

Internet of Things Cybersecurity Platform Benchmark: A Machine Learning Assessment

Robert-Alexandru Craciun, Radu Nicolae Pietraru, Mihnea Alexandru Moisescu

National University of Science and Technology Politehnica Bucharest, Romania (e-mail: robertcraciun28@gmail.com, radu.pietraru@upb.ro, mihnea.moisescu@upb.ro)

Abstract: Internet of Things (IoT) represents the interconnected network of different intelligent objects that contain sensors, software, and connectivity enables them to collect and exchange data over the internet. Ensuring the security of IoT is crucial in the digital realm as devices are all connected and they manage a lot of sensitive data. This paper proposes a benchmark comparison between a traditional IoT device and an AI accelerated device in order to evaluate the performances of each system in the context of IoT security and determine which platform can be used as a robust AI security detection and prevention system.

Keywords: Internet of Things, Artificial Intelligence, Hardware, Platforms, Single Board Computer, Benchmark, Cybersecurity

1. INTRODUCTION

Internet of Things represents the interconnected network of physical devices, vehicles, appliances, and other objects embedded with sensors, software, and connectivity which enables them to collect and exchange data over the Internet. The idea of IoT emerged as a natural progression of the escalating digitization and interconnection of everyday objects. The primary objective of IoT is to facilitate seamless communication between devices and enable them to make intelligent, data driven decisions (Farhan et al., 2018). The potential benefits of IoT include real time monitoring, predictive analytics, and optimization of processes that are leading to increased productivity and resource efficiency. On the other hand, the widespread adoption of IoT raises concerns about privacy, security, and ethical implications due to the massive data collection that the IoT devices gather while used.

IoT emerges as a pivotal force by intricately integrating connectivity into various industries. By establishing complex connections between diverse systems and facilitating seamless information exchange, IoT unlocks substantial benefits across a wide spectrum of industries.

In the realm of healthcare (Yugha et al., 2022), IoT's influence has a considerable impact, enabling real time monitoring of patients' health metrics, enhancing preventive care, and providing timely interventions. Wearable devices, equipped with sensors that continuously track vital signs offer the capability to healthcare providers to remotely monitor patients, revolutionizing the delivery of personalized and preventive care. To enhance the data privacy and security, the authors proposed a Zero-Knowledge Proof authentication, in conjunction with Blockchain's AZTEC protocol to enable secure data exchange in IoT, in healthcare applications.

The agriculture sector (Vishali Priya and Sudha, 2021) benefits from IoT by leveraging sensors that monitor soil conditions, crop health, and livestock wellbeing. Precision agriculture, enabled by IoT, involves the use of drones and sensors to optimize irrigation, fertilizer application, and pest control, increasing the efficiency and reducing resource wastage.

In transportation (Ushakov et al., 2022), IoT contributes to smart logistics and traffic management, optimizing routes, reducing congestion, and enhancing fuel efficiency. Vehicles equipped with IoT sensors can communicate with traffic infrastructure to provide real time traffic updates, enabling drivers to choose optimized routes and reduce overall travel time.

The manufacturing industry (Yang et al., 2019) witnessed increased efficiency and cost effectiveness using IoT applications such as predictive maintenance, which minimizes downtime and maximizes production uptime. Sensors embedded in the machines can collect data on equipment health, enabling proactive maintenance before a breakdown occurs, reducing operational disruptions and maintenance costs.

Smart homes (Al-Mutawa and Eassa, 2020) create interconnected and automated living spaces by using IoT devices, offering residents convenience, energy savings, and enhanced security through surveillance devices.

In essence, the widespread adoption of IoT transforms industries by promoting efficiency, automation, and improved user experiences, with specific applications tailored to the unique needs of each industry. Moreover, it becomes imperative to address the critical aspect of security of the industries in context of IoT devices.

Ensuring the security of IoT is crucial in the digital realm as devices are all connected, and they manage a lot of sensitive data. As IoT systems seamlessly integrate into various environments, the risks associated with inadequate security measures become increasingly significant. Unauthorized access to personal data could lead to privacy breaches, identity theft or even physical security threats (Ahmed and Alheeti, 2022). Additionally, insecure IoT devices can be exploited by cyberattacks that can lead to disruptions in the infrastructure or creation of botnets (Nguyen et al., 2022). Therefore, it is necessary for users and manufacturers to protect against these potential risks and make sure that the IoT ecosystem is secure and resilient to cybersecurity threats. One technology that can

help reinforcing the overall security mechanisms for IoT is Artificial Intelligence.

