$
2	The effect is exponential. The sum of exponential functions for the four criteria is used	$\sum_{i=1}^{4} 2^{k_i}$
3	Interactions make it much more difficult if several different obstacles need to be overcome. The product of the four criteria is used	$\prod_{i=1}^{4} k_i$
4	Access to statistical data (k_4) limits how much the other criteria affect	$\sum_{i=1}^{3} \min(k_i, k_4)$
5	The highest value for the various criteria is decisive/limiting	$\max_{i=1,\ldots,4} k_i$

against an index by dividing every computed value with the overall lowest value given by that model. The computed values are thus comparable relative to each other, with value one (1) being the easiest skill or ability to automate within each model. Then, aggregation with the structure of the NICE framework was done, in order to derive to what extent roles, specialty areas and categories can be automated. The mean values of the normalized skills and abilities were used, to this respect.

In addition to what is mentioned above, it can be concluded that the technology level and the prerequisites to meet the demands for automation, are on different levels today. There is nothing to suggest that they will not diverge even more in the future. It is entirely possible to argue that it is possible to meet the demands of physical work with, e.g., autonomous lawn mowers, and to solve a range of tasks that require fine motor ability, e.g., the peeling of prawns, while at the same time computers still do not have the capability to adequately produce synopses of texts (level three for social interaction in our scale). The present technological level, and thereby the prerequisites to meet the demands for the various criteria, can be uneven and also vary unevenly. In what way such developments affect the evaluation scores is tested in a *what-if* analysis, in which all combinations of the five levels for all the four criteria, is set to be unattainable. For each such combination (a total of 625), binary values indicating whether a skill or an ability is possible to automate or not, were set. The proportion of skills and abilities that could be automated for 21 selected scenarios were then aggregated to the seven categories of the NICE framework.

4 Results

In the following, results describing to what extent roles, specialty areas and categories can be automated according to the five models in Table 5, are presented. The values in the tables can be used for comparisons, e.g., the meaning of the value 2.6 is that a role is twice as hard to automate than a role that has the value 1.3. The assessments are then contrasted to 21 selected scenarios that represent different contexts. The figures in this part (Sect. 4.1) shall be seen as a measure of to what extent, or proportion, it is possible to automate cybersecurity work in each scenario.

The most detailed result presented in this chapter is presented in Table 6, in which it is shown to what extent NICE roles can be automated according to the five models. The presented colors, which are based on a mean (average) of the results of the models, show how easy (green) and hard (red) the different roles are to automate. The role *Database Administrator* is easiest to automate in all five models. The value 1.00 was therefore set as a base (index) for this case. The possibilities to automate other roles are, hence, expressed by how easy or hard they are to automate relative to the *Database Administrator* role. The results should be interpreted as, for example, the role *Systems Developer* being 26 percent (index 1.26) harder to automate than the *Database Administrator* role according to model four, and at the same time 160 percent (index 2.60) harder to automate according to model three. The mean value

indicates that the work of a *Systems Developer* is 65 percent harder (index 1.65) to automate than the work of a *Database Administrator*. Moreover, the table shows that the two roles *Cyber Legal Advisor* and *Executive Cyber Leadership*, are the ones that are hardest to automate according to all models.

The results aggregated into specialty areas are shown in Table 7. In this table some patterns emerge more clearly:

- Tasks that can be regarded as purely technical, such as database administration (e.g., *Data Administration*), network maintenance (e.g., *Cyber Defense Infrastructure Support*), and programming (e.g., *Software Development*), are relatively easy to automate.
- Tasks that require technical knowledge and are also of analytical nature (e.g., *Systems Development*), or that require coordination efforts with other functions (e.g., *Cybersecurity Management*), can be found in the middle of the scale.
- Intelligence work (e.g., *Collection Operations* and *Threat Analysis*) is relatively hard to automate.
- Areas associated with high levels of responsibility (e.g., *Executive Cyber Leadership*), and those that regularly require that a manifold of complex issues need to be weighed together (e.g., *Legal Advice and Advocacy*), are the hardest to automate.

The abovementioned tendencies can be discerned in all models.

At the highest level in the NICE framework, categories of cybersecurity work are listed. Table 8 shows to what extent these categories can be automated. There are less discrepancies at this level compared to the more detailed results presented in Tables 6 and 7. However, it can be noted that the categories *Operate and Maintain*, *Protect and Defend*, and *Investigate*, are generally easier to automate than the other categories.

4.1 Scenarios

There are a total of $5^4 = 625$ possible combinations of the four used criteria. Each and everyone can be regarded as a scenario—a hypothetical situation where the potential for automation has reached a certain level. In the scenarios, this level is expressed by the values of our evaluation criteria (1–5) according to Table 3. To clarify the meaning of this reasoning, we list some examples in Table 9. It could, for example, be imagined that technology has reached, or will reach, a level so that it becomes possible to satisfy the needs of:

- level three (3) on creativity, where multiple variables need to be weighed together and uncertainty has to be accounted for (which, for example, can be manifested by a capability to accurately target specific individuals with advertising),

Table 6 Roles in the NICE framework. The figures are to be read as a measure of how easy/hard it is to automate a role relative to other roles. The easiest role to automate in all models has index 1.0 (the *Database Administrator*)

Role	Mean	Model				
		1	2	3	4	5
Database Administrator (OM-DTA-001)	1.00	1.00	1.00	1.00	1.00	1.00
Data Analyst (OM-DTA-002)	1.13	1.07	1.16	1.27	1.04	1.13
Network Operations Specialist (OM-NET-001)	1.27	1.10	1.21	1.82	1.07	1.16
Cyber Operator (CO-OPS-001)	1.34	1.17	1.29	1.86	1.13	1.24
Software Developer (SP-DEV-001)	1.41	1.21	1.35	1.99	1.16	1.31
Cyber Defense Infrastructure Support Specialist (PR-INF-001)	1.42	1.20	1.34	2.14	1.18	1.23
Cyber Defense Analyst (PR-CDA-001)	1.45	1.22	1.39	2.18	1.19	1.27
Secure Software Assessor (SP-DEV-002)	1.47	1.24	1.41	2.18	1.18	1.34
Cyber Defense Incident Responder (PR-CIR-001)	1.49	1.25	1.41	2.21	1.22	1.33
Cyber Defense Forensics Analyst (IN-FOR-002)	1.53	1.23	1.40	2.51	1.18	1.34
Communications Security (COMSEC) Manager (OV-MGT-002)	1.56	1.26	1.52	2.46	1.22	1.33
Forensics Analyst (IN-FOR-001)	1.57	1.24	1.42	2.63	1.19	1.35
System Administrator (OM-ADM-001)	1.59	1.25	1.46	2.72	1.21	1.30
Systems Developer (SP-SYS-002)	1.65	1.33	1.57	2.60	1.26	1.47
System Test & Evaluation Specialist (SP-TST-001)	1.68	1.33	1.56	2.81	1.28	1.44
Systems Security Analyst (OM-ANA-001)	1.72	1.35	1.66	2.83	1.27	1.49
Cyber Crime Investigator (IN-INV-001)	1.75	1.36	1.62	3.02	1.29	1.48
Exploitation Analyst (AN-EXP-001)	1.76	1.34	1.67	3.04	1.28	1.46
Vulnerability Assessment Analyst (PR-VAM-001)	1.77	1.37	1.64	3.06	1.31	1.46
Research and Development Specialist (SP-TRD-001)	1.78	1.38	1.66	3.03	1.32	1.51
Security Architect (SP-ARC-002)	1.90	1.39	1.76	3.51	1.33	1.50
Knowledge Manager (OM-KMG-001)	1.90	1.37	1.72	3.66	1.35	1.41
Information Systems Security Developer (SP-SYS-001)	1.95	1.43	1.81	3.59	1.38	1.55
Systems Requirements Planner (SP-SRP-001)	1.96	1.45	1.74	3.69	1.43	1.49
Cyber Instructor (OV-TEA-002)	1.98	1.42	1.85	3.77	1.36	1.51
Enterprise Architect (SP-ARC-001)	2.01	1.46	1.84	3.80	1.40	1.57
Product Support Manager (OV-PMA-003)	2.03	1.48	1.83	3.83	1.43	1.58
Technical Support Specialist (OM-STS-001)	2.04	1.47	1.99	3.72	1.38	1.61
Security Control Assessor (SP-RSK-002)	2.04	1.45	1.90	3.89	1.38	1.57
Target Network Analyst (AN-TGT-002)	2.06	1.46	1.97	3.89	1.40	1.57
IT Program Auditor (OV-PMA-005)	2.06	1.50	1.85	3.90	1.47	1.57
Information Systems Security Manager (OV-MGT-001)	2.08	1.46	1.93	4.01	1.42	1.57
All Source-Collection Manager (CO-CLO-001)	2.08	1.46	1.89	4.07	1.41	1.57
All Source-Collection Requirements Manager (CO-CLO-002)	2.09	1.47	1.90	4.06	1.42	1.58

