SoK: Applications and Challenges of using Recommender Systems in Cybersecurity Incident Handling and Response

Martin Husák
Institute of Computer Science
Masaryk University
Brno, Czech Republic
husakm@ics.muni.cz

Milan Cermak
Institute of Computer Science
Masaryk University
Brno, Czech Republic
cermak@ics.muni.cz

ABSTRACT

Incident handling, a fundamental activity of a cybersecurity incident response team, is a complex discipline that consumes a significant amount of personnel's time and costs. There are continuous efforts to facilitate incident handling and response in terms of providing procedural or decision support and processing relevant data. In this paper, we survey the approaches towards (semi-)automated incident handling and response backed by recommender systems that are successful in other domains. We discuss which phases and tiers of incident handling can be automated and to what level while evaluating the maturity of proposed approaches and tools. While we did not find a full-scale recommender system that would guide the user through incident handling and suggest which steps to take, many of them aim at particular problems. The discussed issues are not resolved yet but seem to get the attention of researchers and will likely be investigated in the future.

CCS CONCEPTS

• **Security and privacy → Network security**; **Vulnerability management**; • **Applied computing → Operations research**.

KEYWORDS

Recommender System, Incident Handling, Incident Response

ACM Reference Format:
Martin Husák and Milan Cermak. 2022. SoK: Applications and Challenges of using Recommender Systems in Cybersecurity Incident Handling and Response. In *The 17th International Conference on Availability, Reliability and Security (ARES 2022), August 23–26, 2022, Vienna, Austria.* ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3538969.3538981

1 INTRODUCTION

Incident handling is a crucial service of a Computer Security Incident Response Team (CSIRT), Computer Emergency Response Team (CERT), or Security Operations Center (SOC) [33, 57]. Incident handling covers various tasks beginning with incident triage (initial assessment of severity and scope of the incident), incident analysis, which often includes computer or network forensics, incident response, which often requires interaction with end-users or IT administrators, and gathering lessons learned [13, 33, 38]. The particular tasks of incident handling can further be distributed among several cybersecurity team members with various expertise and skill levels [57]. The high-level procedures are well described in the literature, but the execution is highly dependent on the local environment, available data and tools, and even factors like the organization's capabilities and competence. Although incident handling has been a common practice for over two decades, it is still an interesting research topic with many open challenges [22, 50, 54].

The paramount goal of current research and development efforts with respect to incident handling is the automation of the processes and offloading as much effort from human to computer as possible [3, 57]. Although there are tasks that are hard or impossible to automate and decisions that require a human in the loop, the motivation for automation is manifesting in several directions and growing in importance. First, we are facing a continuous increase in the number and complexity of cyber attacks and disclosed vulnerabilities that are affecting more and more areas of daily lives and human activities. Simultaneously, we face a shortage of skilled personnel. Estimates talk about four million unassigned work positions, while many such positions require solid skills [26]. Automating a significant amount of the incident handling tasks would reduce the burden imposed on current personnel and allow novices to pursue a career in cybersecurity. Incident handling playbooks [30, 43] allow junior incident handlers to resolve incidents faster and more precisely, while automation of laborious tasks allows senior personnel to focus on decision-making and analysis of advanced threats. Less staff would be required to conduct fundamental operations while more capacities could be dedicated to combating advanced threats.

In this paper, we survey the current state of the art in incident handling and response automation. First, we elaborate on procedures and tiers of incident handling and discuss which tasks can be automated and to what extent. Subsequently, we introduce recommender systems [1] in order to enumerate the approaches available to support the tasks of incident handling. It is worth noting that in this paper, we focus specifically on the recommender systems, not the whole landscape of decision support tools for cybersecurity operations. We encountered a plethora of works that propose decision support [22, 25], automation [3], or other facilitation of incident handling and response [39]. In essence, any visualization, data fusion or correlation, calculation, or situation projection could be considered a decision support tool. Although such tools are valuable and helpful, we consider them out of the scope of this work as they only provide additional information or present them in a comprehensive manner but do not implement a recommender system. We focus on approaches that go a step ahead and propose an action to take, pre-select an option, guide the incident handler through the processes, or otherwise actively help the users via the means

ARES 2022, August 23–26, 2022, Vienna, Austria

of collaborative, content-based, or knowledge-based filtering used in the recommender systems [1]. The main contributions of our work are the review of the current status of recommender systems in incident handling, identification of the most important authors and works, and a review of open and resolved challenges.

This paper is structured in six sections. Section 2 recapitulates the fundamentals of incident management, handling, and response and highlights the tasks that can be automated or facilitated by decision support and recommender systems. Section 3 provides a brief introduction to recommender systems. Section 4 reviews the relevant literature and highlights prominent publications, researchers, and research artifacts. Section 5 surveys the available tools and presents their taxonomy. Section 6 concludes the paper.

2 INCIDENT HANDLING AND RESPONSE

In this section, we first introduce the topic of incident handling. Subsequently, we describe the incident handling in two dimensions. First, we follow the workflow and comment on the phases of incident handling. We identify the points in the incident handling procedures in which the incident handlers have to make a decision and suggest how a recommender system could be of use. Second, we describe the incident handling in tiers (levels) corresponding to personal responsibilities. Each tier is focused on a different part of the incident handling, which poses diverse requirements for the recommender systems. The section closes with a brief overview of available tools.

2.1 Fundamentals of Incident Handling

Incident management, handling, and response are the fundamental activities provided by a cybersecurity team (CSIRT/CERT) as listed in RFC 2350 [21] or the CSIRT Services Framework [17]. Recently, the incident response has become a task of third-generation SOC as defined by CISCO [37]. Incident management is an umbrella term for any activities related to cyber incidents. Incident handling is a part of incident management alongside, e.g., vulnerability handling. Incident response is often considered a part of incident handling [33]. Other important materials on the fundamentals of incident handling can be found in the form of guides and handbooks; we may recommend The Computer Security Incident Handling Guide [13] by NIST or The Incident Handler's Handbook [31] by the SANS Institute.

Although the workflow of incident handling is well described in the literature and adopted by practitioners, it is not straightforwardly applicable. Each organization, network, or environment is different, and each team uses different equipment to gather and analyze the data, which makes it extremely difficult to provide guidance in all cases and for all environments with a sufficient level of detail. Spring and Illari [54] reviewed human decision-making in incident response and concluded that the existing guidelines do not give advice on how to obtain an overall picture of an incident with incomplete data, e.g., how to generalize hypothesis in terms of time (among distinct events) and space (among devices in the network or possible cyber victims). These steps are usually accomplished via the expert knowledge of a security team. Poor decision-making was addressed by Webb et al. [58], who proposed a model for security risk management considering such deficiencies. Still, the impact of using decision support tools in incident handling practice is not well investigated. A recent exemption is the work of Happa et al. [22], who studied the impact of using visualization and decision support tools in SOC and stated recommendations for future development, such as including contextual data and enabling collaboration. Another issue is the difficult evaluation and applicability of such tools. Even though a methodology for SOC performance evaluations was recently proposed by Rosso et al. [50], it is often limited to case studies and user evaluations [22, 25].

