Lightweight machine learning for privacy-preserving and secure networked medical devices: The SEPTON project use cases

Sotirios C. Messinis
smessinis@mail.ntua.gr
Institute of Communications and Computer Systems
Athens, Greece
National Technical University of Athens
Athens, Greece

Nicholas E. Protonotarios
np558@cam.ac.uk
University of Cambridge
Cambridge, UK
Institute of Communications and Computer Systems
Athens, Greece

Ioannis N. Tzortzis
itzortzis@mail.ntua.gr
National Technical University of Athens
Athens, Greece

Ioannis Rallis
irallis@central.ntua.gr
National Technical University of Athens
Athens, Greece

Dimitrios Kalogeras
d.kalogeras@noc.ntua.gr
Institute of Communications and Computer Systems
Athens, Greece

Nikolaos Doulamis
ndoulam@cs.ntua.gr
National Technical University of Athens
Athens, Greece

ABSTRACT

Cybersecurity incidents are among the greatest concerns of businesses, government agencies, and private citizens today. In the modern world, the protection of data and information assets has become nearly as important as maintaining the security of physical assets. This creates the need for increased security implementations, leading to improved user acceptance of such applications and, as a consequence, to large-scale adoption of these technologies and full exploitation of their advantages. In healthcare, networked medical devices (NMDs), either referring to hospital medical equipment or wearables, can be vulnerable to security breaches, potentially affecting the safety and effectiveness of each device. In this work, we present the specific areas of recent machine learning research applied to networked medical device security, through the objectives of the Horizon Europe SEPTON research project. State-of-the-art lightweight machine learning approaches are highlighted and the corresponding challenges of cybersecurity applications, ranging from implantable devices to inter-institution medical data exchange use cases, are showcased.

CCS CONCEPTS

• **Security and privacy** → **Systems security**; • **Computing methodologies** → **Machine learning**.

KEYWORDS

neural networks, anomaly detection, lightweight machine learning, networked medical devices

PETRA '23, July 05–07, 2023, Corfu, Greece

ACM ISBN 979-8-4007-0069-9/23/07.
https://doi.org/10.1145/3594806.3596562

ACM Reference Format:
Sotirios C. Messinis, Nicholas E. Protonotarios, Ioannis N. Tzortzis, Ioannis Rallis, Dimitrios Kalogeras, and Nikolaos Doulamis. 2023. Lightweight machine learning for privacy-preserving and secure networked medical devices: The SEPTON project use cases. In *Proceedings of the 16th International Conference on PErvasive Technologies Related to Assistive Environments (PETRA '23), July 05–07, 2023, Corfu, Greece.* ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/3594806.3596562

1 INTRODUCTION

Networked medical devices (NMDs) can be quite complicated and multifaceted; their applications frequently focus on interoperability with existing hospital assets, security and privacy of sensitive data, and end-user usability [27]. The introduction of artificial intelligence into hospitals has led to greater safety and efficiency for patients. Security from attacks is a key issue for critical infrastructures, with several assets being affected, including patients' lives, sensitive personal data, and financial resources. As the number of attacks increases with the introduction of expanded connectivity, the negative potential grows exponentially. This increases the opportunities for cyber attackers to manipulate any networked medical device, from MRI (Magnetic Resonance Imaging) scanners to electric wheelchairs [14].

In these dynamic and security-intensive environments, where proper device functioning must be ensured, device vulnerability and security measures must be properly assessed. Since modern wearable and implantable devices (IMDs) are equipped with wireless transceivers that enable wireless data exchange with external readers, the associated wireless communication may lead to unauthorized device access or even personal data leaks, thus rendering the implant unusable. Due to the fact that modern IMDs are essentially embedded computer systems, specially crafted computer viruses, and malware may infect implantable devices, potentially harming a large patient base. Not only does wireless connectivity transcends the functionality of wearables, but it also makes them visible to the entire healthcare system, increasing the attack surface [6].

Securing NMDs presents unique challenges that may not be efficiently addressed by traditional security methods [24]. One of these challenges corresponds to the low computational and storage capacity of the corresponding sensors, which makes it difficult to apply resource-intensive security measures. A "secure-by-design" system must address the need of stakeholders for rapid access during vital emergencies. Furthermore, the ability to control equipment through external devices increases the risk of security breaches, and necessitates the implementation of lightweight security mechanisms. The effective use of these mechanisms must be ensured through near-optimal anomaly detection and privacy measures, which are expected to be enhanced by recent advances in machine learning [24, 31].

This article is structured as follows: in section 2, we present related work to lightweight machine learning applied to the cybersecurity of NMDs In section 3, we showcase the concept of the SEPTON project, along with its main objectives and contributions. Furthermore, in section 4, we present the development of a lightweight machine learning concept for anomaly detection in NMDs, whereas in section 5 we describe the lightweight cryptographic security primitives, focusing on a special category of encryption algorithms deployed to implantable and wearable devices. In section 6, several aspects of the theory of differential privacy are briefly presented, especially as they relate to medical devices, and in section 7 four use cases of the SEPTON project are outlined.