(continued)

Table 6 (continued)

Multi-Disciplined Language Analyst (AN-LNG-001)	2.10	1.45	1.92	4.14	1.39	1.58
Information Technology (IT) Project Manager (OV-PMA-002)	2.14	1.52	1.90	4.19	1.49	1.59
Program Manager (OV-PMA-001)	2.14	1.52	1.90	4.19	1.49	1.59
Cyber Workforce Developer and Manager (OV-SPP-001)	2.22	1.51	2.08	4.42	1.44	1.65
Target Developer (AN-TGT-001)	2.22	1.52	2.15	4.34	1.44	1.66
IT Investment/Portfolio Manager (OV-PMA-004)	2.32	1.58	2.02	4.82	1.56	1.63
Cyber Instructional Curriculum Developer (OV-TEA-001)	2.33	1.55	2.18	4.76	1.47	1.68
Authorizing Official (SP-RSK-001)	2.37	1.57	2.16	4.90	1.51	1.70
Threat/Warning Analyst (AN-TWA-001)	2.39	1.58	2.32	4.85	1.48	1.73
Mission Assessment Specialist (AN-ASA-002)	2.44	1.60	2.38	4.97	1.51	1.76
All-Source Analyst (AN-ASA-001)	2.47	1.60	2.40	5.05	1.51	1.77
Cyber Policy and Strategy Planner (OV-SPP-002)	2.64	1.65	2.36	5.84	1.61	1.72
Cyber Intel Planner (CO-OPL-001)	2.69	1.67	2.41	5.94	1.61	1.80
Privacy Officer/Privacy Compliance Manager (OV-LGA-002)	2.71	1.68	2.42	6.03	1.63	1.77
Cyber Ops Planner (CO-OPL-002)	2.77	1.70	2.53	6.13	1.62	1.84
Partner Integration Planner (CO-OPL-003)	2.90	1.73	2.64	6.62	1.65	1.86
Executive Cyber Leadership (OV-EXL-001)	3.06	1.78	3.03	6.82	1.69	1.96
Cyber Legal Advisor (OV-LGA-001)	3.16	1.84	2.93	7.25	1.75	2.02

- level three (3) on social interaction, where computers are capable of interpreting simple forms of communication (for example, produce a written synopsis of an oral conversation),
- level three (3) on physical work, where machines are able to maneuver in spaces where there are irregularities and obstacles (for example, warehouse robots that, in addition to the main task, are able to handle unforeseen obstacles),
- level three (3) on availability of statistical data, where structuring of existing data to be useful for machine learning can be performed easily within acceptable time frames (as, for example, when it comes to annotation of large amounts of images).

Table 9 describes the proportion of tasks within each category in the NICE framework that can be automated in 21 systematically chosen scenarios (out of the 625 theoretically possible). This reduced set was chosen to reflect a range of different cases that have plausible conditions. Scenario number three (3)—all 3s—is the one described above. Under these circumstances it can be seen that 61% of the skills and abilities can be automated on average, and a full 83% in the category *Operate and Maintain*.

The table also shows to what extent the four criteria affect the potential for automation:

Table 7 NICE framework specialty areas. The figures are indexed and should be read relative to other specialty areas. The easiest specialty area to automate (index 1.0) is *Data Administration*

Specialty area	Mean	Model				
		1	2	3	4	5
Data Administration (DTA)	1.00	1.00	1.00	1.00	1.00	1.00
Network Services (NET)	1.14	1.04	1.07	1.48	1.04	1.05
Cyber Operations (OPS)	1.20	1.11	1.14	1.51	1.10	1.12
Cyber Defense Infrastructure Support (INF)	1.26	1.13	1.19	1.74	1.15	1.11
Software Development (DEV)	1.28	1.16	1.22	1.68	1.14	1.20
Cyber Defense Analysis (CDA)	1.29	1.15	1.23	1.78	1.15	1.15
Incident Response (CIR)	1.32	1.19	1.25	1.80	1.18	1.20
Digital Forensics (FOR)	1.37	1.17	1.25	2.09	1.15	1.22
Systems Administration (ADM)	1.41	1.18	1.29	2.21	1.18	1.17
Test and Evaluation (TST)	1.49	1.26	1.38	2.28	1.24	1.30
Systems Analysis (ANA)	1.53	1.28	1.46	2.30	1.24	1.35
Cybersecurity Management (MGT)	1.55	1.27	1.48	2.47	1.26	1.28
Cyber Investigation (INV)	1.55	1.29	1.44	2.46	1.25	1.34
Exploitation Analysis (EXP)	1.56	1.27	1.48	2.48	1.24	1.32
Vulnerability Assessment and Management (VAM)	1.57	1.29	1.45	2.49	1.27	1.32
Technology R&D (TRD)	1.58	1.31	1.46	2.46	1.28	1.37
Systems Development (SYS)	1.63	1.32	1.52	2.62	1.29	1.38
Knowledge Management (KMG)	1.67	1.30	1.52	2.98	1.31	1.27
Systems Architecture (ARC)	1.71	1.34	1.58	2.94	1.32	1.38
Systems Requirements Planning (SRP)	1.73	1.38	1.54	3.00	1.38	1.34
Customer Service and Technical Support (STS)	1.80	1.40	1.76	3.03	1.34	1.46
Risk Management (RSK)	1.82	1.38	1.71	3.25	1.35	1.43
Collection Operations (CLO)	1.83	1.39	1.68	3.31	1.37	1.42
Training, Education, and Awareness (TEA)	1.84	1.38	1.73	3.33	1.35	1.42
Language Analysis (LNG)	1.84	1.38	1.70	3.37	1.35	1.43
Project Management/Acquisition and Program (PMA)	1.86	1.43	1.66	3.33	1.43	1.43
Targets (TGT)	1.88	1.41	1.81	3.33	1.38	1.46
Strategic Planning and Policy (SPP)	2.09	1.48	1.94	4.04	1.46	1.51
Threat Analysis (TWA)	2.10	1.49	2.05	3.95	1.44	1.56
All-Source Analysis (ASA)	2.15	1.51	2.11	4.07	1.46	1.59
Cyber Operational Planning (OPL)	2.41	1.60	2.21	5.01	1.57	1.65
Legal Advice and Advocacy (LGA)	2.41	1.61	2.19	5.02	1.59	1.63
Executive Cyber Leadership (EXL)	2.66	1.68	2.68	5.55	1.64	1.77

