

A Natural Language Processing-based Approach for Cyber Risk Assessment in the Healthcare Ecosystems

Stefano Silvestri^{1,*}, Giuseppe Tricomi^{1,2,3}, Giuseppe Felice Russo¹ and Mario Ciampi¹

¹Institute for High Performance Computing and Networking, National Research Council of Italy (ICAR-CNR), via Pietro Castellino 111, Naples, 80131, Italy

²Università degli Studi di Messina, Contrada di Dio 1, Messina, 98166, Italy

³CINI—Consorzio Interuniversitario Nazionale per l'Informatica, Via Ariosto 25, Roma, 00185, Italy

Abstract

The cyber risk in the healthcare sector is constantly increasing, due the large adoption of digital services formed by a complex interconnection of different systems and technologies, which offer a larger attack surface for the attackers. Therefore, the risk assessment of the assets involved in these services is crucial to prevent and mitigate possible critical consequences, which could also affect the health of the patients. A large source of constantly updated information about threats and vulnerabilities of the assets of the healthcare ecosystems is available in natural language text on the Internet (cyber security news, forum, social media, etc.), but it is not easy to fully exploit them for a risk assessment process, due to the complexity of natural language. This paper proposes an AI-based approach for the individual risk assessment of the assets of digital healthcare systems based on the use of NLP and Knowledge Bases, which exploits the information extracted from natural language news from the web. The methodology has been developed within the activities of the EC-funded H2020 AI4HEALTHSEC project, where it has also been successfully tested in real-world scenarios. Moreover, the datasets collected have been made publicly available on the SoBigData research infrastructure.

Keywords

Natural Language Processing, Large Language Models, Cyber Threats, Cyber Vulnerabilities, Impact Assessment, Cyber Risk Assessment

1. Introduction

The healthcare ecosystem is rapidly adopting a growing number of recent technologies, such as Internet of Things (IoT), wearable and implantable devices, Picture Archiving and Communication System (PACS), Electronic Health Records (EHRs), DiCOM images, and others, interconnected to realise and offer innovative healthcare digital services. While their adoption and use improve the quality of service to patients, and support and ease the work of the physicians and the medical professionals, on the other hand, this complex and dynamic inter-connection of several systems offers a larger attack surface for the threat actors interested in attacking the system by exploiting the existing vulnerabilities [1], also taking into account a low level of awareness of the cyber risks by the the healthcare personnel [2], often causing dramatic impacts to the healthcare ecosystem [3]. In example, a cyber-attack on a insecure PACS server

could lead to the web exposure of sensitive information of patients, or an attack to a remote monitoring software of a medical device could damage the equipment of the hospital or change the configuration of the device [4]. This sector has recently suffered several serious cyber attacks: for example, in 2017 and 2021 there were ransomware attacks on U.K. National Health System (NHS) and Ireland's Department of Health and Health Service Executive respectively [5]. Furthermore, inherent vulnerabilities have been found in some medical devices such as Braun's infusion pump and Medtronic's insulin pump [3]. Finally, approximately 90% of healthcare organisations experienced a data breach in 2018 [6]. For these reasons, it is necessary to study the most frequent attacks in healthcare to make the services offered more secure and resilient [4, 7]. Due to the complexity of the healthcare ecosystems, performing an effective cyber risk assessment can help to limit and prevent the cyber security incidents [8]. The cyber risk assessment process has the purpose of identifying, evaluating, and prioritising security risks to the assets of an organisation, allowing to perform the most appropriate action to mitigate the risks and the vulnerabilities.

Internet is a constantly updated source of threat, incident, and vulnerability-related information for healthcare ecosystem assets in the form of unstructured Natural Language (NL) within blogs, specialized Cyber-Security (CS) websites, social media, Knowledge Bases (KBs) and others. Although these sources contain crucial information about

Ital-IA 2024: 4th National Conference on Artificial Intelligence, organized by CINI, May 29-30, 2024, Naples, Italy

*Corresponding author.