2 RELATED WORK

A growing stream of research on lightweight security has emerged in the recent literature, focusing on various aspects, ranging from anomaly detection in the Internet of Things (IoT) to distributed cryptographic and security primitives deployed on various edge devices. In this paper, we concentrate on the implementation of lightweight security solutions in networked medical devices through specific use cases. The suggested solutions rely on a dual nature contribution, divided into lightweight machine learning for anomaly detection and the implementation of machine learning to lightweight security primitives, as they derive from the cryptography literature. In this direction, the authors of [3] simulated the differentials for non-Markov ciphers using machine learning, reducing a differentiation problem to an efficient-to-handle classification problem with implementation on medical devices. The result of their study was the first proof-of-concept on how machine learning can be used as a general tool in the cryptanalysis of symmetric keys.

Another related article [12] presents an experimental study of cryptographic encryption algorithms for their classification as either asymmetric or symmetric. The corresponding algorithms were tested in independent computing devices and, considering their subdivision into symmetric and asymmetric, it was proven that even if the symmetric algorithms were faster, they were not as secure as asymmetric, utilizing a pair of public or private keys. In [37], the importance of implementing the right security measures for the enhancement of medical device security against cyber-attacks is presented. The review of existing solutions was classified based on their cryptographic nature, further analyzed, and compared in computational complexity terms.

Appropriate lightweight cryptographic algorithms and protocols have been broadly highlighted for their capability of limited computational costs and resource utilization. However, the above literature does not seem to emphasize the integration of lightweight machine learning mechanisms in advanced cryptography solutions, especially in specific medical applications, varying from implantable devices to networked hospital infrastructures, and inter-institutional medical data exchange. In [21], the authors present an advanced intrusion detection system based on machine learning, in order to protect shared information in critical networked infrastructures. The approach adopted is split learning, providing higher accuracy, better detection, and classification performance. In [25], a joint blockchain and federated learning secure architecture was developed with primary applications in smart healthcare and smart city settings.

The federated learning components are used for the scalability of machine learning applications ensuring that personal data will be kept safe at the edge. In this context, fast learning schemes have been proposed in the literature, including [7] and [20]. Finally, in [16], the authors proved that privacy-preserving federated learning, based on the cryptographic primitive of homomorphic re-encryption, can efficiently protect user data while training through batch gradient descent (BGD).

3 FRAMEWORK FOR IMMUTABLE IT INFRASTRUCTURE IN HOSPITALS

The SEPTON (SEcurity Protection TOols for Networked medical devices) project is expected to contribute to the cyber-security for IT medical infrastructure by providing immutable end-point solutions. The main action pillar of SEPTON aims to identify the need for scaling and adoption of solutions at the national, regional, or local level through early engagement with patients, healthcare providers, healthcare authorities, and regulators. Specifically, SEPTON aims to develop tools to help healthcare facilities better prepare for and respond to cyber threats to their medical device infrastructure. In this direction, healthcare providers can remain resilient during and after emergencies and minimize business disruptions that affect their ability to provide critical services to the public. By improving these services, protection for networked medical devices, health care, and public health services will be established when a threat occurs.

The project focuses on protecting networked medical devices and associated data exchange. In modern healthcare, NMDs are the points at which health data is generated. This is true for implantable and wearable devices, such as pacemakers, as well as for traditional devices, such as MRI scanners. The individual tools, as well as the SEPTON toolkit itself, provide the foundation for improving the cybersecurity of such devices and reducing the risk of device compromise. The toolkit's inherent flexibility and scalability allow it to be used in a variety of environments, from protecting a few devices in a body area network to monitoring and protecting the assets (devices and network resources) of entire hospitals. The knowledge and experience gained through the project will feed into stakeholder initiatives, which include national, regional, and sectoral bodies, as well as the frameworks created by these bodies.

This way, the impact of SEPTON will be disseminated as widely as possible, and the time for project results to be incorporated into regulations and standards will be shortened. In particular, the project aims to comply with all the General Data Protection Regulation (GDPR) requirements and will follow the strict guidelines of MDCG 2019-16 cybersecurity protocol [10], as well as the post-market regulatory actions associated with it, see Fig. 1. To maximize benefits to society, the project's activities will be extended beyond traditional technical advances. To this end, issues relevant to the acceptance of the project by society and citizens will be appropriately evaluated and the nature of the project will be adapted to the ethical constraints established. Ultimately, the SEPTON project will focus on (a) lightweight and federated anomaly detection, (b) lightweight machine learning, (c) hardware acceleration, (d) lightweight cryptographic security primitives, (e) differential privacy.

4 LIGHTWEIGHT AND FEDERATED ANOMALY DETECTION

Anomaly detection involves discovering events that deviate from typical patterns or behaviors, thus indicating potential problems [22]. There exist numerous methods to detect anomalies and classify biological signals. Among the biological signals studied are electrocardiogram-based (ECG) arrhythmias, where the required classification is a pattern recognition problem that can be solved via machine learning algorithms [26]. The extensive IoT deployment in several verticals has substantially extended the limits of anomaly detection mechanisms. Medical devices are among the top innovative application environments that benefit from the IoT evolution.