- The requirement for creativity appears to be of high importance. *All* roles can be very hard to automate if the requirement for creativity is not fulfilled (Scenario 18).

Table 8 NICE framework categories. The figures are indexed and should be read relative to other categories. The easiest category to automate (index 1.0) is the *Operate and Maintain* category

		Model				
Category	Mean	1	2	3	4	5
Operate and Maintain (OM)	1.00	1.00	1.00	1.00	1.00	1.00
Protect and Defend (PR)	1.07	1.06	1.06	1.12	1.07	1.06
Investigate (IN)	1.09	1.05	1.05	1.22	1.04	1.08
Securely Provision (SP)	1.30	1.18	1.29	1.62	1.17	1.21
Analyze (AN)	1.48	1.27	1.55	2.01	1.25	1.31
Oversee and Govern (OV)	1.49	1.28	1.53	2.10	1.27	1.29
Collect and Operate (CO)	1.56	1.31	1.57	2.30	1.30	1.33

- The requirement for physical work is more or less insignificant. The possibilities for automation remain good even if physical work cannot be carried out at all (Scenario 20).

The table, hence, shows that creativity is more important than the ability to perform physical work. This makes intuitive sense, and should come as no surprise. Further, it can be seen that the requirements for creativity and access to statistical data are the two most important criteria. The requirement for social interaction is the third most important criterion.

5 Discussion

The quantitative measures produced by the panel as well as the different models can be interpreted in many ways. Some noteworthy results that can serve as a starting point for a discussion are listed below:

- The results obtained by the various models differ. The hardest role to automate is 1.75 times more difficult than the easiest role in model four, while it is 7.25 times harder in model three (see Table 6).
- Even if there are differences between the models they yield robust values. The ranking of the proneness for automation for various roles, is more or less equal in all models. That is, there are no major differences in the results depending on whether the criteria interact or not: some tasks are always difficult to automate, while others are always easy.
- Technical or more practically oriented roles, such as database administrators, data analysts, and network operators, seem to be the ones that are easiest to automate. Roles that require formal responsibility, such as legal advisors and executive directors, are more difficult to automate.
- Tasks that deal with systems development and system administration are easier to automate than roles dealing with intelligence issues.

Table 9 Scenarios where the automation criteria have been met to varying degrees. The figures indicate the proportion of skills and abilities that can be automated in each scenario

	Technology and preconditions				Automability							
Scenario	Requirements for creativity	Requirements for social interaction	Requirements for physical work	Existing statistical data	Mean	Protect and Defend (PR)	Securely Provision (SP)	Operate and Maintain (OM)	Oversee and Govern (OV)	Investigate (IN)	Analyze (AN)	Collect and Operate (CO)
1	1	1	1	1	1%	2%	0%	2%	0%	0%	2%	1%
2	2	2	2	2	19%	29%	11%	39%	12%	24%	9%	8%
3	3	3	3	3	61%	82%	62%	83%	47%	67%	49%	37%
4	4	4	4	4	96%	100%	98%	97%	91%	100%	90%	93%
5	5	5	5	5	100%	100%	100%	100%	100%	100%	100%	100%
6	1	2	2	2	2%	2%	2%	4%	1%	4%	3%	1%
7	2	1	2	2	8%	11%	4%	21%	5%	12%	3%	3%
8	2	2	1	2	18%	29%	10%	36%	12%	20%	9%	7%
9	2	2	2	1	1%	2%	0%	3%	0%	0%	2%	2%
10	1	3	3	3	2%	2%	2%	4%	1%	4%	3%	1%
11	3	1	3	3	15%	16%	12%	32%	9%	22%	8%	5%
12	3	3	1	3	54%	77%	54%	72%	46%	47%	47%	35%
13	3	3	3	1	1%	2%	0%	3%	0%	0%	2%	2%
14	1	4	4	4	3%	2%	2%	4%	1%	8%	3%	1%
15	4	1	4	4	17%	16%	13%	36%	9%	33%	8%	5%
16	4	4	1	4	82%	95%	88%	82%	89%	59%	81%	84%
17	4	4	4	1	2%	2%	0%	3%	0%	4%	2%	2%
18	1	5	5	5	3%	2%	2%	4%	1%	8%	3%	1%
19	5	1	5	5	18%	16%	13%	36%	11%	33%	10%	5%
20	5	5	1	5	87%	95%	90%	84%	97%	59%	91%	92%
21	5	5	5	1	2%	2%	0%	3%	0%	4%	2%	2%

- Technological advances or other developments that can automate functions that require creativity, determine how much and how well the cybersecurity area can be automated. Developments of robotics that can meet requirements for physical work, is of comparably limited importance.
- The availability of statistical data is, in parallel with creativity, also decisive for whether, and how, automation within the cybersecurity area can be achieved.
- When the prerequisites for automation, according to our models, are at Levels 2–3, which it is reasonable to assume that they are today, 19–61% of the skills and abilities can in general be automated. Hence, many tasks still remain to automate.

In the remainder of this section we discuss how deficiencies in our methodology may have affected the validity of the results.

5.1 Limitations

The method used for this study can be summarized by the following steps where we:

1. obtained descriptions for various types of cybersecurity work,
2. determined sensible criteria for the purpose of assessing cybersecurity work,
3. carried out assessments by using the criteria developed in step 2, and
4. performed several analyses and syntheses of the results for different abstraction levels.

Every step has inherent weaknesses that affect the result. In the following an assortment of the potentially most serious ones are discussed.

5.1.1 Descriptions of Cybersecurity Work

To encompass the entire scope of cybersecurity work, the analysis was based on the NICE framework. NICE seeks to cover the whole range of tasks that can be related to the field. Here it can be noted that there probably exist discrepancies regarding how roles are described in NICE, and to what extent such roles are prevalent on, for example, the Swedish labor market. Furthermore, the roles in NICE are described with varying granularity. The role *Database Administrator* can probably be defined with a fairly narrow unambiguous process description, while a role such as *Cyber Policy and Strategy Planner* probably does not allow itself to be succinctly defined. But, even if roles were well-defined, which, for example, is the case for the *Software Development* specialty area in which [a developer] "[d]evelops and writes/codes new (or modifies existing) computer applications" [31, p. 12], they can still differ significantly in practice. In a Swedish study about software development it was found that the competence requirements for software developers

in various industrial sectors were extremely heterogeneous. There were wildly varying requirements regarding knowledge in areas such as artificial intelligence, embedded software, app programming, and e-commerce [5].