Integrating Artificial Intelligence (AI) into IoT security enhances the effectiveness of cybersecurity efforts as AI enables real time analysis of larger dataset provided by the IoT devices, identifying patterns that can indicate potential security threats. Machine learning algorithms improve threat detection and response by adapting to new data, automating monitoring, and rapidly detecting and mitigating security breaches proactively (Javed et al., 2023). However, the integration of AI in IoT security raises concerns as not all the capabilities can be effective on IoT devices as there are resource constraints, complex to maintain or interoperability issues. In the evolving domain of IoT security, AI holds promise in fortifying defenses against sophisticated cyber threats. AI protocols can be implemented on a wide array of IoT devices that vary in functionalities, with some of them being built specifically for AI integration.

There is a wide range of IoT device platforms that can be used to enhance the IoT security mechanisms. The inclusion of AI accelerators in these devices boosts the processing capabilities of these devices, enabling them to quickly analyze and respond to security threats in real time. These accelerators empower IoT devices to utilize machine learning algorithms for tasks such as detecting anomalies, recognizing patterns and conducting predictive analysis that allow proactive identification and management of potential vulnerabilities. As the IoT environment evolves, the strategic use of AI accelerators on device platforms proves essential in addressing the cyber security mechanisms that need to be implemented. However, a benchmark of machine learning algorithms in the context of IoT security between a traditional platform and an AI accelerated one needs to be conducted.

This paper will propose a benchmark comparison between a traditional IoT device, and an AI accelerated device in order to evaluate the performances of each system in the context of IoT security and determine which platform can be used as a robust AI security detection and prevention system. The paper will concentrate on data and network security, excluding the topic of physical security. This paper will contain a review of currently available benchmark suites, an overview of the proposed benchmark implementation and the tests that will be conducted, a discussion on the results and how the chosen platform will be used to develop the security system.

2. RELATED WORK

One of the most important aspects of IoT is the security of the systems and the incorporation of machine learning emerges as a vital component in strengthening the interconnected systems. The fusion between IoT and ML offers a powerful solution for protecting devices and data against evolving threats. Machine learning offers a flexible and adaptive strategy for security, allowing systems to learn and recognize behavior patterns and identify potential security breaches. By implementing ML driven intrusion detection mechanisms, IoT ecosystems can identify and respond to malicious activity. Additionally, ML contributes to device identification and authorization ensuring that only legitimate devices have access to the IoT network. The ability to constantly learn from new data patterns enhances

the resilience of IoT networks, safeguarding sensitive information and upholding the integrity of interconnected devices.

Various researchers conducted studies exploring the integration of ML in securing IoT. Their research is driven by the urgent need to address the growing threats to IoT ecosystems. Researchers have focused on creating and testing intrusion detection mechanisms utilizing ML algorithms to quickly identify and respond to unauthorized access or malicious activities within IoT networks. Additionally, efforts have been made to improve device identification and authorization processes, ensuring that only legitimate devices have access to the network. The research community demonstrates efforts to establish robust security framework for IoT ecosystems through application of ML. In the current section, a selection of these studies will be presented to provide an overview of advancements in this domain.

In paper (Banaamah and Ahmad, 2022), the authors underscore the escalating cybersecurity challenges for IoT ecosystems and propose a thorough implementation of deep learning techniques, including Convolutional Neural Networks (CNN), Long Short-Term Memory (LSTM), and Gated Recurrent Units (GRUs) for intrusion detection. The proposed method aims to enhance the accuracy and mitigate false alarms, addressing vulnerabilities associated with denial of service, eavesdropping, and privilege escalation. Experimental results on the Bot-IoT dataset demonstrate promising outcomes, with LSTM having superior performance.

Paper (Kiran et al., 2020) focuses on developing an intrusion detection system for IoT using ML. The study uses a test bed to simulate an IoT environment, capturing data from sensors during normal and attacking phases. With the implementation of man-in-the-middle attacks through ARP poisoning, the collected data is used to train machine learning classifiers, including Naïve Baies, SVM, decision trees, and Adaboost. Results reveal high accuracy in classifying normal and attack data, with the decision tree model achieving perfect accuracy.

In paper (Bouazza et al., 2022), the focus is on addressing security challenges in IoT, specifically vulnerabilities associated with the RPL routing protocol. The research introduces an anomaly intrusion detection system based on machine learning. The custom generated dataset is used for classification algorithms, with Random Forest and Decision Trees offering superior performance, with high accuracy, precision and reduced false alert rate compared to related works.

Paper (Baral et al., 2023) addresses the security concerns associated with compromised IoT devices in wireless networks. The authors propose a device identification mechanism using deep learning approaches like CNN and a combination of LSTM and CNN. The results indicate that CNN performs well in device classification compared to the combined model and the proposed approach shows promise in accurately identifying devices based on their unique communication patterns, providing an alternative to traditional identification methods prone to spoofing.