To this end, lightweight mechanisms combined with distributed learning paradigms prevail in distributed machine learning schemes that facilitate security and privacy-preserving mechanisms [28]. Similarly, federated anomaly detection combines machine learning techniques with federated learning [18]. It involves training anomaly detection models on distributed data sources, without centralizing the data. In traditional anomaly detection, a model is trained on a dataset that contains both normal and anomalous data. This allows the model to learn to distinguish between the two. However, in certain cases, centralizing data in one location for training usually raises privacy and security concerns.

4.1 Lightweight machine learning

Malicious traffic identification using deep learning techniques has become a crucial aspect of anomaly detection research [1]. Currently, the necessity for the application of lightweight machine learning models to an increasing number of edge devices has led to several algorithmic and architectural developments. These include techniques including model compression, quantization, and pruning, which reduce the size and complexity of the model while maintaining its performance. Medical devices, such as implantable devices or wearable sensors, generate large amounts of data, which can be used to identify anomalies and monitor patient health. Lightweight machine learning models can be trained on these data to recognize patterns and identify anomalies in real-time, allowing healthcare professionals to intervene quickly, if necessary.

One advantage of using lightweight machine learning models for anomaly detection in medical devices is that they can be optimized for low-power consumption and run efficiently on resource-constrained devices. This is particularly important for medical devices that may need to operate for extended periods of time on limited battery power [5]. Nevertheless, machine learning models for medical applications must be rigorously tested and validated to assure their safety and efficacy. Concerns about privacy and security must also be addressed in order to secure patient data. Among several deployments and solutions proposed, recurrent neural networks [30] and convolutional neural networks [32] have been widely explored, yielding promising results.

Federated learning can also be proven to be a quite useful approach for anomaly detection in medical devices [34]. This approach can help address privacy and security concerns associated with sharing sensitive medical data. In the context of medical devices, federated learning can be used to train lightweight machine learning models on data generated by a large number of devices. Each device can contribute its own data to the training process, allowing the model to learn from a diverse range of patient populations and clinical scenarios. The model can then be used to identify anomalies in real-time, without the need to transmit sensitive patient data to a central server [2]. Federated learning can also help address challenges related to data heterogeneity, as medical devices may generate data with varying characteristics and formats. By training models on a diverse range of devices, the model can learn to identify anomalies across different types of data generated by different devices.

4.2 Hardware acceleration

Hardware acceleration refers to the use of specialized hardware components in order to speed up certain types of computational tasks. A wide range of novel general-purpose architectures and devices offer certain advantages over traditional processing units. Despite this development, the performance and resource limitations of these platforms pose a significant obstacle to accelerating high-risk medical applications. Recent improvements in deep learning algorithms offer interesting application opportunities for their use in safety-critical biomedical and medical applications [29].

To this end, hardware acceleration can be used to speed up training and inference processes, allowing models to be trained and deployed more quickly and efficiently while enabling privacy-preserving and anomaly-detection mechanisms. Some examples of hardware acceleration include GPUs (graphics processing units) and TPUs (tensor processing units) [15]. Modern wearable health devices require new resource-efficient technologies with high performance, low energy consumption, and high accuracy [19]. Convolutional neural networks (CNNs) have demonstrated high efficiency towards distributed AI functionalities in medical devices, whereas the field-programmable gate arrays (FPGAs) have been extensively used to construct hardware accelerators for CNNs [13].

For applications where the insatiable demands of deep learning methods for computational power and large training data cannot be satisfied, alternative machine learning approaches are necessary; these approaches must have extremely low latency and must be able to work with only a small training data set. Certain AI accelerators,