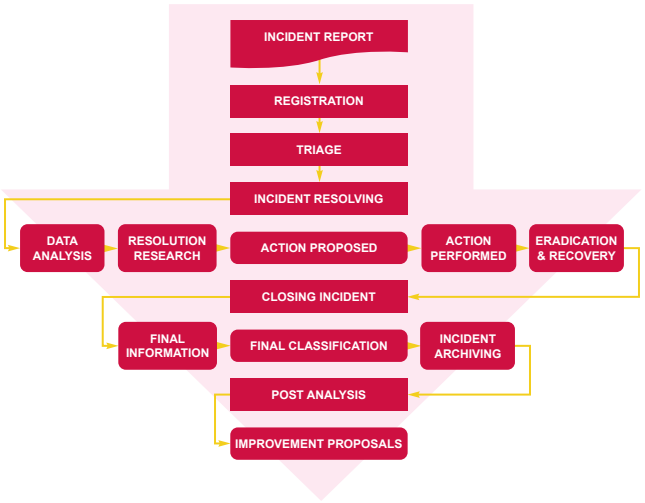

Figure 1: Incident handling workflow elaborated by ENISA [33].

2.2 Workflow and Phases of Incident Handling

The most fundamental workflow of incident handling was introduced by CERT/CC [4]; handling is divided into four phases: detection, triage, analysis, and response. In another well-known example, NIST structures incident handling into four phases: preparation, detection and analysis, eradication and recovery, and post-incident activity [13]. It is worth noting that well-known documents formalizing incident handling are mutually inspired, which makes them, to some extent, similar and interchangeable. In this work, we follow the more elaborated workflow of incident handling described in the guide by ENISA [33] displayed in Figure 1.

In the initial phases, the incident is detected, reported, and registered. Incident detection is out of the scope of incident handling. There are a plethora of tools and approaches to detect an incident described in the literature and used in practice. Although we can proclaim the detection phase as out of the scope of this work, there are several works referenced in this work that also cover the detection phase. We will comment on this issue in the literature review in Section 5. The incident is reported to the cybersecurity team by an intrusion or anomaly detection system, another cybersecurity team (CSIRT/CERT/SOC), IT administrators within the organization, or common users. The incident could be reported by email, phone call, in-person, or via a web form. The cybersecurity team receives the incident report and usually registers it in an incident

handling system under a unique identifier. Reporting and registration are straightforward and do not require any decision support or recommender systems.

Triage is a very interesting phase from the perspective of decision-making; the decisions in this phase have to be taken quickly and often with an insufficient amount of information, whereas the decision affects the rest of the incident handling process. The goal of triage is to prioritize the incident based on its significance and severity and to allocate the resources to handle it. Triage consists of three steps: verification, classification, and assignment.

During the verification, the incident handler checks for three conditions. First, the report must be a real incident report, i.e., not spam or off-topic message. Second, the incident must fall into the team's constituency. For example, it must be related to an IP address in the organization's address range. If this is not true, the incident report should be forwarded to a responsible team. Finally, the incident should not already be reported; otherwise, it is marked as duplicity and not handled further. Checking such conditions is straightforward and can be performed effectively by humans or automatically, so there is no need for decision support in this phase.

The classification is a crucial task that influences all the following procedures. The incident needs to be classified according to the scheme in use by the team (there are many of them, and cybersecurity teams often use their own). The classification then determines which procedures should be done in the incident analysis and response. There are very little data available at this stage, which makes this task very difficult for humans and even for any possible decision support system. Decision support is more than welcome in this phase and approached by researchers [61].

The assignment refers to assigning the incident to a particular handler, either the person on duty at the moment or to a specialist on the particular type of incident. Some incidents may require the assignment of more handlers. We elaborate more on the assignment in the following subsection, where it plays a bigger role.

Incident resolution is the longest and most laborious phase of incident handling. It consists of a cycle of data analysis, resolution research, proposing and performing actions, and eradication and recovery of the incident. Usually, the cycle is iterated several times; new data for analysis may appear, or the performed action might not eradicate the incident. Incident resolution is very situational; each cybersecurity team has a different set of data, data analysis tools, and actions available. Thus, the specific procedures or guidelines are often limited to a certain environment, set of tools, or otherwise restricted. This phase is very rich in problems that could be resolved via decision support and recommender systems. In fact, proposing the action (or, more specifically, selecting the attack mitigation option) is one of the most popular applications of recommender systems in cybersecurity [32, 51, 55]. We comment on the particular examples of tools and approaches in Section 5.

When the incident is resolved, there is a need to close it properly. Incident closure consists of assessing final information (e.g., what happened, who were the actors, how it happened), final classification, and archiving of the incident. The incident handlers can also come back to an archived incident for a post-analysis and improvement proposal, although this is often neglected due to the work overload of cybersecurity teams. Yet, proper final classification and archiving of the incident and its details is a valuable resource for decision support and recommender systems that can use the data to train machine learning models or build knowledge bases.

2.3 Tiers

In addition to phases, incident handling can also be divided according to the roles and responsibilities of security team members. This division is generally referred to as *Tiers* or *Levels*, based on the ITIL (Information Technology Infrastructure Library) methodology [11]. In the initial proposal for cybersecurity teams, members were assigned to two tiers [62]. However, as the number of incidents and their sophistication grew, additional tiers were gradually added, focusing on advanced incident analysis, threat hunting, and security team management. As we have mentioned in previous sections, cybersecurity teams often vary widely in their technologies, processes, and services. Therefore, the division into individual tiers is not precisely defined as it is often determined by the organization's management. The following definitions, which we use in the paper, summarize and generalize the descriptions given in the literature [36, 57] and public posts by cybersecurity companies [6, 12, 46].

Initial incident handling is the responsibility of Tier 1 personnel. This level may be referred to as *triage*, where the *triage specialist* or *alert investigator* works. They manage and configure security monitoring tools, monitor alerts, and review them to determine their relevance and urgency (whether a real security incident is taking place). Besides, they perform triage of the incident and collect relevant data. When they identify a cybersecurity incident, they escalate its resolution to Tier 2. Employees at this initial level are typically the least experienced analysts with system administration, programming, and scripting skills.

The main part of handling a security incident happens at Tier 2, referred to as *investigation*, where *incident responders* and *analysts* work. They receive incidents, perform deep analysis, and correlate collected evidence with threat intelligence to identify the threat actor, nature of the attack, and systems or data affected. In more detail, they leverage emerging threat intelligence (e.g., IoCs, updated rules) to identify affected systems and the scope of the attack, review and collect asset data (e.g., configurations, running processes) on these systems, and direct remediation and recovery efforts. These activities have a higher impact and generally require more experienced analysts. At the knowledge level, the same skills are expected at Tier 1, together with knowledge of advanced forensics, malware assessment, threat intelligence, and ethical hacking.

Tier 3 is used to deal with more serious incidents and is referred to as *threat hunting*, where *expert security analysts* and *threat hunters* work. They support complex incident response, cooperate with Tier 2 on major incidents, and spend any remaining time looking through forensic and telemetry data for threats that detection software may not have identified as suspicious. In more detail, they conduct vulnerability assessments, penetration tests, and review alerts, industry news, threat intelligence, and security data. Besides, they actively hunt for threats that have made their way into the network, as well as unknown vulnerabilities and security gaps. Personnel at this level are expected to have the same knowledge as those on the lower levels but with even more experience, including high-level incident analysis, knowledge of penetration testing tools,

cross-organizational data visualization, malware reverse engineering, and identifying and developing responses to new threats and attack patterns.