The roles in NICE also appear to be heavily influenced by tasks that are common in the U.S. (federal) defense and intelligence sector. This sector, however, probably only constitutes a minor part of the commercial sector as a whole. As a consequence, there is a risk that government-specific roles are given too much relative weight, given that all roles in NICE are "treated equally". In other words, an exaggeration of the importance of government-only jobs risks to influence the overall validity of the results negatively.

A final remark concerning NICE is that it contains a number of skills and abilities that become obsolete when tasks are automated. One example is skills related to the use of man-machine interfaces. There is no need for such skills when machines solely communicate with other machines.

5.1.2 The Assessments

As has already been brought forward, the variance of the judgments of the panel does not necessarily have to be seen as something negative. The variations can to some part be explained by their different backgrounds and knowledge levels. It was a deliberate goal of the research design to allow for variance and the use of mean values, as described, in order to produce well-balanced assessments. The panel members' scoring on skills and abilities relative to the four criteria was positively correlated, i.e., the assessments generally pointed in the same direction: if one assessor judged a value to be high, the others typically also thought so. The mean deviations between the scores of the assessors were moderate (on average 0.4 on a five-point scale). There were also systematic differences where, for example, the average scores for creativity differed between the assessors. For creativity the difference was almost a whole point (0.8) with mean values of 2.1, 2.4, 2.8, and 2.9 for the four assessors.

None of the factors mentioned this far are judged to be problematic. There were, however, also skills and abilities that showed discrepancies in judgments, that were probably not only due to the different backgrounds and knowledge levels of the assessors. During the seminar it was discovered that the actual meaning of certain skills and abilities were interpreted differently, because they were described in vague terms. One such example was the skill to perform intelligence collection on the so-called *dark web*. This skill was the one that showed the largest assessor score deviance of all skills and abilities. When this was discussed, it turned out that the assessors had made different assumptions about the context. This, in turn, led to that the skill was assessed to require anything between no social skills at all, to excellent skills. The former would be reasonable if the task was primarily about sifting through large amounts of text in search for specific information. The latter would be more sensible if social interactions were thought to be part of tricking informants to give up information. This example illustrate, again, that there can

be different interpretations of the NICE descriptions. We assert that it would be hard, not to mention *extremely* time-consuming, to completely unify the views of all assessors for all the descriptions of the 541 analyzed tasks. In light of these remarks, the mean values that we used can be seen as a good approximation.

5.1.3 The Aggregation

The aggregation, in our opinion, was made in a straightforward and transparent manner. However, it also involved some simplifications. Two major ones are that:

1. all skills and abilities had the same weight, even if some probably are more important for successfully carrying out work in a role, than others,
2. all roles and specialties also had the same weight, even though there is a larger number of people working with, e.g., *System Administration* than with, e.g., *Executive Cyber Leadership* in the labor market.

Consequently, the analysis would be improved if the *portions* of the requirements of skills and abilities were judged for each role. Further, the assignment of weights to the various roles and specialties would, likewise, be helpful for improving analyses of market potential and addressing questions involving automation of multiple roles. This remains as future work.

Another deficiency in the aggregation process was that it relied on the underlying assumption that all tasks within a role can be fully automated, and not that only certain parts of it can. Further, some specific skills and abilities are more determinate than others for whether a role can be successfully automated. A survey filled out by participants at the large cybersecurity conferences DEF CON and Black Hat, can serve to illustrate this point. Respondents judged how important certain skills and abilities were for the specialty *Vulnerability Assessment and Management* [20]. Here, respondents judged a skill "in conducting vulnerability scans and recognizing vulnerabilities in security systems" to be the most important, which would be easy to automate. At the same time they judged that a skill "in the use of social engineering techniques" was less important. This skill would be significantly harder to automate.

The five models must be seen as gross simplifications that serve the purpose of representing recurring patterns, and not as a rigid method to exactly predict to what extent cybersecurity work can be automated. It is a bit surprising that the models, despite their fundamentally different constructions (e.g., linear vs. exponential, and minima vs. maxima), yield such a consistent ranking (as indicated by the heat map coloring in Tables 6, 7, and 8). One reason for this is that the values of the four criteria are correlated, i.e., high demands for creativity seldom come without equally high demands for another criterion, like perhaps social interaction. An advantage here is that the models display small differences, and that any model can be used. This, in turn, leads to the results of the study being, again, more robust, less susceptible to assumptions, and easier to interpret.

5.2 Other Important Variables

The results discussed this far have been about possibilities for automation as of *today*. This study can therefore provide some guidance about what role to automate first if there are several options to consider. There are, however, many other variables that also affect decisions about preferred automation solutions (as has been briefly touched upon in Sect. 2.1). In the following we discuss some of these.

5.2.1 Market Potential

It is reasonable to assume that the roles that are the easiest and most profitable to automate will be automated first [6]. Profitability can be reached in different ways, though. One could either automate jobs that are relatively simple, but occupy large groups of people, or more complicated jobs that occupy fewer, but better paid personnel. To roughly estimate the market potential for automation of various IT jobs, data from Statistics Sweden, the Swedish governmental agency for official statistics, has been used. Table 10 shows the grand sums of all monthly salaries that are paid out in various IT occupations.[3] As an example, it can be seen that the savings in salaries would be eight (8) times greater if the jobs of *Software- and system developers* were to be automated instead of *ICT operations technicians*, on the Swedish labor market. We did not, however, investigate if the allotment of portions of cybersecurity professions on the labor market is similar in other countries.

We used the data in Table 10 due to that we do not have access to statistics related to specific NICE roles, which means that the jobs displayed in the table do not fully relate to the jobs in the NICE framework. Comparisons are therefore hard to make. The website CyberSeek,[4] however, indicates the proportions of commonality of various IT jobs in terms of the NICE categories in the United States. There are great differences. In the beginning of April 2019 there were 207,190 jobs listed in the *Operate and Maintain* category, while a mere 49,825 were listed in *Collect and Operate*. To summarize, data from both Sweden and the United States suggest that the market potential for the various specialties in the NICE framework differ. The question of automation should therefore also be seen in light of this fact. Besides the number of specific jobs on the market—the *commonality* of jobs—as discussed, *competence requirements* also affect salaries [19]. Jobs with high demands yield high salaries and vice versa, and as mentioned jobs with high responsibility requirements appear to be difficult to automate.

[3] Compiled based on Statistics Sweden's table entitled "Employees 16–64 years at national level by occupation (4-digit SSYK 2012), sector, and sex. Year 2014–2017" in the Swedish occupational register, http://www.statistikdatabasen.scb.se/, and "Salary search", https://www.scb.se/en/finding-statistics/sverige-i-siffror/salary-search/.

[4] https://www.cyberseek.org/.