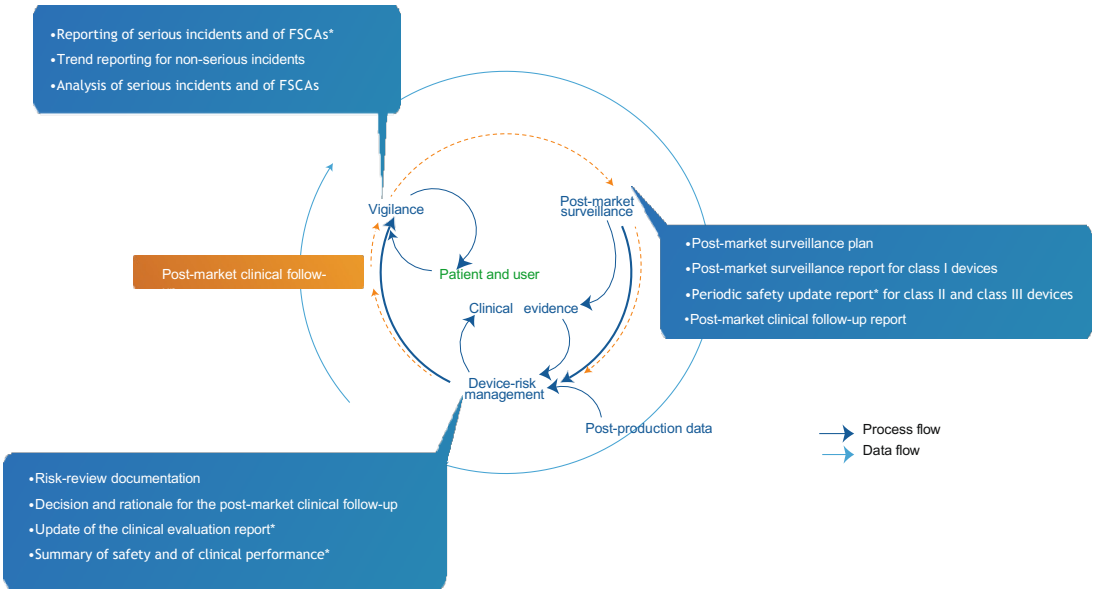

Figure 1: The patient-centric and user-centric approaches of the post-market regulatory cycle for a high-risk medical device (MDCG 2019-16 protocol). [4]

including several optical ones, enable classification on short time scales for fast optical imaging, sensing, and metrology without increasing the data size for further enabling the reduction of the latency referring to the nonlinear classification of certain data by several orders of magnitude [39].

5 LIGHTWEIGHT CRYPTOGRAPHIC SECURITY PRIMITIVES

Lightweight cryptographic security primitives are algorithms or protocols designed to provide security in resource-constrained environments, such as low-power embedded devices or sensors, where traditional cryptographic algorithms may not be feasible due to their high computational and memory requirements [9]. These primitives are typically optimized for low computational overhead and small memory requirements, making them suitable for devices with limited computational power and memory. In the relevant literature, the ciphers are groups of encryption algorithms that work with fixed-length blocks of data. Due to their specific nature, lightweight ciphers are optimized for low power consumption, low latency, and small code size. Hash functions, in turn, are mathematical functions that convert data into a fixed-size output called a hash.

Lightweight hash functions are optimized for low power consumption and a small memory footprint. In addition, the lightweight version of the popular public-key cryptography refers to a cryptographic method that uses a public key and a private key for encryption and decryption [35]. Until now, typical lightweight cryptographic primitives and protocols seem to have been vulnerable to quantum attacks. Due to its resistance to quantum assaults, advanced technologies such as post-quantum cryptography are vital for the IoT [17]. Furthermore, conventional cryptanalysis of modern ciphers can be impractical or demonstrate apparent limitations. Deep learning techniques are becoming more and more successful against potential attacks, verifying the security of emerging lightweight ciphers, while classifying their complex differences.

6 DIFFERENTIAL PRIVACY

Differential privacy (DP) is a technique used to protect the privacy of individuals in large datasets by adding random noise to the data, in such a way that statistical analysis of the data remains accurate without revealing the identity of individual data points. This technique is particularly useful in healthcare and medical devices, where sensitive information about individuals' health and medical conditions is often collected and analyzed. Medical devices such as wearable fitness trackers, glucose monitors, and heart monitors collect data about an individual's health in real time. DP can be used to protect this data from being accessed by unauthorized third parties [8]. In addition to protecting the privacy of individuals, DP can also be used to improve the accuracy of medical research by allowing researchers to access large datasets without compromising the privacy of the individuals within those datasets.

This can lead to more accurate research results and better treatment options for patients. Privacy-preserving mechanisms for physiological signals collected by intelligent wearable devices can be also enhanced by the use of DP. For example, in the data publishing stage, new models based on the combination and optimization of k-anonymity and differential privacy are presented in [11]. Furthermore, temporal differential privacy mechanisms are successfully deployed for real-time data to suppress privacy leakage while updating data. To this end, the vanilla differentially private stochastic gradient descent (DP-SGD) shapes the way for the development of more advanced differentially private algorithms deployed in wearable and IoT devices [23]. Apart from its combination with other very promising technologies, DP can also be used to make the machine learning algorithms differentially private themselves shaping the way for emerging applications in the state-of-the-art quantum machine learning [33].

7 THE SEPTON APPLICATION SCENARIOS

The EU-funded SEPTON project, under Horizon Europe, will develop an advanced cybersecurity toolkit, along with a connected detailed dashboard referring to all the anomaly detection and security measurement aspects of modern medical devices and infrastructures. In this section, we present the 4 distinctive medical device use cases of the SEPTON project highlighting the potential of the aforementioned cybersecurity technologies in combination with applied lightweight machine learning in different scale demonstrations. For details, see Fig. 2.