Tier 4 is no longer focused directly on incident handling but covers all the activities of the security management team. This typically includes a SOC manager who recruits and evaluates employees, manages the escalation process, and reviews incident reports. They develop and execute crisis communication plans for the chief information security officer and other stakeholders, run compliance reports, support the audit process, and measure SOC performance metrics. In terms of knowledge, personnel at this level should understand the tools and approaches used in lower tiers together with project management skills, incident response management training, and strong communication skills.

2.4 Tools

There are numerous tools supporting the incident handling workflow. A comprehensive, community-maintained list of them can be found in a popular form of an "awesome list" on GitHub[1]. Although there are several implementations of incident management software, many of them implementing full incident handling workflow or incident life cycle, only a few of them are widely used in the community. Many cybersecurity teams also develop their own solutions and publish their source codes, but it is often problematic to transfer such complex systems into a different environment. A well-known incident handling system is RTIR[2], a modification of the popular Request Tracker (RT) specifically designed for incident response needs. RTIR is a popular choice among established cybersecurity teams that do not intend to build their own solution but welcome the options of enhancing the existing one with custom modules and connectors. Another interesting and well-known tool is TheHive[3], which facilitates mostly the incident analysis phase by connecting to MISP (Malware Information Sharing Platform), which enables investigation with many data from cyber threat intelligence providers. Such features sparked interest in TheHive project in recent years and made it a popular choice among newly founded teams.

An interesting example of a commercial tool by Basis Technology appeared recently. The Cyber Triage[4] features a recommendation engine to assist in digital forensics. Although this is a commercial product and no implementation details are publicly available, it seems to be a variation of a classical product recommender system based on the author's quote "you may like this process based on your interest in this file" on the project blog[5].

3 RECOMMENDER SYSTEMS

The recommender systems are information filtering systems that not only researchers' attention due to vast application options but also become recognized by the general public. Recommender systems are part of information systems accessed by millions of users on a daily basis, including social networks, news sites, e-shops, and multimedia streaming services. Such systems are able to recommend goods to buy, people to connect with, or movies to watch based on users' preferences and history of action. Business applications are a significant driving force in the development of recommender systems [1, 47]. The recommender system used by Netflix [20] has become so well-known that many researchers admit to being inspired by it, even in cybersecurity [52]. We do not intend to delve into the details of recommender systems in this paper; there are many books and surveys on this wide area of research. We may recommend an exhaustive textbook by Charu C. Aggarwal [1]. An interesting introduction to the topic intended for cybersecurity experts was presented by Pawlicka et al. [47]. Herein, we recapitulate the most important features and classification of such systems.

There are two versions of the recommender problem and four goals to achieve by the recommender systems [1]. The problem is either a prediction or a top-k recommendation. The prediction problem (also known as the matrix completion problem) is predicting an unknown value using a training dataset. For example, based on a user's ratings of similar movies, the recommender system predicts how the user will rate a new movie. There are not many possible applications of this problem in cybersecurity because there are often not that many users and items (or similar combinations of entities). However, it could be applied to a classification of threat intelligence data or recommending possible attack paths. The top-k recommendation problem (or ranking problem) can be illustrated as recommending top-k items to a particular. Such a problem could be applied to incident handling to recommend a list of actions to perform in response to an incident.

The four generic goals of recommender systems are relevance, novelty, serendipity, and diversity. Relevance is obvious and most important; only the relevant recommendations are usable, but in many cases, the other three goals influence the final recommendation when there are multiple relevant options. Novelty means recommending items that are somewhat novel; a recent movie could be favored over an older classic to support the new movie's sales, or a recommendation of a recently discovered vulnerability might get attention over an older, already patched one. Serendipity corresponds to recommending more surprising results, unexpected by a user. When considering the vulnerability scenario again, a user interested in web security might be aware of recent vulnerabilities in popular web frameworks but could be surprised by a recommendation of a relevant vulnerability in OS or web server. Finally, the diversity ensures that the recommendations are not repeated or monothematic; the user may lose interest in such recommendations over time. Diverse recommendations increase the chance of novel or serendipitous findings, which would make this business-oriented criterion relevant also to the cybersecurity field. Recommending unusual or counter-intuitive action in the analysis of exploitation using a zero-day attack could indeed speed up the incident handling.

There are four major types of recommender systems, including their combinations, namely [1]: collaborative filtering, content-based, knowledge-based, and hybrid and ensemble-based. Let us illustrate their differences on a classic example of product recommendation. Collaborative filtering is the most well-known approach to making recommendations. In essence, it predicts the interest of a user based on other users' interests. The item is recommended to the user if it received a good rating from users who gave similar

[1] https://github.com/meirwah/awesome-incident-response
[2] https://bestpractical.com/rtir
[3] https://thehive-project.org/
[4] https://www.cybertriage.com/
[5] https://www.cybertriage.com/new-features/incident-response-recommendation-engine/

ratings to other items as the user in question. Content-based filtering is similar to collaborative filtering but uses attributes instead of ratings. The user is recommended the items with similar attributes as the items the user already expressed interest in. Knowledge-based recommender systems appeared to resolve the issues of items that are purchased rarely or have many variations, which makes it difficult to assess a sufficient amount of ratings. Examples are real estate, automobiles, or financial services. Such systems are based on structured domain knowledge. Finally, hybrid and ensemble-based recommender systems often use a combination of approaches to achieve the best results in a particular application.

It is also worth mentioning that the recommender systems face two main problems, commonly referred to as a new user problem and a cold start problem [1]. The new user problem describes a situation in which a new user starts using the system without a history of actions. Therefore, the recommender system has no input from the particular user and cannot make predictions or rankings even though it has data from other users. The cold start describes the situation where the whole system has no data on user preferences. The cold start is approached by content-based filtering and knowledge-based systems. Nevertheless, such systems require a set of attributes or building a knowledge base, which can be difficult in certain domains.

4 RESEARCH OVERVIEW

In this section, we present the findings of a research and literature review we conducted to enumerate the research publications on the topics of this paper. We used Google Scholar to search for relevant literature using the keywords of *recommender system* or *decision support* combined with *incident handling* or *cybersecurity*. Herein, we enumerate existing surveys, pinpoint the most influential papers and venues, and comment on influential researchers.

4.1 Publication Venues and Existing Surveys

The related work on recommender systems in incident handling or cybersecurity, in general, can be found in many venues, although there is no journal or conference dedicated to this topic. ACM Recommender Systems (RecSys)[6] is a flagship conference on recommender systems. Unfortunately, the applications in cybersecurity are rarely discussed there. The same applies to Elsevier's Decision Support Systems journal[7]. The relevant papers are more likely to be found in Elsevier's Computers & Security[8] or ACM Digital Threats: Research and Practice[9] journals, which recently covered the issues of SOC operations and incident handling. Recent work in progress could be found in the ARES conference[10], namely in its EU workshops. A plethora of short papers on closely related topics could also be found at the Cyber Science conference[11].

To the best of our knowledge, there is only one literature survey on the topic, although very recent. Pawlicka et al. [47] first presented an overview of recommender system types, their advantages and disadvantages, and their possible application in cybersecurity. Subsequently, the paper surveys the state of the art in using recommender systems in cybersecurity. We consider the survey as an excellent introduction to the decision support and recommender system for cybersecurity researchers and practitioners. Readers interested in an exhaustive literature review are also recommended to read the survey. However, Pawlicka et al. do not analyze particular use cases and applications of recommender systems in cybersecurity. The authors only state that such systems are used in cybersecurity and that there is no existing taxonomy. An interesting observation is that many of the authors of the surveyed research works claim to be the first or one of the first to apply a specific method in cybersecurity [47]. This may indeed be caused by a lack of understanding of the recommender systems landscape; it is additional motivation for the systematization of knowledge in this area.