The paper (Koball et al., 2023) proposes an unsupervised machine learning approach for device identification. Traditional methods, such as MAC addresses, face challenges like easy spoofing and supervised ML approaches require extensive labeled datasets. The proposed method utilizes a K-Nearest Neighbor classifier, incorporating preprocessing, outlier removal, feature selection, and clustering. The study evaluates performance of both supervised and unsupervised ML on a dataset comprising 8 IoT devices. While supervised ML models generally have higher accuracy, unsupervised models demonstrate reasonable results and offer advantages in dynamic network scenarios.

Paper (Meidan et al., 2017) presents an approach to accurately identify IoT devices connected to a network using ML algorithms applied to network traffic data. The study involves the collection and labeling of network traffic from nine different IoT devices, PCs, and smartphones. Through supervised learning, a complex classifier is trained using Gradient Boosting Machine, Random Forest, and XGBoost models. The overall accuracy of the proposed IoT classification is reported as 99.28%, which demonstrates a good option for distinguishing and identifying IoT devices based on network analysis.

While machine learning algorithms offer promising solutions for IoT device classification based on network traffic data, the absence of suitable platform benchmarks in this domain poses a significant challenge. Based on the reviewed papers that offer good solutions for creating a secure gateway for IoT devices using ML, the current paper will propose a customized benchmark for machine learning assessment on different hardware platforms.

3. AI PLATFORMS FOR IOT CYBERSECURITY

In the dynamic domain of IoT, the utilization of AI accelerated platforms represents the next step in the context of cybersecurity within IoT ecosystems. Devices like Orange Pi, equipped with a Neural Processing Unit (NPU) or the Coral Dev Board, equipped with a Tensor Processing Unit (TPU), exemplify the next generation of edge computing. These specialized hardware components offer IoT devices the capability to execute intricate machine learning algorithms locally, without relying on centralized cloud resources. In the context of cybersecurity of IoT ecosystems, AI models can analyze data at the edge in real time, spotting anomalies and potential security threats. By leveraging dedicated AI accelerators, these devices can enhance the efficiency and response of security systems, raising the security of IoT networks. The decentralized approach not only ensures real time threat identification but also mitigates the risks associated with transmitting sensitive data to external platforms, contributing to a more resilient and secure IoT infrastructure.

The comparison between traditional IoT devices, exemplified by the widely known Raspberry Pi, and their more advanced counterparts featuring AI accelerators like NPUs and TPUs is a crucial exploration in the context of cybersecurity of IoT ecosystems. While traditional IoT devices are appreciated for their versatility and cost effectiveness (Ariza and Báez, 2022), they encounter challenges related to computational limitation when used to execute complex intrusion detection algorithms.

On the other hand, AI accelerated devices that have a dedicated hardware component in the form of NPUs or TPUs offer an appealing alternative by significantly enhancing the speed and effectiveness of machine learning tasks (Lee, 2021).

Traditional IoT devices are widely seen as adaptable and affordable for a wide range of use cases. They serve as versatile platforms for diverse applications within IoT ecosystems. However, their computational power may not be sufficient when used for complex intrusion detection algorithms due to high computational demand. These algorithms are designed to distinguish anomalies and potential security breaches and require substantial computing power for real time execution where traditional IoT devices may encounter constraints.

In contrast, NPUs and TPUs introduce a new paradigm by providing dedicated AI hardware acceleration components to IoT ecosystems. This augmentation significantly strengthens the speed and efficiency of machine learning tasks, making them proficient in addressing the computational requirements for real time intrusion detection mechanisms. NPUs, with their efficient handling of neural network workloads (Tan and Cao, 2021), and TPUs, calibrated for accelerating matrix operations for deep learning models (Shahid and Mushtaq, 2020), provide a powerful computational toolkit that innovates IoT cybersecurity mechanisms.

The specialization embedded in NPUs and TPUs empowers AI accelerated devices to execute complex intrusion detection algorithms with a blend of speed and accuracy directly at the edge. The edge computing paradigm, where computational tasks are performed close to the data source, substantially benefits from the accelerated processing offered by these dedicated hardware units. This leads to real time and precise identification of potential security threats within the IoT ecosystem.

While traditional IoT devices retain their accessibility and versatility, AI accelerated devices offer a balance between computational power and cost effectiveness. The performance boost facilitated by NPUs and TPUs positions these devices as robust solutions for performant intrusion detection systems in IoT networks. This comparison highlights the evolving landscape where the strong security mechanisms intersect with the necessity for efficient and cost effective computational solutions in the expanding domain of IoT.