✉ stefano.silvestri@icar.cnr.it (S. Silvestri);

giuseppe.tricomi@icar.cnr.it (G. Tricomi);

giuseppefelice.russo@icar.cnr.it (G. F. Russo);

mario.ciampi@icar.cnr.it (M. Ciampi)

🆔 0000-0002-9890-8409 (S. Silvestri); 0000-0003-3837-8730

(G. Tricomi); 0009-0001-2090-9647 (G. F. Russo);

0000-0002-7286-6212 (M. Ciampi)

© 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).



risk management and assessment, on the other hand, it is difficult to fully leverage them, due to the inherent complexity (polysemy, irony, long and complex sentences, non-standardized abbreviations, acronyms) of NL. Therefore, extracting relevant information from this mass of data becomes a demanding task [9]. The information extraction from NL text issues is currently addressed in literature adopting AI-based Natural Language Processing (NLP) models, usually implementing Named Entity Recognition (NER) systems [10, 11, 12, 13] using Large Language Models (LLMs) and CS KBs. However, there is a lack of focus in the literature on analyzing and prioritizing threats and vulnerabilities about the most frequent threats in healthcare. In this context, this paper extends the ideas previously presented in [14, 15, 16], combining NLP-based threat and vulnerability approaches to define an impact and risk assessment for the healthcare ecosystems, evaluating it by exploiting CS textual sources available on the Internet, presenting the final NLP cyber risk assessment methodology developed within the activities of the EC-funded H2020 AI4HEALTHSEC research project, as well as the collection of a textual CS dataset related to the “SoBigData.it” research project.

The paper is organized as follows: in Section 2, the most recent and related studies in the literature are outlined; subsequently, the details of the proposed approach are described in Section 3.5; afterwards, Section 4 shows the implementation of the proposed solution, a description of the datasets used and the research project where the approach was tested in real-world scenarios. Finally, Section 5 provides conclusions and future works.

2. Related Works

There are several recent works in the literature dealing with risk assessment and CS information extraction from NL documents. The authors of [8] reviewed and compared different generic cyber risk assessment frameworks in the healthcare field, comparing them, discussing the methodology of assessment and the limitations associated with them. A threat and mitigation model tailored for the IoT health devices is presented in [17], combining STRIDE and DREAD models: threats are identified using STRIDE model on the device access points, and then ranked using DREAD. This approach is suitable for both the designers and users of health IoT devices.

The security and privacy challenges in Medical Cyber-Physical Systems (MCPS) are discussed in [18], highlighting that trust and threat models usually consider MCPS stakeholders, including healthcare practitioners, system administrators and non-medical staff, with incorrect levels of trust. Also, in [2], the issues related to the CS awareness of the healthcare personnel are underlined, reviewing the existing gaps in CS strategies adopted by

healthcare organizations and the risk assessment methodologies adopted. The authors demonstrated that in this domain, there is often a lack of adequate training for healthcare workers and a lack of specialized figures, such as a chief information officer, highlighting the need to have security protocols updated to the latest standards.

Also, AI-based information extraction from CS textual documents has been recently developed and presented in the literature. In [13] is presented SecureBERT, a Bidirectional Encoder Representations from Transformers (BERT) model trained on CS-domain large NL corpora, which outperforms other similar models in NLP tasks in the CS domain. The authors of [10] collected a large corpus of labeled sequences from Industrial Control Systems device’s documentation to pre-train and fine-tune a BERT language model, named CyBERT. Also [12] proposed another interesting CS NER system, which exploits an architecture based on BERT, an LSTM, Iterated Dilated Convolutional Neural Networks (ID-CNNs), and Conditional Random Field, to improve the obtained performances.

The main innovation of the proposed approach is the use of CS information extracted from NL texts to calculate the threat, vulnerability, and impact levels, allowing the risk assessment for the various assets involved in digital healthcare services to be finally obtained.

3. Methodology

The proposed risk assessment methodology is composed of the following five steps: i) *Healthcare Ecosystem Assets Identification and Categorisation*; ii) *Threat Identification and Assessment*; iii) *Vulnerability Assessment*; iv) *Impact Assessment*; and v) *Risk Assessment*.