7.1 Implantable devices

In the special case of implanting a deep-brain neurostimulator device, with proper placement of the monitoring and stimulating electrodes, the implant is capable of suppressing seizures or mitigating their impact by electrical stimulation in the cortex [36]. Along with the implant, at least one reader device, communicating with the implant wirelessly through a proprietary protocol, is usually present. In modern versions of the system, the reader is either an Android or an Apple app installed on a smartphone. With this reader, a patient or doctor is able to (a) monitor the progress, occurrence rate, and other aspects of the patient's seizures, and (b) modify the stimulation program delivered by the implant to better suit his needs and daily life. Except for the patient reader, a bedside reader that resides always at the patient's home could also be tasked with communicating with the implant to dump data logs and forward them to a proprietary healthcare network and cloud.

The implant and the corresponding reader are required to employ certain security primitives in order to protect private data communicated between the two entities. Blocking access to serve security may be as hazardous as permitting indiscriminate access to serve safety. Properly discriminating between the two situations is a hard problem to solve that epitomizes the security challenge in modern implantable devices and becomes only more exacerbated when more devices are gradually being exposed to the broader healthcare ecosystem. Dynamic biometrics is a feature with the benefit of being cheap to access on the implantable device side (in vivo), and of being capable of high entropy, thus enabling the construction of lightweight security primitives.

Traditionally, these security primitives and protocols consume a lot of battery power. However, this is not suitable for modern implantable devices, since these devices are typically designed to operate for up to a decade or so while implanted in the human body. Therefore, lightweight machine learning mechanisms can be specifically tailored to the security requirements of implantable devices considering the battery consumption requirements.

7.2 Wearable devices

A very common medical wearable device is an ECG recording. Most patients with cardiovascular diseases exhibit cardiac arrhythmia which is one of the most common problems in cardiology. Cardiac arrhythmias are irregular heartbeats caused by the improper functioning of the heart. Some pathological conditions can be diagnosed at an early stage using ECG recordings, which may lead to better outcomes and save lives. Such long-term ECG monitoring is a tedious task as it generates huge amounts of data that has to be analyzed by well-trained medical professionals. Therefore, there is a need for recording devices should be portable or wearable in order to improve the efficiency of the diagnosis process.

The portable (or wearable) ECG device can easily have its battery recharged or replaced; therefore, power budgets can also be somewhat higher. With regards to implantable devices, this implies opting for lightweight security on the device, though the frequent uplink of data over the Internet is expected to be the major contributor to battery depletion that needs to be addressed.

7.3 Networked hospital equipment

When a new MRI scanner is installed at the radiology department of a hospital, a PACS (Picture Archiving and Communication System) server, along with an associated medical imaging client software are usually installed as well. Several existing modalities will communicate with the new PACS in order to store their imagery and the new client software will be installed onto several computers to allow the doctors to consult on the medical exams of the patients. The IT system administrator must ensure that the newly acquired medical equipment, as well as the accompanying server and associated software, will not pose any security threat to the existing systems, while the assets need to be protected from possible unauthorized access or malicious attacks, both external and internal.

Currently, there are no known tools to detect vulnerabilities in NMDs, other than the conventional ones used by network administrators and security specialists. Legacy operating systems and software and the incompatibility between systems leave vulnerabilities such as misconfiguration and security holes. These include vulnerabilities from non-negotiated interfaces with third-party software, often through web interfaces. However, devices may not be fully owned and operated by the care center (usually on lease from a system provider). Furthermore, a device that was given to a patient as part of a treatment plan will hence re-enter the hospital's realm at a later point and may intentionally or unintentionally cause issues within the care center or when given out to the next patient. This means either that the care center's own systems are at risk or that a device provided by the hospital threatens the client.

With developments such as polymorphic malware, current methods for detecting intrusions at the end-point and in the network

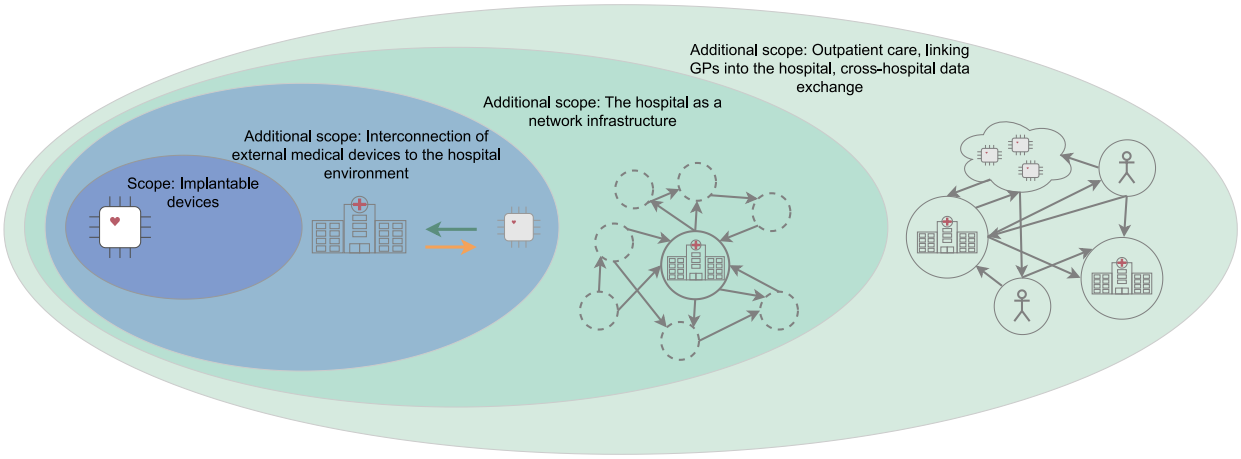

Figure 2: SEPTON project use cases: scope and scale.