Table 10 IT jobs and the total amount of monthly salary paid out, based on average salaries and the number of people occupied within each job

Code	Name	Total salary (Swedish crowns, SEK)
2512	Software- and system developers	3,208,320,900
2511	System analysts and ICT-architects	666,564,200
3512	ICT support technicians	570,362,400
3514	Computer network and systems technicians	476,825,800
3511	ICT operations technicians	403,930,800
2614	Business and company lawyers	344,061,300
2515	System administrators	279,847,900
2514	System testers and test managers	249,895,000
3513	System administrators	154,556,200
2516	Security specialists (ICT)	93,998,000
3515	Webmasters and web administrators	77,634,000

5.2.2 Intent and Ability

Up til this point we have been reasoning about the feasibility of automation and market potential, but a few additional prerequisites also have to be fulfilled before automated solutions can be launched. There has to be a *will* to automate.

Already in 1957, Cowan noted that human resistance could be an obstacle to successful automation efforts [12]. Factors such as habit and fear were seen as prohibitive, including aspects related to, for example, personal complacency and the fear of losing a well paid or high status position. As a solution to turn people who felt threatened by automation, Cowan suggested either to seek their active involvement in the automation efforts, or extensive training. This would increase the chances for them to accept or even embrace automation [12]. These suggestions ring equally true today.

First, it seems to be hard to more precisely quantify the level of resistance towards machines by individuals in the cybersecurity field. Despite the quantification difficulties, this problem appears to be an important one for organizations to handle. We have previously put forth (see Sect. 2.1) that both usefulness and usability are important factors for prospective users when new technology is introduced [27]. All automation is probably not met with equal resistance, though; software testing is a task that may be perceived as less stimulating than, for example, systems architecture work, and may therefore encounter less resistance. A type of stronger resistance, where whole groups rather than individuals are united in "anti-automation" alliances, can also be imagined; on an organizational level, it is not too far-fetched that whole IT departments might fight efforts to outsource their tasks by, for example, buying cloud computing capabilities from an external supplier.

Second, both individuals and organizations have to be ready to not only accept automation changes, but also to *embrace* them. De Zan [15] argues that organizational changes generally follow when IT investments have been made, and new technology is introduced. Such changes can consist of restructured workflows

and new ways of sharing data. Data can then be used for machine learning. A high degree of automation probably leads to a *more* data-intensive work environment. The increased amount of data could then lead to difficulties with the existing organizational structures and IT architecture. That is, it is no trivial task to actually make use of all (new) data to bring forth more efficient decisions [6].

5.2.3 Ethical and Legal Issues

Given the content of this chapter, one realizes that many tasks can be automated from a purely technical standpoint. Some of those tasks, however, will not be automated due to ethical or legal considerations. Self-driving cars is an example where the question of legal liability after an accident poses a great obstacle for automation: it is not permitted to abdicate the responsibility to a machine [23]. Other ethical and legal issues may also arise in cybersecurity, although the majority of cases are likely to be unproblematic. To have automated systems scan employees' e-mails for malware is probably not controversial. Such a solution, where no humans view personal information, may be preferred over a manual process. An investigation after abusive use of a company's IT resources, on the other hand, can probably not be fully automated. Decisions about filing formal criminal charges or firing an individual, cannot be delegated to a machine. For such cases there need to be humans in the loop. In line with this reasoning, it seems reasonable to believe that tasks in the NICE category *Investigate* are more suitable for automation than tasks that involve, e.g., executive leadership or judgmental components. It can also be speculated that there are probably considerably less ethical problems involved in tasks such as construction of new systems (e.g., *Systems Development*), than for tasks that involve collection of sensitive data related to systems already in use (e.g., *Collection Operations*).

Another example is the use of facial recognition technology. Such use would be perfectly feasible and serve a greater good by its ability to identify and catch criminals. There are, however, examples where use has been prohibited. The city of San Francisco chose to ban facial recognition in 2019 due to concerns regarding potential privacy issues and the possibility for abusive use.[5] Yet other aspects to keep in mind, are the introduction of potential biases in AI systems, and the opacity of the employed algorithms. Bias can be either deliberately or accidentally built into the algorithm itself, or arise from already biased training data [52]. Lack of algorithmic transparency can result in both ethical and practical problems. Research with the aim to interpret and explain *black box* decision support systems is ongoing [24].

It can be concluded that both laws and ethics are factors to be considered with regard to whether automation is allowed or not, but also with regard to whether it should be strived for or not. However, laws and ethics also affect the very *conditions* for creating automated solutions. The availability of data for research,

[5] https://www.nytimes.com/2019/05/14/us/facial-recognition-ban-san-francisco.html.

systems development, and alignment of automated solutions, can be regulated by laws. Mandatory reporting of cyber-related incidents, as an example, would greatly enhance the knowledge about cyberattacks, which could be used to extract knowledge for the benefit of cybersecurity efforts, not least for the handling of threat intelligence [45]. More restrictive laws related to personal integrity, on the other hand, would make the collection of large datasets for machine learning significantly harder.

5.3 Will Automation Improve Cybersecurity?

A highly relevant question to ponder is whether extensive AI automation of cybersecurity work actually improves the level of cybersecurity. In this section we speculate about possible outcomes. Coombs et al. [11] found no studies that evaluated changes in overall organizational performance even after extensive automation. We earlier brought up that it is more likely that specific tasks within a role are automated than all aspects of that role. Zhong et al. [53], for example, argue that cybersecurity work in security operations centers cannot be expected to be fully automated anytime in the near future. They stress that there are in principle two main paths to achieve increased efficiency in such centers. The first would be to identify areas that can be automated, and strive for the automation of those. The second would be to ensure viable coexistence conditions for humans and machines, so that maximum gains from the automated parts can be extracted. It is important to achieve synergies in man-machine teams and draw on the strengths of both.

There are factors that point towards improved cybersecurity due to automation. Automation will cut the time consumption for several tasks, and increased speed will be important against nonautomated adversaries. Another factor is that automation can help reduce the required amount of information that eventually needs to be considered by human analysts. A major (cognitive) problem for cybersecurity analysts is the need to process overwhelming amounts of information [13], e.g., to handle information overload [18]. Much time is spent on removing false positives from alert systems. At the same time it is important for human operators to maintain a high degree of attention and ability to focus over extended periods of time. Such "cyber vigilance" [35] has been shown to affect the results of maintaining good cybersecurity. Automation will, thus, help analysts to maintain their vigilance longer by removing tedious work and by reducing their cognitive load.

Hitherto we have discussed automation for cyber *defense*, which this chapter is about. It should be noted, though, that the potential for automation is equally present for other areas within the cyber domain. Technologies such as, e.g., automatic vulnerability scans [49], and machine learning approaches for improving intrusion detection systems [8], can just as well be used for offensive purposes. In fact, there are fears that ever-increasing use of AI for cybersecurity will lead to a "cyber arms race" that even risks to escalate into conventional war with actual physical attacks [44].

It is next to impossible to foresee the actual gains of future automation of cybersecurity functions. The main factor that makes such predictions highly uncertain is that the underlying automation technology is available for potential adversaries as well. Potential improvements on the cybersecurity side therefore risk being levelled out by equal improvements on the adversarial side. But if this balance is tilted, and highly automated cybersecurity functions are up against nonautomated adversaries, chances are that the level of cybersecurity will be improved compared to today. The nature of cybersecurity work for humans may also change. If many tedious and time-consuming tasks are automated, the capabilities of human experts can be "saved" for particularly hard tasks.