3.1. Healthcare Ecosystem Assets Identification and Categorisation

The preliminary step of the methodology provides a list of the assets of the considered digital complex healthcare system by identifying the corresponding services involved and their assets, with the final purpose of measuring their criticality within the healthcare system. For instance, the assets of a remote patient consultation service could include a Database, a Linux Server, communication software, and a web server. After their identification, the assets are also categorized, using the Common Platform Enumeration (CPE)¹ catalogue to map them with the corresponding area (based on their type) and category (depending on their functionalities), as shown in the next Table 1. This step allows us to understand the importance of each asset within the ecosystem and to provide a list of the assets that require risk assessment.

¹<https://nvd.nist.gov/products/cpe>

Table 1
Assets areas and categories.

Area	Name
1	User interactions with implants and sensors
2	Medical equipment and IT devices
3	Services and processes
4	Interdependent HCILs – Ecosystem
Category	Functionalities
Influence	Found in most organizations, distinct
Type	Software, hardware, Operating System (OS), Information Sensitivity
Sensitivity	Restricted, unrestricted
Criticality	Essential, required, deferrable

These classifications are used to evaluate the criticality of each asset of the healthcare system, by measuring the dependency level that an asset has with other system components. We defined our dependency levels:

- **Independent** assets have a distinct operation and exhibit no dependency on other assets. If the asset fails, no cascading events occur.
- **Incoming** dependency, if syntactically, another asset uses its data or functionality. If such an asset fails, the operation of all related assets that use its data or functionality may be disrupted.
- **Outgoing** dependency, if syntactically it uses data or functionality of another asset. Therefore, if the latter asset fails, the operation of the former asset will be affected as well.
- **Coupling relationship** reveals that two assets have both incoming and outgoing dependencies. Thereupon, failures in one of the assets will affect the functionality of the other.

Thus, the *criticality level* of an asset can be determined by the number of services and relevant business flows it participates in. Specifically, the General Asset Criticality level based on running services (GAC) is calculated as the weighted summation of their interdependencies, normalized by the total number of services in the examined healthcare ecosystem. Thereupon, the Asset Criticality for a specific service (ACS) is equal to its GAC value divided by the number of relevant/redundant assets that co-exist in the service. Finally, based on the ACS range values, it is possible to assign a criticality level to each asset, as shown in Table 2.

Table 2
Asset Criticality Levels.

ACS Value Range	Asset Criticality Level
[0,1]	Low
(1,2]	Medium
(2,3]	High

3.2. Threat Identification and Assessment

Once the assets have been identified, the next step aims to assess the threats that could affect those assets, following the approach previously described in [14, 15, 16]. Firstly,

a threat identification phase is performed by exploiting the Common Attack Pattern Enumeration and Classification (CAPEC)², which also provides a detailed set of the characteristics of the threats, such as *Likelihood of Attack*, *Related Attack Patterns*, *Execution Flow*, *Prerequisites* and others. In this way, we obtain the list of the threats for each asset that operates in the considered healthcare service/system (identified in the previous step). Each threat also includes the CAPEC ID, a CAPEC category that will be used to rate the threat, and the corresponding characteristics.

Then, it is possible to assess the threats, assigning them a severity level. Our methodology exploits the NL history of reported incidents related to those threats, extracted from large CS domain collections available online, such as forums, social media, news, and others, using an AI-based NLP approach. In detail, we use a Named Entity Recognition (NER) architecture based on SecureBERT [13], a BERT model pre-trained on a very large CS domain text collection (more than 2.2 million documents), preprocessed with a CS customized tokenizer, and fine-tuned for the NER task, to extract the mentions of the pairs threat and asset found in each sentence of the NL source. In this case, we produced a custom training set, annotated with the entity types of interest (*Asset*, and *Threat*) using the semi-supervised approach described in [19]. Then, the threat level is calculated based on the percentage of the occurrence of the mentions of that threat within the considered dataset, following the ranges shown in Table 3. The assessment is finally performed through a mapping between the assets of the services of the healthcare system and the pairs asset and threat with the corresponding threat level.