are being questioned. The number of new threats is so large, that scanners now keep a time window and have to discard threats more than two months old. Existing machine-learning approaches within intrusion detection systems that find a hard boundary between samples do not deal with these situations reliably. However, current advances in deep clustering approaches, such as density-based clustering to network traffic analysis, are able to cope with slight variations in user behavior yet detect outliers from the norm. The expected goal is a universal toolkit that learns patterns by itself by observing network traffic. Based on the assumption that the entire network is not infected entirely with a threat at the same time and that there exists an exact time when a threat is introduced to the system, this also helps to automatically derive a secure baseline.

7.4 Medical data exchange between institutions

When a patient needs a referral to a specialist, there exist different procedures, depending on the patient's choice of hospital. A hospital might subscribe to a secure portal service that is integrated with both the hospital information system (HIS) and the GP information system (GIS). The specialist who makes a diagnosis and suggests a treatment sees the patient. This information can be sent back to the patient and can be received from both the GIS and the pharmacy information system (PIS), both of which are installed on the GP's computer. Another hospital may have subscribed to a different service. In this case, it must create the referral letter in a Web environment by manually inserting the medical information from GIS into the browser.

Another facility may even use an outdated system that is incompatible with the rest of the hospital system. As a result, it may require that all referrals be sent by fax. A major challenge is ensuring secure operations and data protection via a mix of third-party software and online services, even fax machines in some cases. Due to incompatibilities and inconsistencies between the various moving parts and the overall manual transmission of information, not necessarily fully understood even by physicians, pharmacists, and other support staff, the reliability of the system is at risk while existing security vulnerabilities can easily go undetected. Sharing of medical data among different healthcare institutions allows for the creation of a large dataset for healthcare analysis that can be used for effective decision-making regarding diseases and treatment plans, as well as for the production of general health statistics.

To enable privacy-friendly medical data exchange and privacy-friendly analysis of medical data exchanged between the aforementioned parties, we introduce a differential privacy-based approach to the data exchange process, theoretically guaranteeing privacy and limiting information loss for sensitive data. We develop a system model that includes users, cloud servers, and healthcare facilities. In our system, interactive and non-interactive DP approaches are compared. Homomorphic encryption methods are used to compute the statistical functions of encrypted data. Although performing fully homomorphic encryption (FHE) is very costly, somewhat homomorphic encryption (SHE), which performs only a limited number of multiplications on encrypted data, is more efficient than FHE [38]. However, the computational cost of SHE is high and should be implemented with efficient algorithms. By combining lightweight machine learning with cryptographic and differential privacy techniques, we will develop private tools that allow medical professionals and scientists to share useful statistical information about sensitive data.

8 CONCLUSION

This work concentrates on the potential of integration and effective implementation of advanced lightweight machine learning algorithms to the security aspects of networked modern medical devices. Beginning with an overview of the relevant literature we emphasized specific medical cybersecurity use cases, addressed by the SEPTON EU Horizon research program, and how the aforementioned state-of-the-art technologies affect them referring to their scale. The ongoing effectiveness of lightweight machine learning architectures is presented with promising results in these critical healthcare settings. In the near future, we plan to share the main scientific measurements derived from the use cases investigated.

ACKNOWLEDGMENTS

This research is supported by the Horizon Europe project "Security Protection Tools For Networked Medical Devices (SEPTON)" funded under grant agreement No. 101094901.

REFERENCES

[1] Adel Abusitta, Glaucio H.S. de Carvalho, Omar Abdel Wahab, Talal Halabi, Benjamin C.M. Fung, and Saja Al Mamoori. 2023. Deep learning-enabled anomaly detection for IoT systems. *Internet of Things* 21 (2023), 100656. https://doi.org/10.1016/j.iot.2022.100656

[2] Mohammed Adnan, Shivam Kalra, Jesse C Cresswell, Graham W Taylor, and Hamid R Tizhoosh. 2022. Federated learning and differential privacy for medical image analysis. *Scientific Reports* 12, 1 (2022), 1953.