5.4 Effects on the Labor Market

In this section we discuss possible effects on the (cybersecurity) labor market in an imaginary future scenario after most cybersecurity tasks have been automated.

It is evident that automation efforts this far have not resulted in a diminished need for IT work [15]. In 2014, it was reported that programmer was the most common profession in the (Swedish capital) Stockholm region [42], and that the need for programmers on the labor market remained at a high level. In 2019, *Software- and system developers* were ranked as the eighth largest profession in Sweden, according to Statistics Sweden.[6] In fact, the IT business as a whole continues to grow, and cybersecurity personnel are also in high demand [42]. Further, on the international level, multiple reports point to severe shortages of cybersecurity personnel. In addition, it has been reported that the shortage cannot be remedied in the near future, because the proposed solutions take time to implement [15].

Automation will increasingly affect work that is carried out in the cybersecurity field. The OECD [32] points out that personnel with the lowest level of training or salaries, are the ones that first risk being replaced by automated solutions. Here it is important to remember that whole professions are not necessarily going to be automated in their entirety [15]. What will remain are subtasks that computers (machines) are worse at solving than humans. These include unusual unscripted tasks, and tasks that require social interaction and problem-solving that cannot easily be described by algorithms.

In general, Brynjolfsson and McAfee [7] argue that technological advances have meant that the gap between the most competent part of the workforce, and those who perform on, or below, average, has increased. The most competent personnel tend to be more sought-after, indispensable, and therefore better paid, at the expense of the group that we here call the average performers. An indication of this phenomenon is that military commanders in the U.S. have been reported to claim that the most productive IT specialists produce up to a factor hundred more worth of effect, than

[6] https://www.scb.se/.

the least productive ones [7]. With this increased focus on specialized competencies and high productivity, it is therefore the average performers who risk being made redundant on the labor market. They will have to accept lowered salaries, or even permanent unemployment. In this chapter we have not dug into any details of this reasoning, but note that the same kind of reasoning probably applies to the cybersecurity labor market as well.

6 Conclusions

In this section we answer the three research questions that were outlined in Sect. 1. The reader should note that the answers represent the outcome of the analysis given the analytical process. That is to say that there are uncertainties regarding the universal validity of the conclusions. Moreover, the results apply for the wider cybersecurity field, that is, it is obviously possible to obtain better and more precise results for research questions targeting more isolated questions. For example, in the study presented by Sommer and Paxson [37] discussed in Sect. 2.2, the development of intrusion detection systems directly relates to automation of work performed specifically by the NICE framework role *Cyber Defense Incident Responder*.

6.1 What Variables Affect How Hard a Cybersecurity Role Is to Automate?

From the point of envisioning models for automation of generic work tasks, three variables of interest were identified. In the beginning of our analysis, however, four variables emerged: requirements for creativity, social interaction, physical work, and the availability of statistical data related to the role in question. These variables in themselves represent gross simplifications, as they can be described in many equally sensible ways. Creativity, as an example, can be seen both as having a capability to invent novel solutions (originality), and to identify viable (but already existing) solutions to problems. The latter seems more common. It is therefore a problem to isolate and succinctly define variables that affect the possibilities for automation. The exception is the requirement for physical work, which clearly involves physical movement. Out of the original four variables, physical work was determined to play an insignificant role (see Sect. 4.1), and was removed. The answer to the question is, thus:

- requirements for creativity,
- requirements for social interaction, and
- the availability of statistical data.

The importance of the variables vary with roles (see Table 4). There are also differences in the types of requirements for different roles. Tasks carried out by a *Program Manager* and an *IT Program Auditor* have requirements for creativity and for social interaction on roughly the same level, whilst a *Systems Developer* and a *Systems Security Analyst* are performing tasks that have significantly higher requirements for creativity than for social interaction. It could be noted that the requirements covary to a large extent, i.e., if there is a requirement for a high value for one variable (like creativity), there is often a similarly high requirement for another variable (like the availability of statistical data):

- requirements for social interaction and requirements for creativity for various skills and abilities have a correlation coefficient of 0.38,
- requirements for social interaction and availability of statistical data for various skills and abilities have a correlation coefficient of 0.58,
- requirements for creativity and availability of statistical data for various skills and abilities have a correlation coefficient of 0.73.

The covariation demonstrated above is not particularly unexpected. For example, social skills often require creative ways of communicating, and it is natural that tasks without documented historical examples/data to rely on, require creative solutions. This means, however, that the answer to the question becomes less clear, i.e., the fact that the variables covariate indicates that there are underlying variables that play a role, and that the criteria used for the study presented herein are in this sense not an optimal (orthogonal) breakdown of the variables that affect automation. What these underlying variables might be, and which criteria might be more suitable, is left for future research.

6.2 How Likely Is It That Current Cybersecurity Roles Will Be Automated?

The tables in Sect. 4 show significant differences between how hard it seems to be to automate various roles. They show that it is much easier to automate tasks that database administrators, data analysts, and network specialists perform, than those performed by intelligence analysts and senior executives. There are indicators that support this conclusion. For example, there are studies that show that the budgets allotted for database administrators decreased between 2013 and 2017.[7] However, as already mentioned in Sect. 5, it is unlikely that whole professions will be automated simultaneously, and the cybersecurity market as a whole does not necessarily seem to be moving in a direction where a majority of its personnel will be deemed redundant or need retraining. On the contrary, the U.S. Bureau of Labor Statistics,

[7] Computer Economics, Inc., "Database Administrator Ranks Show Steady Decline", https://www.computereconomics.com/article.cfm?id=2439.

for example, suggests that database administrators are heading towards a bright future, where the rate of increase is larger than that of the average profession.[8] This, in turn, might in part be explained by the increase of the whole IT sector, making it possible that database administrators face a bright future even if their relative percentage of the IT budget decreases.

6.3 *What Variables Constrain the Potential for Automation of Today's Cybersecurity Roles?*

As has been described in Sect. 6.1, the variables covary. This makes it hard to conclusively answer the question about the most constraining variables. However, some indications are that:

- the low demand for physical work, tells us that this requirement is not a significant obstacle (see, for example, Scenarios 3 and 12 in Table 9),
- the requirement for social interaction is a minor obstacle than the requirement for creativity and the lack of statistical data (see, for example, Scenarios 6, 7, and 9 in Table 9),
- the requirement for creativity and the need for existing statistical data covaries strongly (with a correlation coefficient of 0.73).

Based on the above remarks, a conclusion is that an ability to produce machines that suggest creative solutions to hard problems, would enable many roles to be automated. Another possibility to enhance the chances for automation would be to focus on collecting data to be used for machine learning. In practice, these suggestions are not mutually exclusive, though: recent creative advances in AI hinge upon the availability of historical data with correct solutions (labeled datasets), or the possibility to create relevant data by, for example, letting machines compete against each other. The lack of data, hence, seems to be a significant obstacle for further automation within the cybersecurity field.