Table 3
Threat Levels and corresponding percentage of occurrence.

Threat Level	Occurrence Percentage	Description
Very High	[80-100]	Severe impact on critical services and assets
High	[60-80]	Significant impact on critical services and assets
Medium	[40-60]	Intermediate impact on services and assets and no critical service would be affected
Low	[20-40]	Low impact and no critical service would be affected
Very Low	[1-20]	Significant low impact

3.3. Vulnerability Assessment

The next step has the purpose of building a vulnerability exploit prediction scoring system specifically tailored for the healthcare domain. To this end, we adopted the NLP and Machine Learning (ML) approach described in [15], which leverages CS domain textual data sources to train a supervised ML classification model able to predict the

²<https://capec.mitre.org>

vulnerability score, obtaining in this way the vulnerability assessment. In summary, this method uses the textual data included in the CVE (the *Report* column of this KB) and the corresponding exploitability and impact metrics, namely the *attack vector*, *attack complexity*, *privileges required*, *user interaction*, *scope*, *confidentiality impact*, *integrity impact* and *availability*, to obtain a vector representation with the corresponding labels related to exploitability and impact metrics, used to train a set of ML XGBoost classifiers, which are able to predict the labels of the Attack Vector (*Network*, *Adjacent Network*, *Local*, *Physical*) and of the exploitability and impact metrics, summarised in the next Table 4.

Table 4
Exploitability and impact metrics and corresponding labels.

Exploitability and Impact metrics	Labels
Attack Complexity	Low, High
Privileges Required	None, Low, High
User Interaction	None, Required
Scope	Unchanged, Changed
Confidentiality	None, Low, High
Integrity	None, Low, High
Availability	None, Low, High

Then, an extension of CVE Exploit Prediction Scoring System (EPSS) is adopted [20], defining a Common Vulnerability Scoring System (CVSS)-like score using the labels predicted by the trained ML models on the NL texts, and following the specifications provided by [21]. The vulnerability level is based on the ranges of the computed CVSS-like score, as shown in Table 5.

Table 5
CVSS score ranges and corresponding vulnerability levels

CVSS-like Score Range	Vulnerability Level
8.0, 10	Very High
6.0, 8.0	High
4.0, 6.0	Medium
2.0, 4.0	Low
0.0, 2.0	Very Low

3.4. Impact Assessment

The next step of the proposed methodology is the Individual Impact Assessment, where the *impact level* is calculated to measure the effect that can be expected as the result of the successful exploitation of a vulnerability that resides in a critical asset. In this case, the methodology leverages the CVE KB used in conjunction with the same NER module used in the case of Threat Assessment fine-tuned to extract the assets and vulnerabilities entity types (see Section 3.2). This methodology exploits an additional set of adjectives related to the vulnerabilities and belonging to a predefined dictionary. These adjectives, such as *severe*, *serious*, *dangerous*, *etc.*, tend to indicate via a weight coefficient the severity level of the vulnerability. In detail, this dictionary is the result of the processed fea-

tures evaluated with two different classifiers that output scores to predict relevancy and severity, following the approach described in [22]. Each adjective is associated with a coefficient, calculated by taking through the log-odd ratio, then computing the exponential function on the log-odd, and converting odds to probability, using the formula: $probability = odds / (1 + odds)$. In this way, it is possible to associate the vulnerability to a scale *Low*, *Medium*, and *High*, where *Low* corresponds to [0, 33) (meaning that there is an 0-33% impact assessment probability), *Medium* corresponds to [33, 66), i.e., and *High* corresponds to [66, 100].

For vulnerabilities expressed in CVSS (obtained in the previous step), the three security criteria *Confidentiality* (*C*), *Integrity* (*I*), and *Availability* (*A*) are rated on a three-tier-scale: *None*, *Low*, and *High* (see previous Table 4). We can define a mapping from this three-tier scale onto a five-tier scale ranging from Very Low (*VL*) to Very High (*VH*) combining these characteristics, as shown in Table 6, providing in this way an initial impact level of a specific asset/vulnerability combination.