[3] Anubhab Baksi. 2022. *Machine Learning-Assisted Differential Distinguishers for Lightweight Ciphers*. Springer Singapore, Singapore, 141–162. https://doi.org/10.1007/978-981-16-6522-6_6

[4] Shiko Ben-Menahem, Raymond Nistor, Gloria Macia, Georg Krogh, and Joerg Goldhahn. 2020. How the new European regulation on medical devices will affect innovation. *Nature Biomedical Engineering* 4 (03 2020). https://doi.org/10.1038/s41551-020-0541-x

[5] John Carter, Spiros Mancoridis, and Erick Galinkin. 2022. Fast, Lightweight IoT Anomaly Detection Using Feature Pruning and PCA. In *Proceedings of the 37th ACM/SIGAPP Symposium on Applied Computing* (Virtual Event) *(SAC '22)*. Association for Computing Machinery, New York, NY, USA, 133–138. https://doi.org/10.1145/3477314.3508377

[6] Liezel Cilliers. 2020. Wearable devices in healthcare: Privacy and information security issues. *Health information management journal* 49, 2-3 (2020), 150–156.

[7] Nikolaos Doulamis and Athanasios Voulodimos. 2016. FAST-MDL: Fast Adaptive Supervised Training of multi-layered deep learning models for consistent object tracking and classification. In *2016 IEEE International Conference on Imaging Systems and Techniques (IST)*. IEEE, New York, NY, USA, 318–323. https://doi.org/10.1109/IST.2016.7738244

[8] Amalie Dyda, Michael Purcell, Stephanie Curtis, Emma Field, Priyanka Pillai, Kieran Ricardo, Haotian Weng, Jessica C Moore, Michael Hewett, Graham Williams, et al. 2021. Differential privacy for public health data: An innovative tool to optimize information sharing while protecting data confidentiality. *Patterns* 2, 12 (2021), 100366.

[9] Mohammed El-hajj, Hussien Mousawi, and Ahmad Fadlallah. 2023. Analysis of Lightweight Cryptographic Algorithms on IoT Hardware Platform. *Future Internet* 15, 2 (2023), 15. https://doi.org/10.3390/fi15020054

[10] Medical Device Coordination Group et al. 2019. *MDCG 2019-16 Guidance on Cyber Security for Medical Devices*. Technical Report. European Commission.

[11] Junqi Guo, Minghui Yang, and Boxin Wan. 2021. A practical privacy-preserving publishing mechanism based on personalized k-anonymity and temporal differential privacy for wearable IoT applications. *Symmetry* 13, 6 (2021), 1043.

[12] Mohammad Kamrul Hasan, Muhammad Shafiq, Shayla Islam, Bishwajeet Pandey, Yousef A Baker El-Ebiary, Nazmus Shaker Nafi, R Ciro Rodriguez, and Doris Esenarro Vargas. 2021. Lightweight cryptographic algorithms for guessing attack protection in complex internet of things applications. *Complexity* 2021 (2021), 1–13.

[13] Ahmed K Jameil and Hamed Al-Raweshidy. 2022. Efficient CNN Architecture on FPGA Using High-Level Module for Healthcare Devices. *IEEE Access* 10 (2022), 60486–60495.

[14] Brendan Kelly, Conor Quinn, Aonghus Lawlor, Ronan Killeen, and James Burrell. 2023. Cybersecurity in Healthcare. In *Trends of Artificial Intelligence and Big Data for E-Health*. Springer International Publishing, Schweizer Verlag, 213–231.

[15] Georgios Kornaros. 2022. Hardware-Assisted Machine Learning in Resource-Constrained IoT Environments for Security: Review and Future Prospective. *IEEE Access* 10 (2022), 58603–58622. https://doi.org/10.1109/ACCESS.2022.3179047

[16] Hanchao Ku, Willy Susilo, Yudi Zhang, Wenfen Liu, and Mingwu Zhang. 2022. Privacy-Preserving federated learning in medical diagnosis with homomorphic re-Encryption. *Computer Standards & Interfaces* 80 (2022), 103583.

[17] Adarsh Kumar, Carlo Ottaviani, Sukhpal Singh Gill, and Rajkumar Buyya. 2022. Securing the future internet of things with post-quantum cryptography. *Security and Privacy* 5, 2 (2022), e200.

[18] Beibei Li, Shang Ma, Ruilong Deng, Kim-Kwang Raymond Choo, and Jin Yang. 2022. Federated anomaly detection on system logs for the internet of things: A customizable and communication-efficient approach. *IEEE Transactions on Network and Service Management* 19, 2 (2022), 1705–1716.

[19] Lin Lu, Jiayao Zhang, Yi Xie, Fei Gao, Song Xu, Xinghuo Wu, Zhewei Ye, et al. 2020. Wearable health devices in health care: narrative systematic review. *JMIR mHealth and uHealth* 8, 11 (2020), e18907.

[20] Konstantinos Makantasis, Anastasios D Doulamis, Nikolaos D Doulamis, and Antonis Nikitakis. 2018. Tensor-based classification models for hyperspectral data analysis. *IEEE Transactions on Geoscience and Remote Sensing* 56, 12 (2018), 6884–6898.