In 2016, Gadepally et al. [19] published an article on using recommender systems in cybersecurity, namely by the Department of Defense, since the authors are affiliated with the Lincoln Laboratories. The authors suggested how could the recommender systems be used and in what context. Moreover, they proposed possible directions for future work. While the technical paths seem to be approached by researchers and developers, the socio-technical aspects do not. For example, we did not encounter any examples of preserving privacy, as discussed by Gadepally et al., in the surveyed works. More importantly, there does not seem to be any fruitful interaction between academia, industry, and other actors who are interested in this topic and could advance the research and development.

4.2 Researchers and Research Groups

Unfortunately, very few researchers or research groups seem to focus on the application of recommender systems in incident handling or cybersecurity in general. Most researchers contribute with only a few publications on the topic. However, there are several notable exceptions. Chen Zhong of the University of Tampa[12] has authored several papers that are within the scope of the paper. Moreover, her contributions to the analysis of incident triage [59–61] are one of the most valuable in this area. We referenced only the most significant works by Chen Zhong; we recommend the readers to check her other works on closely related topics. Some other notable researchers include Aleksandra Pawlicka[13] of ITTI, who authored the survey on recommender systems in cybersecurity [47] but not yet any other related work. Nikolaos Polatidis[14] of the University of Brighton is a researcher in recommender systems with overlaps to cybersecurity, who contributed with several relatively impactful papers [48, 49].

Although we did not identify many influential or productive researchers, the masters' and doctoral students seem to be interested in the topic. We encountered several theses on closely relevant topics. The theses of Katherine B. Lyons [32] from 2014 or Erion Sula [55] from 2019 have already gained attention in the literature [47]. The other examples include the master thesis of Milad Asgari Mehrabadi on a recommender system for predicting privacy leaks in mobile traffic [5] from 2019, the doctoral thesis of Linus Karlsson on preventive measures in cybersecurity [28] from 2019

[6] https://recsys.acm.org/
[7] https://www.journals.elsevier.com/decision-support-systems
[8] https://www.journals.elsevier.com/computers-and-security
[9] https://dl.acm.org/journal/dtrap
[10] https://www.ares-conference.eu/
[11] https://c-mric.org/
[12] https://utweb.ut.edu/chen.zhong/
[13] https://orcid.org/0000-0003-4380-014X
[14] https://research.brighton.ac.uk/en/persons/nick-polatidis

(including the author's paper on recommender system for vulnerability scoring [29]), or the master's thesis of Suzy Edith Moukala Both on personalized question-based recommender system for cybersecurity [35] from 2021. We are pleased to see junior researchers working on such topics. This could be an indicator that the topic is going to be further explored in the future.

5 OVERVIEW OF TOOLS AND APPROACHES

In this section, we provide an overview of the recommender systems applied to resolving issues related to incident handling. The papers and tools we found in the research overview are grouped into several categories loosely based on incident handling phases and tiers. Some of the works discussed in this section were also discussed in the survey by Pawlicka et al. [47]. We comment on such papers only briefly; the readers are kindly referred to the survey for more details, especially those related to the recommender systems in use. We do not delve deep into the implementation details since we are more interested in the applications and use cases in relation to incident handling.

5.1 Tier 1 and Incident Triage

The first groups of papers are relevant to the initial phases of incident handling on lower tiers, namely the problem of triage and Tier 1 incident handling. The problems tackled in this section are recommending the most suitable workflow for handling a new incident. This is not yet the analysis; the incident handlers at this phase have to decide if the incident is real and relevant, how to categorize it, which and how many analysts it should be assigned to, or which guidelines or playbook they should follow. While some teams may have a dedicated triage specialist or all the incident handlers (junior and senior) may participate in triage, it is also possible that the initial phases are offloaded to the IT helpdesk or another generic contact point in the organization, where the operators might not even have any skills in cybersecurity. Nevertheless, most of the discussed works assume this is not the case, and incident handlers are trained in cybersecurity, at least on an elementary level.

The problem of incident triage was studied by Zhong et al., who in 2015 proposed ARSCA, a tool for tracing triage actions [59]. The idea behind ARSCA is to collect the traces of triage actions conducted by incident handlers for a deeper understanding of the triage and decisions made in it. Such traces were proposed to be used for incident handling automation [60]. As proposed by the authors, the traces of senior incident handlers' actions can be used to train junior handlers [61]. We point out that such data are highly valuable for decision support or recommender system for incident handling. In this regard, the authors developed a smart retrieval system that looks up historical traces of senior handlers that are the most similar to the traces actually conducted by a junior handler [61]. We consider this work as one of the closest to the idea of a recommender system for incident handling.

Esposte et al. [14] approached the issue of so-called alert flooding in 2016. Indeed, incident handlers very often face information overload and the influx of an overwhelming number of alerts; many of them are also not relevant. The authors proposed a recommender system that collects alerts from external sources and recommends them to any person profiled as a network administrator. The alerts are filtered via the administrator's preferences and ratings of previous alerts. On the one hand, this work seems highly relevant and applies a full-scale recommender system. On the other hand, the authors do not seem to follow any incident handling workflow and approach the cybersecurity perspective quite vaguely. A similar problem was approached by Ayala et al. [7, 8] in 2021. The authors proposed a hybrid recommender system combining collaborative and knowledge-based filtering to recommend cybersecurity incidents (namely anomalies and vulnerabilities) to the users. The goal of the system is to leverage expert knowledge to prioritize the incidents and recommend the most urgent ones to resolve. An interesting feature of this work is the use of knowledge-based filtering that approaches the cold start problem of a recommender system.

In 2021, Kraeva and Yakhyaeva [30] approached the automatic selection of the most relevant incident playbook by which the incident should be handled. Instead of common rule-based methods, the proposed system uses metric learning and neural networks and does not require the involvement of domain experts or additional training. Indeed, this seems like a highly relevant contribution to the triage problem. Unfortunately, the paper is very brief and focuses on the learning and recommendations; it is only proposed to use such an approach.

We did not find any other works on applying recommender systems to the initial incident handling phases. However, we would like to raise attention to two works that might be relevant in the future. Both could be used to support the triage phase and significantly reduce the volume of data to process. ASSERT by Okutan and Yang [42] assigns individual intrusion detection alerts to empirical attack models and highlights the critical activities and aggregated statistics that can be used to predict future attack actions. PATRL by Moskal and Yang [34] translates intrusion detection alert descriptions into an intuitive description of an attack campaign in the form of MITRE ATT&CK[15]. The authors developed a machine learning process that maps over 64,000 Suricata alerts in a small set of Action-Intent-Stages derived from ATT&CK. Filtering a large number of entries into a small set of options is a common objective with recommender systems. We can see a potential application of a tool similar to PATRL that would translate alerts into a comprehensible explanatory note. The incident handlers would then not have to spend valuable time finding out what a particular alert description means. PARTL can save time spent by looking up details for a CVE number (if included in the alert), while ASSERT facilitates learning about the potential impact. Both tools are built on solid theoretical backgrounds and seem ready for evaluation in an operational environment. We believe they can be a valuable contribution to decision-making in incident handling.