Then, the final impact level per asset is obtained by combining the initial impact with the asset criticality level (see Table 2), with the previous scale related to the adjectives and the corresponding vulnerabilities extracted by the NER module, as stated in next Table 7.

3.5. Risk Assessment

Finally, the Risk assessment is obtained by combining the Threat, Vulnerability, and Impact levels obtained in the previous steps, calculating the individual risk level for each asset following the next Table 8.

4. Implementation and Experiments

To implement the Threat and Impact assessment methods, we firstly needed a large and updated CS domain textual document collection. To this end, we collected the news published by The Hacker News website³, a CS news platform that attracts over 8 million readers monthly, which is daily updated with attacks, threats, vulnerabilities, and other CS news. A Python web crawler and scraper for this website has been specifically developed to retrieve, extract, collect, and normalise the text of each posted news. The scraping task is performed bi-weekly, making this dataset constantly updated also increasing its size. Moreover, this corpus is also made publicly on the SoBigData research infrastructure⁴. The NER module is based on SecureBERT [13], a BERT model pre-trained

³<https://thehackernews.com>

⁴Available at https://data.d4science.org/ctlg/ResourceCatalogue/the_hackernews_dataset

Table 6
Initial Impact Level calculation.

C \ I		None			Low			High		
		None	Low	High	None	Low	High	None	Low	High
None	VL	VL	L	L	L	M	M	M	H	
Low	VL	L	M	L	M	H	M	H	VH	
High	L	M	M	M	H	H	H	VH	VH	

Table 7
Final Impact Level calculation.

Asset Criticality	Low			Medium			High		
	Low	Medium	High	Low	Medium	High	Low	Medium	High
Initial Impact Level	Final Impact Level								
VL	VL	VL	L	VL	L	L	L	L	M
L	VL	L	M	L	M	M	L	M	H
M	L	L	M	M	M	M	M	M	H
H	L	M	M	M	M	H	M	H	H
VH	M	M	H	M	H	H	H	H	VH

on a very large CS domain text collection (more than 2.2 million documents), preprocessed with a CS customised tokenizer to improve its performance. This model has been fine-tuned for the NER task, to extract the mentions of the pairs of threat and asset found in each corpus sentence for the threat assessment, the mentions of vulnerabilities, the corresponding adjectives, and the assets for the impact assessment. To this end, we created two custom training sets, annotated with the entity types of interest (*Asset*, and *Threat* in the first case and *Asset*, *Vulnerability* and *Adjectives* in the latter case) using the semi-supervised approach described in [19]. The implementation of this module is based on the Huggingface Transformers Python library. The vulnerability assessment ML classifiers have been implemented using the Dmlc XGBoost library, a distributed gradient boosting library designed to be highly efficient and flexible.

The proposed methodology has been developed and implemented within the activities of the EC-funded H2020 project “AI4HEALTHSEC—A Dynamic and Self-Organised Artificial Swarm Intelligence Solution for Security and Privacy Threats in Healthcare ICT Infrastructures”. In this project, the proposed approach has been tested in real-world pilot scenarios provided by the Fraunhofer Institute for Biomedical Engineering (IBMT), a partner of the project. The pilots tested three different complex healthcare systems scenarios, namely Implantable Medical Devices, Wearables, and Biobank. The results of the tests, reported in [14, 15, 16], confirmed the effectiveness and the applicability of our method.

5. Conclusion and Future Works

The paper proposes an AI-based approach for the individual risk assessment of the assets of digital healthcare systems. The approach, after the classification of the criticality of the assets using CS KBs, leverages NER and ML systems to extract and classify relevant information from textual CS sources, allowing to calculate the threat, vulnerability and impact levels, which are finally combined to obtain the risk level of each asset. The methodology was successfully tested in real-world pilot scenarios of the EC-funded H2020 AI4HEALTHSEC project, demonstrating its applicability and effectiveness. Moreover, the datasets, which are constantly updated, are made publicly available on the SoBigData research infrastructure.