The current paper will provide a comprehensive comparison of different platforms, both traditional and AI accelerated, to determine which platform can be used to create a powerful secure IoT gateway to protect IoT devices that communicate using the TCP/IP protocol. The following chapter will showcase the comparison methodology and what metrics will be used to compare the IoT edge platforms.

4. METHODOLOGY

The conducted study contains various stages that includes a review of literature that has been previously analyzed, benchmark description for the tested components related to device identification and intrusion detection, an examination of individual test results, why were the specific tests selected

for the benchmark, what platforms are used for the comparison and why the benchmark focuses on the CPU for each platform.

A different study (Crăciun et al., 2023a) assessed various ML algorithms using a Python implementation using various machine learning algorithms focusing on intrusion detection in IoT networks. The study utilized the publicly available BoT-IoT dataset from the University of New South Wales Canberra, which included both normal and botnet traffic. To overcome hardware limitations, a 5% subset of the dataset was extracted with an 80%-20% split for training and testing the ML models. The training file contains a dataset of 2.934.817 records and the testing file contains a similar dataset, but with a smaller number of records with just 733.706 records. The preprocessed dataset contains ten crucial features for attack classification and four ML algorithms were selected for model development: Random Forest (RF), Decision Trees (DT), Naïve Bayes (NB) and eXtreme Gradient Boost (XGB). The evaluation metrics used are the accuracy, precision, sensitivity, hamming loss and F1-score for the comparison between all these algorithms. The tests showed that the XGB algorithm exhibited superior performance across all metrics, proving to be the most accurate classifier for predicting attacks, categories, and subcategories. The findings were validated through k-fold Cross Validation, confirming the effectiveness of XGB for predicting intrusions in IoT ecosystems for IoT devices that communicate using the TCP/IP protocol.

XGB integrates Lasso (L1) and Ridge (L2) regularization terms into its cost function. This strategic integration serves the purpose to prevent overfitting by penalizing complex models, enhancing robustness and reducing susceptibility to memorizing the training data (Chen and Guestrin, 2022). In essence, XGB's incorporation of regularization terms not only prevents overfitting but also fine tunes the model, ensuring a more resilient and adaptable predictive performance.

The objective function combines a loss term and regularization terms. The general form of the objective function for a regression task is as follows:

$$Obj = \sum_{i=1}^n L(y_i, \hat{y}_i) + \sum_{k=1}^K \Omega(f_k) \quad (1)$$

Here:

- $L(y_i, \hat{y}_i)$ represents the loss term, measuring the difference between the predicted and actual values
- $\Omega(f_k)$ represents the regularization term for each tree, incorporating both L1 and L2 regularization

The objective for learning each tree is a sum of the per-node loss and regularization terms. For a given tree f_k , the objective is defined as follows:

$$Obj = \sum_{j=1}^T \left[gain \cdot \frac{G_j^2}{H_j + \lambda} + \alpha \cdot penalty \right] \quad (2)$$

Here:

- G_j and H_j are the sum of the gradients and Hessians in node j
- λ and α are the regularization parameters

XGBoost introduces the regularized learning objective to guide the optimization process, striking a balance between fitting the training data and preventing overfitting. It is a versatile algorithm and can be applied to various types of machine learning tasks including classification, regression, and ranking. Its flexibility makes it a valuable tool across a wide range of domains. The current study will use this benchmark to test the platforms, as well as the following study that focuses on device identification.

An additional study (Crăciun et al., 2023b) similarly assessed various ML algorithms using a Python implementation by creating different models for device identification of IoT hardware platforms. Utilizing a dataset provided by Emmanuel Tsukerman, containing 1000 records for training and, respectively, testing with pre cleaned and labeled data featuring device categories, the study explores 298 network metrics between IoT devices and external network, organized into different categories. Multiple popular ML algorithms are evaluated and compared using similar metrics like the previous study. The results indicate that Gradient Boost (GB) performed the best across all metrics and despite the limited dataset GB proves to be a reliable classifier for identifying IoT devices that communicate using the TCP/IP protocol.

GB minimizes the loss function by moving in the direction of steepest descent. The weak learners used are typically decision trees with shallow depths (He et al., 2020). These trees are fit to the residuals or errors of the previous models. The choice of loss function is crucial as the functions include mean squared error for regression problems and log loss for classification problems. The algorithm minimizes the loss function $L(y, F(x))$ by adjusting the parameters of the weak learners in the direction opposite to the gradient of the loss. The loss function measures the difference between the true target y and the current ensemble prediction $F(x)$.