5.2 Tier 2 and Incident Analysis and Response

In the next group of approaches, we present the works on incident analysis and response corresponding to Tier 2 of incident handling. It is assumed that at this point in the incident handling workflow, the incident handler already has certain knowledge of the incident (e.g., its category and impact) and investigates technical details and response options. Typical goals include collecting evidence,

[15] https://attack.mitre.org/

conducting basic forensic analysis, finding other impacted systems or users, and selecting mitigation and recovery actions.

The KRAKEN recommender system, proposed by Brisse et al. [10] in 2021, helps incident handlers by suggesting exploration paths. KRAKEN is using the knowledge of advanced attack descriptions observed in the real world and confronts it with the investigated logs. The goal of KRAKEN is to dynamically recommend relevant parts of the dataset to explore during the investigation. KRAKEN aims at the investigation of advanced persistent threats and was evaluated with the help of senior cybersecurity analysts. An interesting fact is that the authors build upon existing tools for visualization of and navigation in the investigated data. An interesting older work cited by Brisse et al. is NAVSEC proposed by Nunnaly et al. [41] in 2013. NAVSEC is a network security visualization tool that features a recommendation engine to guide a user through the data.

In 2021, Martin Husák proposed a recommender system that would help handle the incidents like ransomware [24]. This work builds on the assumption that ransomware and other types of malware spread in the network, and when an infection is reported, there is a risk that the hosts in the proximity of the infected host are in danger or already infected, too. The proposed approach is quite simple from the perspective of recommender systems; it looks up the hosts in the proximity of a reported infected host and sorts them by their similarity to the infected one. Several metrics of proximity and similarity are proposed and combined. However, implementation and assignment of weight to particular metrics were left for future work. Nevertheless, the work poses an interesting application of recommender systems in the incident analysis phase.

Sworna et al. [56] recently proposed APIRO, a framework that helps in incident handling by recommending which security tool to use. Their work is very technical and goes down to the recommendation of a particular API to call, e.g., use MISP to download IoCs and inspect the network traffic via Snort. The very promising feature of APIRO is the single unified interface to communicate with heterogeneous data and tools. The user could interact with only one system, which then forwards them to the specialized tools. The manuscript seems to be valuable but is available only on arXiv.

A popular task in cybersecurity research is finding an optimal countermeasure to a cyber attack. There is a plethora of research works on this topic based on different approaches, from game theory to recommender systems. We may recommend a survey by Nespoli et al. [39]. A prime example is the work of Ossenbühl et al. [44], who proposed REASSESS, a countermeasure selection system aligned with the NIST incident life cycle [38], to automate incident handling. Although it is not based on a recommender system, it shows a common issue of such works; there are not many countermeasures to select from. The authors list only three options, IP traffic filtering, traffic limiting, and no action. Even if we consider more options for manipulating network traffic, such as routing the traffic to a honeypot, we still have only a few options; not all of them could be available in practice. Thus, the question here should not be which action to select but rather whether to act or not.

Nevertheless, three proposals of recommender systems connected to incident response selection are worth mentioning. In 2014, Katherine B. Lyons [32] proposed a combination of the intrusion detection system (IDS) and a recommender system in her doctoral thesis. When the IDS system raises an alert, the recommender system reacts with a prioritized list of actions to do. In 2017, Sayan et al. [51] proposed ICSA, combined attack detection and defense recommendation system. Contrary to the other works, ICSA is really a hybrid solution. Instead of only reacting to what is detected by the IDS, ICSA makes predictions of upcoming attacks and simultaneously recommends their mitigation, which makes it an interesting intelligent intrusion prevention system.

A highly specific problem was approached by Erion Sula [55] in 2019. The authors proposed a recommender system that assists in DDoS mitigation by selecting the cheapest option. Although it aims solely at one type of attack, it also illustrates the depth of cybersecurity problems. The proposed system considers a variety of features, ranging from attack type to the region and location of the victim to the budget available to the defenders. Indeed, there is a lot to consider even when mitigating one specific type of attack.

Recommender systems are also considered in vulnerability assessment and scoring. Although vulnerability management is a separate task from incident handling, it is typically approached in a similar manner by the cybersecurity teams. The vulnerabilities are mostly recognized by their CVE number, and their potential impacts and other scores are usually filled into a CVSS record. However, such records are global and aim for general use, not considering specifics of particular environments or user preferences. Therefore, in 2020, Karlsson et al. [29] implemented a recommender system that considers users' history and preferences and recommends vulnerabilities close to the user's preference or interests. This work is an excellent example of a recommender system in cybersecurity because the approach is very similar to the well-known product recommendation systems. A more sophisticated recommender system was proposed by Huff et al. [23] in 2021. The authors propose a system that collects software names and versions in the form of CPE strings, links them to the CVE records, and recommends the vulnerabilities that are the most relevant for the given network. Such a tool has the potential to facilitate vulnerability and patch management by recommending a small set of relevant CVEs from the huge and continuously growing database.

5.3 Tier 3 and Intelligence

At Tier 3, advanced analyses are conducted. The focus of the analysts is directed toward threat intelligence, risk assessment, and attack prevention rather than towards handling a particular incident. While such tasks are usually conducted by skilled analysts, they are complex and include the processing of large volumes of heterogeneous data from many sources, which calls for the use of recommender systems. Collaborative filtering can be applied to global cyber threat intelligence (CTI) data, while knowledge-based systems are promising in approaching heterogeneous inputs.

In 2019, Du et al. [15] proposed People-Readable Threat Intelligence (PRTI), a highly condensed knowledge graph in which the noise is filtered, and the items are comprehensive to human analysts. Subsequently, they propose a method recommending PRTI to the users based on their interests. However, the core of the work is in the construction of the knowledge graph. In 2020, Ou et al. [45] proposed building a large knowledge graph and applying collaborative filtering to it. First, they construct a knowledge graph from the data

in four public knowledge bases (vulnerability, weakness, exploit, and attack pattern databases) and apply text mining to discover threat information from the descriptions in natural language. Subsequently, the paths in the knowledge graph are recommended. The authors claim their approach provides diverse recommendations (see Section 3) and is more accurate than content-based filtering.

The application of recommender systems can also be found in cyber attack prevention. One of the first applications of recommender systems in cybersecurity is the work by Polatidis et al. [48] from 2017. The authors proposed a system that finds all possible attack paths in complex infrastructure and then recommended the paths that are most likely to be exploited. The approach is inspired by product recommendation systems, although they use a combination of several approaches and state they are open to using other algorithms if the task demands it. The authors revisited their approach in a more recent publication from 2020 [49].

Another use of recommender systems in attack prevention is for predictive blacklisting, as proposed by Soldo et al. [52, 53]. The authors proposed an implicit recommendation system that is used to build a blacklist of network entities that are likely to behave maliciously in the future. The motivation for this work is that while some network entities are detected to behave maliciously over and over, many more are detected a few times in a short period of time and not again. The problem of effective blacklisting is selecting (or predicting) those entities that will likely continue in malicious activities. A small blacklist of 100 IP addresses may achieve higher hit rates (preventing more attacks) than a raw blacklist of thousands of IP addresses [9]. Soldo et al. [52] were first inspired by the recommender system of Netflix [20] and then enhanced their work with spatio-temporal features to improve its effectiveness. Although the topic was further investigated by other authors, most recently by Bartoš et al. [9], they do not fall under the scope of recommender systems.