[21] Safa Otoum, Nadra Guizani, and Hussein Mouftah. 2022. On the Feasibility of Split Learning, Transfer Learning and Federated Learning for Preserving Security in ITS Systems. *IEEE Transactions on Intelligent Transportation Systems* 2022, 1 (2022), 1–9. https://doi.org/10.1109/TITS.2022.3159092

[22] Guansong Pang, Chunhua Shen, Longbing Cao, and Anton Van Den Hengel. 2021. Deep learning for anomaly detection: A review. *ACM computing surveys (CSUR)* 54, 2 (2021), 1–38.

[23] Tran Thi Phuong et al. 2023. Differentially private stochastic gradient descent via compression and memorization. *Journal of Systems Architecture* 135 (2023), 102819.

[24] Ayoub Si-Ahmed, Mohammed Ali Al-Garadi, and Narhimene Boustia. 2023. Survey of Machine Learning based intrusion detection methods for Internet of Medical Things. *Applied Soft Computing* 140 (2023), 110227. https://doi.org/10.1016/j.asoc.2023.110227

[25] Saurabh Singh, Shailendra Rathore, Osama Alfarraj, Amr Tolba, and Byungun Yoon. 2022. A framework for privacy-preservation of IoT healthcare data using Federated Learning and blockchain technology. *Future Generation Computer Systems* 129 (2022), 380–388.

[26] Gawsalyan Sivapalan, Koushik Kumar Nundy, Soumyabrata Dev, Barry Cardiff, and Deepu John. 2022. ANNet: A lightweight neural network for ECG anomaly detection in IoT edge sensors. *IEEE Transactions on Biomedical Circuits and Systems* 16, 1 (2022), 24–35.

[27] Hamed Soroush, David Arney, and Julian Goldman. 2016. Toward a safe and secure medical Internet of Things. *IIC J. Innov* 2, 1 (2016), 4–18.

[28] Philip Treleaven, Malgorzata Smietanka, and Hirsh Pithadia. 2022. Federated learning: the pioneering distributed machine learning and privacy-preserving data technology. *Computer* 55, 4 (2022), 20–29.

[29] Jai Narayan Tripathi, Binod Kumar, and Dinesh Junjariya. 2022. Hardware Accelerator Design for Healthcare Applications: Review and Perspectives. In *2022 IEEE International Symposium on Circuits and Systems (ISCAS)*. IEEE, New York, NY, USA, 1367–1371. https://doi.org/10.1109/ISCAS48785.2022.9937920

[30] Imtiaz Ullah and Qusay H Mahmoud. 2022. Design and development of RNN anomaly detection model for IoT networks. *IEEE Access* 10 (2022), 62722–62750.

[31] Athanasios Voulodimos, Nikolaos Doulamis, Anastasios Doulamis, and Eftychios Protopapadakis. 2018. Deep Learning for Computer Vision: A Brief Review. *Computational Intelligence and Neuroscience* 2018 (02 2018), 1–13. https://doi.org/10.1155/2018/7068349

[32] Zumin Wang, Jiyu Tian, Hui Fang, Liming Chen, and Jing Qin. 2022. LightLog: A lightweight temporal convolutional network for log anomaly detection on the edge. *Computer Networks* 203 (2022), 108616.

[33] William M Watkins, Samuel Yen-Chi Chen, and Shinjae Yoo. 2023. Quantum machine learning with differential privacy. *Scientific Reports* 13, 1 (2023), 2453.

[34] Brett Weinger, Jinoh Kim, Alex Sim, Makiya Nakashima, Nour Moustafa, and K. John Wu. 2022. Enhancing IoT anomaly detection performance for federated learning. *Digital Communications and Networks* 8, 3 (2022), 314–323. https://doi.org/10.1016/j.dcan.2022.02.007

[35] Susila Windarta, Suryadi Suryadi, Kalamullah Ramli, Bernardi Pranggono, and Teddy Surya Gunawan. 2022. Lightweight Cryptographic Hash Functions: Design Trends, Comparative Study, and Future Directions. *IEEE Access* 10 (2022), 82272–82294.

[36] Stephen Wong, Ram Mani, and Shabbar Danish. 2019. Comparison and selection of current implantable anti-epileptic devices. *Neurotherapeutics* 16 (2019), 369–380.

[37] Jean-Paul A Yaacoub, Mohamad Noura, Hassan N Noura, Ola Salman, Elias Yaacoub, Raphaël Couturier, and Ali Chehab. 2020. Securing internet of medical things systems: Limitations, issues and recommendations. *Future Generation Computer Systems* 105 (2020), 581–606.

[38] Mohammad Sadegh Yousefpoor, Efat Yousefpoor, Hamid Barati, Ali Barati, Ali Movaghar, and Mehdi Hosseinzadeh. 2021. Secure data aggregation methods and countermeasures against various attacks in wireless sensor networks: A comprehensive review. *Journal of Network and Computer Applications* 190 (2021), 103118. https://doi.org/10.1016/j.jnca.2021.103118

[39] Tingyi Zhou, Fabien Scalzo, and Bahram Jalali. 2022. Nonlinear Schrödinger Kernel for Hardware Acceleration of Machine Learning. *Journal of Lightwave Technology* 40, 5 (2022), 1308–1319.