Acknowledgments This work was originally commissioned to the Swedish Defence Research Agency, FOI, by the Swedish Civil Contingencies Agency, MSB. Some initial results were reported on in a short memo [38]. Stefan Varga was sponsored by the Swedish Armed Forces. The authors would like to acknowledge and thank Erik Zouave for his part in the assessment work.

[8] U.S. Department of Labor Bureau of Labor Statistics, Occupational Outlook Handbook, "Database Administrators and Architects", https://www.bls.gov/ooh/computer-and-information-technology/database-administrators.htm.

References

1. Arntz, M., Gregory, T., Zierahn, U.: Revisiting the risk of automation. Econ. Lett. **159**, 157–160 (2017). https://doi.org/10.1016/j.econlet.2017.07.001
2. Autor, D.H., Levy, F., Murnane, R.J.: The skill content of recent technological change: an empirical exploration. Q. J. Econ. **118**(4), 1279–1333 (2003). https://doi.org/10.1162/003355303322552801
3. Avgerinos, T., Cha, S.K., Rebert, A., Schwartz, E.J., Woo, M., Brumley, D.: Automatic exploit generation. Commun. ACM **57**(2), 74–84 (2014). https://doi.org/10.1145/2560217.2560219
4. Barr, E.T., Harman, M., McMinn, P., Shahbaz, M., Yoo, S.: The oracle problem in software testing: a survey. IEEE Trans. Softw. Eng. **41**(5), 507–525 (2015). https://doi.org/10.1109/TSE.2014.2372785
5. Borg, M., Wernberg, J., Olsson, T., Franke, U., Andersson, M.: Illuminating a blind spot in digitalization: software development in Sweden's private and public sector. In: Proceedings of the IEEE/ACM 42nd International Conference on Software Engineering Workshops (ICSEW 2020), pp. 299–302. ACM, New York (2020). https://doi.org/10.1145/3387940.3392213
6. Bresnahan, T.F., Brynjolfsson, E., Hitt, L.M.: Information technology, workplace organization, and the demand for skilled labor: firm-level evidence. Q. J. Econ. **117**(1), 339–376 (2002). https://doi.org/10.1162/003355302753399526
7. Brynjolfsson, E., McAfee, A.: The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies. W. W. Norton & Company, New York (2014)
8. Buczak, A.L., Guven, E.: A survey of data mining and machine learning methods for cyber security intrusion detection. IEEE Commun. Surv. Tutorials **18**(2), 1153–1176 (2016). https://doi.org/10.1109/COMST.2015.2494502
9. Chen, T.R., Shore, D.B., Zaccaro, S.J., Dalal, R.S., Tetrick, L.E., Gorab, A.K.: An organizational psychology perspective to examining computer security incident response teams. IEEE Secur. Privacy **12**(5), 61–67 (2014). https://doi.org/10.1109/MSP.2014.85
10. Connell, J.: Kitanai, kitsui and kiken: the rise of labour migration to Japan. ERRRU Working Paper No 13, Economic and Regional Restructuring Research Unit, University of Sydney, Sydney (1993)
11. Coombs, C., Hislop, D., Taneva, S.K., Barnard, S.: The strategic impacts of intelligent automation for knowledge and service work: an interdisciplinary review. J. Strategic Inform. Syst. **29**(4), 1–30 (2020). https://doi.org/10.1016/j.jsis.2020.101600
12. Cowan, G.R.: The human side of automation. Electr. Eng. **76**(9), 768–771 (1957). https://doi.org/10.1109/EE.1957.6442704
13. D'Amico, A., Whitley, K., Tesone, D., O'Brien, B., Roth, E.: Achieving cyber defense situational awareness: a cognitive task analysis of information assurance analysts. Proc. Hum. Fact. Ergon. Soc. Annu. Meeting **49**(3), 229–233 (2005). https://doi.org/10.1177/154193120504900304
14. Dasgupta, P., Collins, J.B.: A survey of game theoretic approaches for adversarial machine learning in cybersecurity tasks. AI Mag. **40**(2), 31–43 (2019). https://doi.org/10.1609/aimag.v40i2.2847
15. De Zan, T.: Mind the gap: the cyber security skills shortage and public policy interventions. Tech. rep., Global Cyber Security Center, Rome (2019)
16. Deloitte: 2018 Luxembourg cyber security technology adoption survey (2018)
17. D'Silva, V., Kroening, D., Weissenbacher, G.: A survey of automated techniques for formal software verification. IEEE Trans. Comput.-Aided Design Integr. Circuits Syst. **27**(7), 1165–1178 (2008). https://doi.org/10.1109/TCAD.2008.923410
18. Eppler, M.J., Mengis, J.: The concept of information overload: a review of literature from organization science, accounting, marketing, MIS, and related disciplines. Inform. Soc. **20**(5), 325–344 (2004). https://doi.org/10.1080/01972240490507974