Acknowledgments

This work is supported by the European Union—NextGenerationEU—National Recovery and Resilience Plan (Piano Nazionale di Ripresa e Resilienza, PNRR)—Project: “SoBigData.it—Strengthening the Italian RI for Social Mining and Big Data Analytics”—Prot. IR0000013—Avviso n. 3264 del 28/12/2021.

We thank Simona Sada and Giuseppe Trerotola for the administrative and technical support provided.

Table 8
Individual Risk Level calculation.

Threat Vulnerability Impact	Very Low					Low					Medium					High					Very High				
	VL	L	M	H	VH	VL	L	M	H	VH	VL	L	M	H	VH	VL	L	M	H	VH	VL	L	M	H	VH
VL	VL	VL	L	L	L	VL	L	L	L	M	VL	L	M	M	M	L	L	M	M	M	L	L	M	M	M
L	VL	L	L	L	M	L	L	L	M	M	L	L	M	M	M	L	M	M	M	H	L	M	M	H	H
M	L	L	L	M	M	L	L	M	M	M	L	M	M	M	H	M	M	M	H	H	M	M	M	H	H
H	L	L	M	M	M	L	M	M	M	H	M	M	M	H	H	M	M	H	H	H	M	M	H	H	H
VH	L	M	M	M	H	M	M	M	H	H	M	M	H	H	H	M	H	H	H	VH	M	H	H	VH	VH