5.4 Tier 4 and Incident Management

In the last group of surveyed works, we present the approaches to applying recommender systems to incident management on the top tier. While previously presented works dealt with current incidents and other urgent needs, the works presented here are aimed at the management, planning, and investments.

Earlier, we presented recommender systems applied to incident response selection. A similar issue can be resolved by recommender systems on a higher level as the problem of recommending the most suitable solution to cyber infrastructure protection. Briefly, instead of recommending which action to take in response to an attack, the goal is to recommend equipment and tools that are worth purchasing and deploying. The issues of optimal investment into defense tools [16] or their optimal placement [40] were studied in the past without the notion of recommender systems.

In 2019, Franco et al. [18] proposed MENTOR, a tool for recommending the most suitable solution to network protection. MENTOR considers the customer's profile, budget, and requirements. It assesses the customer's infrastructure and attacks against it. Subsequently, it recommends a solution (list of products and services) to protect the infrastructure and evaluates if it fulfills the customer's requirements. Factors like the availability of a product in a region and its pricing are also considered. The authors further discuss how various properties (e.g., price) may drastically change the recommendation when specific algorithms are used, which is a valuable outcome for researchers in the field.

Decision support and recommender systems on the top level are not intended only for the large organizations but also for small and medium enterprises (SMEs), who very often lack expertise in this area. While this is an open issue, Ahmed et al. [2] preliminarily mapped the cybersecurity awareness and capabilities among the SMEs in the Middle East and proposed a recommender system that suggests the most suitable solution for a particular SME. Similar works typically provide only the best practices [27].

5.5 Taxonomy

The proposed taxonomy of recommender systems in incident handling is presented in Figure 2. First, we assign the recommender systems under the phases of incident handling in which it is applied. A simple 4-phase workflow of incident handling [4, 33] is used. Simultaneously, we categorize the works by the incident handling tiers, where they are supposed to be used. Even though there are similar tasks that are resolved at multiple levels differently, we found a mapping in which the phases and tiers and exchangeable. Not all recommender systems resolve only one issue. In such cases, we chose the dominant application, such as in the case of the work of Lyons [32], who integrated an incident response selector with an intrusion detection system. The work of Sayan et al. [51] remains the only hybrid approach as it provides recommendations simultaneously for detection and response.

At Tier 1, triage is the most interesting phase from the decision-making perspective, but the number of works that approached this topic is lower than expected (apart from a number of publications by Zhong et al. [59–61] that mostly do not propose a recommender system). Tier 2 includes incident analysis and response, currently the most popular recommender system application in cybersecurity. The tools that would assist the incident handlers in the analytical process from a wider perspective are scarce. The surveyed works mostly aim at recommending which path to take in the investigation, which tool to use, or which target to inspect but not on their combinations. The other works cover response selection and vulnerability management. Tier 3 covers highly specialized tools and advanced approaches to threat intelligence, predictive analytics, and preemptive measures. The analysts may benefit namely from collaborative, content-based, and knowledge-based filtering of threat intelligence data that are large in volume and difficult to analyze. At Tier 4, we identified several works that facilitate planning and investments in cyber defense and complement the response selection tools.

6 CONCLUSION

In this paper, we provided an overview of using recommender systems in incident handling. Our work was motivated by the needs of cybersecurity teams and operation centers (CSIRT, CERT, SOC), which have to deal with a scarce workforce and information overload while their members need to make many decisions when handling an incident [22, 50, 57]. Under such conditions, any means of automation or decision support are appreciated [3, 25]. We first

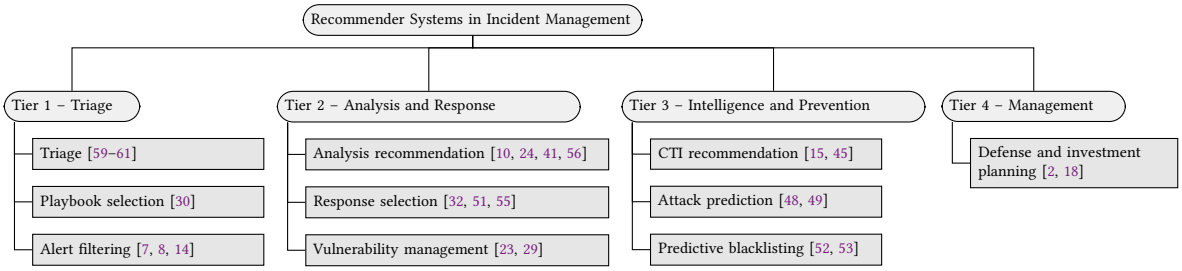

Figure 2: Taxonomy of recommender systems in cybersecurity.

introduced the incident handling and recommender systems. Subsequently, we presented the outcomes of the literature review and overview approaches discussed in the literature. We also proposed a taxonomy of recommender systems in incident handling.

The foremost finding of our work is that there are almost no recommender systems that would approach the incident handling holistically; the researchers and developers rather focus on particular tasks. Moreover, not all the works that propose a recommender system actually implement any filtering method (collaborative, content-based, or knowledge-based) but instead use other mechanisms to support decision-making. The most popular topics include recommending the most suitable method of incident suppression or mitigation [32, 51, 55]. Recommender systems were also proposed to approach particular tasks of incident detection [51], prevention [48, 49, 52, 53], and analysis [24]. However, the incident triage, which would benefit from the recommender systems the most, was not approached very often and seems to be a very difficult problem [60, 61]. Even though triage should be prompt and steers the whole process, it is suggested that it takes too much time and is often conducted by the least skilled personnel. Simultaneously, the existing recommender systems approach highly specific tasks conducted by skilled personnel.

We argue that one of the factors in this discrepancy is that researchers might prefer niche problems that are interesting from a scientific perspective, i.e., they were not studied before, and they can be experimentally approached and evaluated. On the contrary, the common problems, such as triage, would be very difficult to evaluate in a controlled environment conforming to research practices [22, 50]. We also admit that in the triage phase, there are insufficient data to process by automated tools, which brings high risks in the research and experimental development of such tools [60, 61]. If the data are available, they are often large in volume, and there is a need to filter them [7, 8, 14]. Heterogeneity of the data is also an important factor that hinders progress [10]. Therefore, new scientific results and progress could be expected in the niche areas, while the common tasks like triage will be more likely approached by practitioners or developers of widely used tools, who may facilitate the decision-making by simple yet effective improvements like tooltips, visualizations, and auxiliary data retrieval. A mature recommender system for incident handling that would resolve the triage and assist in incident analysis and response is not expected to appear shortly; we guess it may take several more years. On the contrary, the good news is that we encountered several masters' and doctoral theses on similar topics, which indicates there is an increasing interest in the research community.

ACKNOWLEDGMENTS

This research was supported by ERDF "CyberSecurity, CyberCrime and Critical Information Infrastructures Center of Excellence" (No. CZ.02.1.01/0.0/0.0/16_019/0000822).