19. Erbacher, R.F., Frincke, D.A., Wong, P.C., Moody, S., Fink, G.: A multi-phase network situational awareness cognitive task analysis. Inform. Visualization **9**(3), 204–219 (2010). https://doi.org/10.1057/ivs.2010.5
20. Ferguson-Walter, K.J., Shade, T.B., Rogers, A.V., Niedbala, E.M., Trumbo, M.C., Nauer, K., Divis, K.M., Jones, A.P., Combs, A., Abbott, R.G.: The Tularosa study: an experimental design and implementation to quantify the effectiveness of cyber deception. In: Proceedings of the 52nd Annual Hawaii International Conference on System Sciences (HICSS 2019), pp. 7272–7281 (2019). https://doi.org/10.24251/HICSS.2019.874
21. Frey, C.B., Osborne, M.A.: The future of employment: how susceptible are jobs to computerisation? Technol. Forecasting Soc. Change **114**, 254–280 (2017). https://doi.org/10.1016/j.techfore.2016.08.019
22. Goodall, J.R., Lutters, W.G., Komlodi, A.: I know my network: collaboration and expertise in intrusion detection. In: Proceedings of the 2004 ACM Conference on Computer Supported Cooperative Work (CSCW 2004), pp. 342–345. ACM, New York (2004). https://doi.org/10.1145/1031607.1031663
23. Greenblatt, N.A.: Self-driving cars and the law. IEEE Spectr. **53**(2), 46–51 (2016). https://doi.org/10.1109/MSPEC.2016.7419800
24. Guidotti, R., Monreale, A., Ruggieri, S., Turini, F., Giannotti, F., Pedreschi, D.: A survey of methods for explaining black box models. ACM Comput. Surv. **51**(5), 1–42 (2018). https://doi.org/10.1145/3236009
25. Guo, M., Wang, G., Hata, H., Babar, M.A.: Revenue maximizing markets for zero-day exploits. Auton. Agents Multi-Agent Syst. **35**(2), 1–29 (2021). https://doi.org/10.1007/s10458-021-09522-w
26. Harreld, J.B.: Foreword: automation is at the center of human progress. In: Nof, S.Y. (ed.) Springer Handbook of Automation, Springer Handbooks, pp. XI–XII. Springer, Berlin/Heidelberg (2009). https://doi.org/10.1007/978-3-540-78831-7
27. King, W.R., He, J.: A meta-analysis of the technology acceptance model. Inform. Manag. **43**(6), 740–755 (2006). https://doi.org/10.1016/j.im.2006.05.003
28. Kordy, B., Piètre-Cambacédès, L., Schweitzer, P.: DAG-based attack and defense modeling: don't miss the forest for the attack trees. Comput. Sci. Rev. **13–14**, 1–38 (2014). https://doi.org/10.1016/j.cosrev.2014.07.001
29. Liu, B., Shi, L., Cai, Z., Li, M.: Software vulnerability discovery techniques: a survey. In: Proceedings of the 2012 Fourth International Conference on Multimedia Information Networking and Security (MINES 2012), pp. 152–156. IEEE, Piscataway (2012). https://doi.org/10.1109/MINES.2012.202
30. Manyika, J., Chui, M., Miremadi, M., Bughin, J., George, K., Willmott, P., Dewhurst, M.: A future that works: automation, employment, and productivity. Report, McKinsey Global Institute, San Francisco (2017)
31. Newhouse, W., Keith, S., Scribner, B., Witte, G.: National initiative for cybersecurity education (NICE) cybersecurity workforce framework. NIST Special Publication 800-181, National Institute of Standards and Technology, U.S. Department of Commerce (2017). https://doi.org/10.6028/NIST.SP.800-181
32. OECD: Artificial Intelligence in Society. OECD Publishing, Paris (2019). https://doi.org/10.1787/eedfee77-en
33. Parasuraman, R., Riley, V.: Humans and automation: use, misuse, disuse, abuse. Hum. Fact. **39**(2), 230–253 (1997). https://doi.org/10.1518/001872097778543886
34. Patton, R.D., Patton, P.C.: What can be automated? What cannot be automated? In: Nof, S.Y. (ed.) Springer Handbook of Automation, Springer Handbooks, chap. 18, pp. 305–313. Springer, Berlin/Heidelberg (2009). https://doi.org/10.1007/978-3-540-78831-7_18
35. Sawyer, B.D., Finomore, V.S., Funke, G.J., Mancuso, V.F., Funke, M.E., Matthews, G., Warm, J.S.: Cyber vigilance: effects of signal probability and event rate. Proc. Hum. Fact. Ergon. Soc. Annu. Meeting **58**(1), 1771–1775 (2014). https://doi.org/10.1177/1541931214581369

36. Scandariato, R., Walden, J., Joosen, W.: Static analysis versus penetration testing: a controlled experiment. In: 2013 IEEE 24th International Symposium on Software Reliability Engineering (ISSRE 2013), pp. 451–460. IEEE, Piscataway (2013). https://doi.org/10.1109/ISSRE.2013.6698898
37. Sommer, R., Paxson, V.: Outside the closed world: on using machine learning for network intrusion detection. In: Proceedings of the 2010 IEEE Symposium on Security and Privacy, pp. 305–316. IEEE, Piscataway (2010). https://doi.org/10.1109/SP.2010.25
38. Sommestad, T., Brynielsson, J., Varga, S.: Möjligheter för automation av roller inom cybersäkerhetsområdet [Opportunities for automation of cybersecurity roles]. FOI Memo 6737, Swedish Defence Research Agency, Stockholm (2019)
39. Sommestad, T., Franke, U.: A test of intrusion alert filtering based on network information. Secur. Commun. Netw. **8**(13), 2291–2301 (2015). https://doi.org/10.1002/sec.1173
40. Sommestad, T., Hunstad, A.: Intrusion detection and the role of the system administrator. Inform. Manag. Comput. Secur. **21**(1), 30–40 (2013). https://doi.org/10.1108/09685221311314400
41. Sommestad, T., Sandström, F.: An empirical test of the accuracy of an attack graph analysis tool. Inform. Comput. Secur. **23**(5), 516–531 (2015). https://doi.org/10.1108/ICS-06-2014-0036
42. Stockholm Chamber of Commerce: Programmerare: vanligaste yrket i Stockholmsregionen [Programmer: the most common profession in the Stockholm region]. Analys 2014:3, Stockholms Handelskammare, Stockholm (2014)
43. Suta, C., Barbieri, L., May-Gillings, M.: Future employment and automation. In: Hogarth, T. (ed.) Economy, Employment and Skills: European, Regional and Global Perspectives in an Age of Uncertainty, pp. 17–43. Fondazione Giacomo Brodolini, Rome (2018)
44. Taddeo, M., Floridi, L.: Regulate artificial intelligence to avert cyber arms race. Nature **556**(7701), 296–298 (2018). https://doi.org/10.1038/d41586-018-04602-6
45. Tounsi, W., Rais, H.: A survey on technical threat intelligence in the age of sophisticated cyber attacks. Comput. Secur. **72**, 212–233 (2018). https://doi.org/10.1016/j.cose.2017.09.001
46. van de Weijer, S.G.A., Leukfeldt, R., Bernasco, W.: Determinants of reporting cybercrime: a comparison between identity theft, consumer fraud, and hacking. Eur. J. Criminol. **16**(4), 486–508 (2019). https://doi.org/10.1177/1477370818773610
47. van der Zande, J., Teigland, K., Siri, S., Teigland, R.: The substitution of labor: from technological feasibility to other factors influencing the potential of job automation. In: Larsson, A., Teigland, R. (eds.) The Digital Transformation of Labor: Automation, the Gig Economy and Welfare, Routledge Studies in Labour Economics, chap. 3, pp. 31–73. Routledge, London (2019)
48. Verendel, V.: Quantified security is a weak hypothesis: a critical survey of results and assumptions. In: Proceedings of the 2009 Workshop on New Security Paradigms Workshop (NSPW 2009), pp. 37–49. ACM, New York (2009). https://doi.org/10.1145/1719030.1719036
49. Wen, T., Zhang, Y., Wu, Q., Yang, G.: ASVC: an automatic security vulnerability categorization framework based on novel features of vulnerability data. J. Commun. **10**(2), 107–116 (2015). https://doi.org/10.12720/jcm.10.2.107-116
50. Werlinger, R., Hawkey, K., Beznosov, K.: An integrated view of human, organizational, and technological challenges of IT security management. Inform. Manag. Comput. Secur. **17**(1), 4–19 (2009). https://doi.org/10.1108/09685220910944722
51. Werlinger, R., Muldner, K., Hawkey, K., Beznosov, K.: Towards understanding diagnostic work during the detection and investigation of security incidents. In: Proceedings of the Third International Symposium on Human Aspects of Information Security & Assurance (HAISA 2009), pp. 119–132 (2009)
52. Wilner, A.S.: Cybersecurity and its discontents: artificial intelligence, the internet of things, and digital misinformation. Int. J. **73**(2), 308–316 (2018). https://doi.org/10.1177/0020702018782496

53. Zhong, C., Yen, J., Liu, P.: Can cyber operations be made autonomous? An answer from the situational awareness viewpoint. In: Jajodia, S., Cybenko, G., Subrahmanian, V.S., Swarup, V., Wang, C., Wellman, M. (eds.) Adaptive Autonomous Secure Cyber Systems, pp. 63–88. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-33432-1_4