References

- [1] P. Ribino, M. Ciampi, S. Islam, S. Papastergiou, Swarm intelligence model for securing healthcare ecosystem, *Procedia Computer Science* 210 (2022) 149–156. doi:<https://doi.org/10.1016/j.procs.2022.10.131>.
- [2] S. Nifakos, K. Chandramouli, C. K. Nikolaou, P. Papachristou, S. Koch, E. Panaousis, S. Bonacina, Influence of human factors on cyber security within healthcare organisations: A systematic review, *Sensors* 21 (2021). doi:[10.3390/s21155119](https://doi.org/10.3390/s21155119).
- [3] D. McKee, P. Laulheret, McAfee Enterprise ATR uncovers vulnerabilities in globally used B. Braun infusion pump, 2021. URL: <https://www.trellix.com/blogs/research/mcafee-enterprise-atr-uncovers-vulnerabilities-in-globally-used-b-braun-infusion-pump/>.
- [4] S. Islam, S. Papastergiou, H. Mouratidis, A dynamic cyber security situational awareness framework for healthcare ICT infrastructures, in: *Proceedings of the 25th Pan-Hellenic Conference on Informatics, PCI '21*, ACM, Volos, Greece, 2022, p. 334–339. doi:[10.1145/3503823.3503885](https://doi.org/10.1145/3503823.3503885).
- [5] D. Rees, Cyber attacks in healthcare: the position across europe, 2021. URL: <https://www.pinsentmasons.com/out-law/analysis/cyber-attacks-healthcare-europe>.
- [6] Sixth annual benchmark study on privacy & security of healthcare data, 2016. Ponemon Institute.
- [7] K. S. Bhosale, M. Nenova, G. Iliev, A study of cyber attacks: In the healthcare sector, in: *2021 Sixth Junior Conference on Lighting (Lighting)*, 2021, pp. 1–6. doi:[10.1109/Lighting49406.2021.9598947](https://doi.org/10.1109/Lighting49406.2021.9598947).
- [8] S. Memon, S. Memon, L. Das, B. R. Memon, Cyber security risk assessment methods for smart healthcare, in: *2024 IEEE 1st Karachi Section Humanitarian Technology Conference (KHI-HTC)*, 2024, pp. 1–6. doi:[10.1109/KHI-HTC60760.2024.10481961](https://doi.org/10.1109/KHI-HTC60760.2024.10481961).
- [9] M. Tikhomirov, N. Loukachevitch, A. Sirotina, B. Dobrov, Using BERT and augmentation in named entity recognition for cybersecurity domain, in: *25th International Conference on Applications of Natural Language Processing and Information Systems*, Springer, Saarbrücken, Germany, 2020, pp. 16–24.
- [10] K. Ameri, M. Hempel, H. Sharif, J. Lopez Jr., K. Perumalla, Cybert: Cybersecurity claim classification by fine-tuning the bert language model, *Journal of Cybersecurity and Privacy* 1 (2021) 615–637. URL: <https://www.mdpi.com/2624-800X/1/4/31>. doi:[10.3390/jcp1040031](https://doi.org/10.3390/jcp1040031).
- [11] S. Zhou, J. Liu, X. Zhong, W. Zhao, Named entity recognition using bert with whole world masking in cybersecurity domain, in: *2021 IEEE 6th International Conference on Big Data Analytics (ICBDA)*, volume 26, IEEE, Xiamen, China, 2021, pp. 316–320. doi:[10.1109/ICBDA51983.2021.9403180](https://doi.org/10.1109/ICBDA51983.2021.9403180).
- [12] Y. Chen, J. Ding, D. Li, Z. Chen, Joint bert model based cybersecurity named entity recognition, in: *2021 The 4th International Conference on Software Engineering and Information Management, ICSIM*, Yokohama, Japan, 2021, pp. 236–242. doi:[10.1145/3451471.3451508](https://doi.org/10.1145/3451471.3451508).
- [13] E. Aghaei, X. Niu, W. Shadid, E. Al-Shaer, SecureBERT: A domain-specific language model for cybersecurity, in: *Security and Privacy in Communication Networks*, Springer, Cham, 2023, pp. 39–56.
- [14] S. Islam, S. Papastergiou, S. Silvestri, Cyber threat analysis using natural language processing for a secure healthcare system, in: *2022 IEEE Symposium on Computers and Communications (ISCC)*, 2022, pp. 1–7. doi:[10.1109/ISCC55528.2022.9912768](https://doi.org/10.1109/ISCC55528.2022.9912768).
- [15] S. Silvestri, S. Islam, S. Papastergiou, C. Tzagkarakis, M. Ciampi, A machine learning approach for the nlp-based analysis of cyber threats and vulnerabilities of the healthcare ecosystem, *Sensors* 23 (2023). doi:[10.3390/s23020651](https://doi.org/10.3390/s23020651).
- [16] S. Silvestri, S. Islam, D. Amelin, G. Weiler, S. Papastergiou, M. Ciampi, Cyber threat assessment and management for securing healthcare ecosystems using natural language processing, *International Journal of Information Security* 23 (2024) 31–50. doi:[10.1007/s10207-023-00769-w](https://doi.org/10.1007/s10207-023-00769-w).
- [17] A. Omotosho, B. A. Haruna, O. M. Olaniyi, Threat modeling of internet of things health devices, *Journal of Applied Security Research* 14 (2019) 106–121. doi:[10.1080/19361610.2019.1545278](https://doi.org/10.1080/19361610.2019.1545278).
- [18] H. Almohri, L. Cheng, D. Yao, H. Alemzadeh, On threat modeling and mitigation of medical cyber-physical systems, in: *2017 IEEE/ACM International Conference on Connected Health: Applications, Systems and Engineering Technologies (CHASE)*, 2017, pp. 114–119. doi:[10.1109/CHASE.2017.69](https://doi.org/10.1109/CHASE.2017.69).
- [19] G. Aracri, A. Folino, S. Silvestri, Integrated use of KOS and deep learning for data set annotation in tourism domain, *Journal of Documentation* 79 (2023) 1440–1458. doi:[10.1108/JD-02-2023-0019](https://doi.org/10.1108/JD-02-2023-0019).
- [20] J. Jacobs, S. Romanosky, B. Edwards, I. Adjerid, M. Roytman, Exploit prediction scoring system (EPSS), *Digital Threats* 2 (2021). doi:[10.1145/3436242](https://doi.org/10.1145/3436242).
- [21] A.A.V.V., Common Vulnerability Scoring System version 3.1 Specification Document, Technical Report, FIRST.Org, 2019. URL: https://www.first.org/cvss/v3-1/cvss-v31-specification_r1.pdf.
- [22] L. Breiman, Random forests, *Machine learning* 45 (2001) 5–32.