REFERENCES

[1] Charu C. Aggarwal. 2016. *Recommender Systems: The Textbook.* Springer, Cham.
[2] Nadir Naveed Ahmed and Krishnadas Nanath. 2021. Exploring Cybersecurity Ecosystem in the Middle East: Towards an SME Recommender System. *Journal of Cyber Security and Mobility* 10 (2021), 511–536. Issue 3.
[3] Massimiliano Albanese, Hasan Cam, and Sushil Jajodia. 2014. *Automated Cyber Situation Awareness Tools and Models for Improving Analyst Performance.* Springer, Cham, 47–60.
[4] Chris Alberts, Audrey Dorofee, Georgia Killcrece, Robin Ruefle, and Mark Zajicek. 2004. Defining incident management processes for CSIRTs: A work in progress. Software Engineering Institute, Carnegie Mellon University.
[5] Milad Asgari Mehrabadi. 2019. *A Recommendation System for Predicting Privacy Leaks in Mobile Traffic.* Master's thesis. UC Irvine.
[6] AT&T. 2021. The security operations center (SOC) team: operations & responsibilities. https://cybersecurity.att.com/solutions/security-operations-center/building-a-soc/soc-team.
[7] Carlos Ayala, Kevin Jiménez, Edison Loza-Aguirre, and Roberto O. Andrade. 2021. *A Hybrid Recommender for Cybersecurity Based on Rating Approach.* Springer, Cham, 445–462.
[8] Carlos Ayala, Kevin Jimenez, Edison Loza-Aguirre, and Roberto O. Andrade. 2021. A Hybrid Recommender System for Cybersecurity Based on a Rating Approach. In *Advances in Security, Networks, and Internet of Things.* Springer, Cham, 397–409.
[9] Vaclav Bartos, Martin Zadnik, Sheikh Mahbub Habib, and Emmanouil Vasilomanolakis. 2019. Network entity characterization and attack prediction. *Future Generation Computer Systems* 97 (2019), 674–686.
[10] Romain Brisse, Simon Boche, Frédéric Majorczyk, and Jean-François Lalande. 2021. KRAKEN: A Knowledge-Based Recommender system for Analysts, to Kick Exploration up a Notch. In *14th International Conference on Security for Information Technology and Communications.* 17 pages.
[11] Judith M. Brown, Steven Greenspan, and Robert Biddle. 2016. Incident response teams in IT operations centers: the T-TOCs model of team functionality. *Cognition, Technology & Work* 18, 4 (2016), 695–716.
[12] Cassetto, Orion. 2022. Security Operations Center Roles and Responsibilities. https://www.exabeam.com/security-operations-center/security-operations-center-roles-and-responsibilities/.
[13] Paul Cichonski, Tom Millar, Tim Grance, and Karen Scarfone. 2012. Computer Security Incident Handling Guide: Recommendations of the National Institute of Standards and Technology. NIST Special Publication 800-61.
[14] Arthur de Moura Del Esposte, Rodrigo Campiolo, Fabio Kon, and Daniel Batista. 2016. A Collaboration Model to Recommend Network Security Alerts Based on the Mixed Hybrid Approach. (2016), 14 pages.
[15] Ming Du, Jun Jiang, Zhengwei Jiang, Zhigang Lu, and Xiangyu Du. 2019. PRTIRG: A Knowledge Graph for People-Readable Threat Intelligence Recommendation. In *Knowledge Science, Engineering and Management.* Springer, Cham, 47–59.
[16] Andrew Fielder, Emmanouil Panaousis, Pasquale Malacaria, Chris Hankin, and Fabrizio Smeraldi. 2016. Decision support approaches for cyber security investment. *Decision Support Systems* 86 (2016), 13–23.

[17] FIRST – Forum of Incident Response and Security Teams. 2019. Computer Security Incident Response Team (CSIRT) Services Framework. https://www.first.org/standards/frameworks/csirts/csirt_services_framework_v2.1.

[18] Muriel Figueredo Franco, Bruno Rodrigues, and Burkhard Stiller. 2019. MENTOR: The Design and Evaluation of a Protection Services Recommender System. In *2019 15th International Conference on Network and Service Management (CNSM)*. IEEE, New York, NY, USA, 7 pages.

[19] Vijay N Gadepally, Braden J Hancock, Kara B Greenfield, Joseph P Campbell, William M Campbell, and Albert I Reuther. 2016. Recommender systems for the department of defense and intelligence community. *Lincoln Laboratory Journal* 22, 1 (2016), 74–89.

[20] Carlos A. Gomez-Uribe and Neil Hunt. 2016. The Netflix Recommender System: Algorithms, Business Value, and Innovation. *ACM Transactions on Management Information Systems* 6, 4, Article 13 (12 2016), 19 pages.

[21] Erik Guttman and Nevil Brownlee. 1998. Expectations for Computer Security Incident Response. RFC 2350.

[22] Jassim Happa, Ioannis Agrafiotis, Martin Helmhout, Thomas Bashford-Rogers, Michael Goldsmith, and Sadie Creese. 2021. Assessing a Decision Support Tool for SOC Analysts. *Digital Threats: Research and Practice* 2, 3, Article 22 (6 2021).

[23] Philip Huff, Kylie McClanahan, Thao Le, and Qinghua Li. 2021. A Recommender System for Tracking Vulnerabilities. In *The 16th International Conference on Availability, Reliability and Security* (Vienna, Austria) *(ARES 2021)*. ACM, New York, NY, USA, Article 114, 7 pages.

[24] Martin Husák. 2021. Towards a Data-Driven Recommender System for Handling Ransomware and Similar Incidents. In *2021 IEEE International Conference on Intelligence and Security Informatics (ISI)*. IEEE, New York, NY, USA, 6 pages.

[25] Martin Husák, Lukáš Sadlek, Stanislav Špaček, Martin Laštovička, Michal Javorník, and Jana Komárková. 2022. CRUSOE: A toolset for cyber situational awareness and decision support in incident handling. *Computers & Security* 115 (2022), 102609.

[26] International Information System Security Certification Consortium. 2019. Strategies for Building and Growing Strong Cybersecurity Teams. https://www.isc2.org/Research/2019-Cybersecurity-Workforce-Study.

[27] Keshav Kapoor, Karen Renaud, and Jacqueline Archibald. 2018. Preparing for GDPR: helping EU SMEs to manage data breaches. In *Symposium on Digital Behaviour Intervention for Cyber Security*. AISB, 13–20.

[28] Linus Karlsson. 2019. *Contributions to Preventive Measures in Cyber Security*. Ph. D. Dissertation. Lund University, Sweden.

[29] Linus Karlsson, Pegah Nikbakht Bideh, and Martin Hell. 2020. A Recommender System for User-Specific Vulnerability Scoring. In *Risks and Security of Internet and Systems*. Springer, Cham, 355–364.

[30] Irina Kraeva and Gulnara Yakhyaeva. 2021. Application of the Metric Learning for Security Incident Playbook Recommendation. In *2021 IEEE 22nd International Conference of Young Professionals in Electron Devices and Materials (EDM)*. IEEE, New York, NY, USA, 475–479.

[31] Patrick Kral. 2012. The Incident Handler's Handbook. https://www.sans.org/reading-room/whitepapers/incident/incident-handlers-handbook-33901.

[32] Katherine B. Lyons. 2014. *A recommender system in the cyber defense domain*. Master's thesis. Air Force Institute of Technology.

[33] Miroslaw Maj, Roeland Reijers, and Don Stikvoort. 2010. Good Practice Guide for Incident Management. https://www.enisa.europa.eu/publications/good-practice-guide-for-incident-management.

[34] Stephen Moskal and Shanchieh Jay Yang. 2021. Translating Intrusion Alerts to Cyberattack Stages using Pseudo-Active Transfer Learning (PATRL). In *2021 IEEE Conference on Communications and Network Security (CNS)*. IEEE, New York, NY, USA, 110–118.

[35] Suzy Edith Moukala Both. 2021. *Personalized question-based cybersecurity recommendation systems*. Master's thesis. University of Montreal.

[36] Joseph Muniz, Aamir Lakhani, Omar Santos, and Moses Frost. 2021. *The Modern Security Operations Center*. Addison-Wesley Professional, Boston, MA, USA.

[37] Joseph Muniz, Gary McIntyre, and Nadhem AlFardan. 2016. *Security Operations Center*. Cisco Press.

[38] National Institute of Standards and Technology (NIST). 2018. *Framework for Improving Critical Infrastructure Cybersecurity*.

[39] Pantaleone Nespoli, Dimitrios Papamartzivanos, Félix Gómez Mármol, and Georgios Kambourakis. 2018. Optimal Countermeasures Selection Against Cyber Attacks: A Comprehensive Survey on Reaction Frameworks. *IEEE Communications Surveys Tutorials* 20, 2 (Secondquarter 2018), 1361–1396.

[40] Steven Noel and Sushil Jajodia. 2008. Optimal IDS Sensor Placement and Alert Prioritization Using Attack Graphs. *Journal of Network and Systems Management* 16, 3 (2008), 259–275.

[41] Troy Nunnally, Kulsoom Abdullah, A. Selcuk Uluagac, John A. Copeland, and Raheem Beyah. 2013. NAVSEC: A Recommender System for 3D Network Security Visualizations. In *Proceedings of the Tenth Workshop on Visualization for Cyber Security* (Atlanta, GA, USA) *(VizSec '13)*. ACM, New York, NY, USA, 41–48.

[42] Ahmet Okutan and Shanchieh Jay Yang. 2019. ASSERT: attack synthesis and separation with entropy redistribution towards predictive cyber defense. *Cybersecurity* 2, 1 (2019).

[43] Cyril Onwubiko and Karim Ouazzane. 2020. SOTER: A Playbook for Cybersecurity Incident Management. *IEEE Transactions on Engineering Management* (2020), 21 pages. Early access.

[44] Sven Ossenbühl, Jessica Steinberger, and Harald Baier. 2015. Towards Automated Incident Handling: How to Select an Appropriate Response against a Network-Based Attack?. In *2015 Ninth International Conference on IT Security Incident Management & IT Forensics*. IEEE, New York, NY, USA, 51–67.

[45] Yunjia Ou, Tianyang Zhou, and Junhu Zhu. 2020. Recommendation of cyber attack method based on knowledge graph. In *2020 International Conference on Computer Engineering and Intelligent Control (ICCEIC)*. IEEE, New York, NY, USA, 60–65.

[46] Palo Alto Networks. 2022. What Is a SOC? https://www.paloaltonetworks.com/cyberpedia/what-is-a-soc.

[47] Aleksandra Pawlicka, Marek Pawlicki, Rafał Kozik, and Ryszard S. Choraś. 2021. A Systematic Review of Recommender Systems and Their Applications in Cybersecurity. *Sensors* 21, 15 (2021), 25 pages.

[48] Nikolaos Polatidis, Elias Pimenidis, Michalis Pavlidis, and Haralambos Mouratidis. 2017. Recommender Systems Meeting Security: From Product Recommendation to Cyber-Attack Prediction. In *Engineering Applications of Neural Networks*. Springer, Cham, 508–519.

[49] Nikolaos Polatidis, Elias Pimenidis, Michalis Pavlidis, Spyridon Papastergiou, and Haralambos Mouratidis. 2020. From product recommendation to cyber-attack prediction: Generating attack graphs and predicting future attacks. *Evolving Systems* 11, 3 (2020), 479–490.

[50] Martin Rosso, Michele Campobasso, Ganduulga Gankhuyag, and Luca Allodi. 2022. SAIBERSOC: A Methodology and Tool for Experimenting with Security Operation Centers. *Digital Threats: Research and Practice* 3, 2, Article 14 (2 2022).

[51] Carla Sayan, Salim Hariri, and George Ball. 2017. Cyber Security Assistant: Design Overview. In *2017 IEEE 2nd International Workshops on Foundations and Applications of Self* Systems (FAS*W)*. IEEE, New York, NY, USA, 313–317.

[52] Fabio Soldo, Anh Le, and Athina Markopoulou. 2010. Predictive blacklisting as an implicit recommendation system. In *2010 Proceedings IEEE INFOCOM*. IEEE, New York, NY, USA, 9 pages.

[53] Fabio Soldo, Anh Le, and Athina Markopoulou. 2011. Blacklisting Recommendation System: Using Spatio-Temporal Patterns to Predict Future Attacks. *IEEE Journal on Selected Areas in Communications* 29, 7 (2011), 1423–1437.

[54] Jonathan M. Spring and Phyllis Illari. 2021. Review of Human Decision-Making during Computer Security Incident Analysis. *Digital Threats: Research and Practice* 2, 2, Article 11 (4 2021).

[55] Erion Sula. 2019. ProtecDDoS: A recommender system for distributed denial-of-service protection services. https://files.ifi.uzh.ch/CSG/staff/franco/extern/theses/BA-Erion-Sula.pdf. Bachelor thesis, University of Zurich, Switzerland, Supervisors: Muriel Franco, Bruno Rodrigues.

[56] Zarrin Tasnim Sworna, Chadni Islam, and Muhammad Ali Babar. 2022. APIRO: A Framework for Automated Security Tools API Recommendation. https://arxiv.org/abs/2201.07959

[57] Manfred Vielberth, Fabian Böhm, Ines Fichtinger, and Günther Pernul. 2020. Security Operations Center: A Systematic Study and Open Challenges. *IEEE Access* 8 (2020), 227756–227779.

[58] Jeb Webb, Atif Ahmad, Sean B. Maynard, and Graeme Shanks. 2014. A situation awareness model for information security risk management. *Computers & Security* 44 (2014), 1–15.

[59] Chen Zhong, John Yen, Peng Liu, Rob Erbacher, Renee Etoty, and Christopher Garneau. 2015. ARSCA: a computer tool for tracing the cognitive processes of cyber-attack analysis. In *2015 IEEE International Multi-Disciplinary Conference on Cognitive Methods in Situation Awareness and Decision*. IEEE, New York, NY, USA, 165–171.

[60] Chen Zhong, John Yen, Peng Liu, and Robert F. Erbacher. 2016. Automate Cybersecurity Data Triage by Leveraging Human Analysts' Cognitive Process. In *2016 IEEE 2nd International Conference on Big Data Security on Cloud (BigDataSecurity), IEEE International Conference on High Performance and Smart Computing (HPSC), and IEEE International Conference on Intelligent Data and Security (IDS)*. IEEE, New York, NY, USA, 357–363.

[61] Chen Zhong, John Yen, Peng Liu, Rob F. Erbacher, Christopher Garneau, and Bo Chen. 2017. *Studying Analysts' Data Triage Operations in Cyber Defense Situational Analysis*. Springer, Cham, 128–169.

[62] Carson Zimmerman. 2014. *Ten Strategies of a World-Class Cybersecurity Operations Center*. The MITRE Corporation.