For regression problems, Mean Square Error is commonly used:

$$L(y, F(x)) = \frac{1}{2} (y - F(x))^2 \quad (3)$$

For classification problems, Log Loss is often defined as:

$$L(y, F(x)) = -[y \log(F(x)) + (1 - y) \log(1 - F(x))] \quad (4)$$

The iterative process involves minimizing the loss function by updating the ensemble at each step:

$$F_{m+1}(x) = F_m(x) - \eta \cdot \nabla L(y, F(x)) \quad (5)$$

The Python script from both studies will be used to compare the platforms from the current paper. Only the best model from each paper will be used for benchmarking as there is no need to train all the models again to come to the same conclusion. Both studies were selected as benchmarking options due to the fact that both provide robust solutions for device identification and intrusion detection mechanisms using machine learning algorithms for IoT devices that communicate using the TCP/IP protocol with the edge device.

There are 3 different platforms used to train and compare algorithm training. Each device has a unique hardware design as they have unique AI accelerator types (no accelerator, NPU,

and TPU). The 3 platforms have the following hardware configuration:

Table 1. Tested platforms specifications.

Spec	Raspberry Pi 5	Orange Pi 5	Coral Dev Board
Socket	BCM2712	RK3588S	NXP i.MX 8M SoC
CPU	ARM-Cortex A76	ARM-Cortex A76 + ARM-Cortex A55	ARM-Cortex A53 + ARM-Cortex M4F
RAM	4GB LPDDR4X	4GB LPDDR4/4x	1GB LPDDR4
Storage	SD Card	SD Card	eMMC
External Storage	USB SSD SATA3	USB SSD SATA3	USB SSD SATA3
AI Accelerator	N/A	NPU Accelerator – 6TOPS	TPU Accelerator – 4TOPS

The platforms will be compared using the following metrics: dataset reading time, models training time, CPU time, RAM utilization and CPU utilization.

The proposed benchmarking tests are using the CPU of the devices as both device identification and intrusion detection are classification problems and both are a mismatch for AI accelerators (He et al., 2020). Even though only the CPU is used, the current benchmark still shows a powerful result since the detection algorithms need to be retrained constantly to be aware of both new devices added to the gateway and new network vulnerabilities that need to be addressed.

5. RESULTS

After running the scripts on each platform, the models produced a set of metrics for each model type. The scope is to compare the performance of each hardware configuration to find a suitable platform that can be used as a secure IoT gateway to identify the connected devices and secure them against intrusions. The performance is measured by the following metrics: dataset reading time, models training time, CPU time, RAM utilization and CPU utilization.

5.1 Device identification results

The model trained for device identification is created by using the Gradient Boost algorithm.

a. Data reading times: Figure 1. contains reading times of both training and testing data for each platform.




Fig. 1. Dataset reading times [s].

The reading times are good for all platforms. Even though the dataset has just 1000 rows for each dataset, testing and

training, the reading is done fast on all the platforms. The slowest platform is the Coral Dev Board as the eMMC has slower reading speeds than the SD Card used on the other platforms. Scaling the dataset means that it will take longer to access the data, but each platform should take fairly acceptable times to read a bigger dataset.

b. Training time GB: Figure 2. contains training times for the GB model for each platform.

The model creation time has acceptable times for the Raspberry Pi 5 and the Orange Pi 5, but the Coral Dev Board is more than 3 times slower to create the model, even considering the small dataset. If the dataset has more records, the Coral Dev Board will struggle to create the model and will face delayed deployment of the model or even failing to create it. The other boards will be able to keep up with the demand, due to their powerful CPU and extended RAM capacity. In order to keep the Coral Dev Board as a good candidate for the IoT gateway, other algorithms can be used to create the model (Craciun et al., 2023b), but it will come with decreased efficiency in detecting the devices.




Fig. 2. Training time GB [s].

c. CPU times: Figure 3. contains CPU times for the GB model training for each platform.




Fig. 3. CPU times [s].

The CPU time represents the time spent by the CPU when executing the script. Since the Coral Dev Board has the slowest CPU out of the three tested platforms, the time spent by the CPU to read the dataset and create the model is the highest. The best platform is the Orange Pi 5 which is more than 3 times faster than the Coral Dev Board for the time spent processing the script.

d. Average RAM usage: Figure 4. contains the average RAM usage for the GB model training for each platform.




Fig. 4. RAM usage [MB].

RAM usage has similar average amount on all the platforms due to the same dataset that is used in all experiments. The lower RAM utilization is represented by the 2000 entries for the training and testing datasets, which don't use a lot of RAM to store the information.

e. Average CPU usage: Figure 5. contains the average CPU usage for the GB model training for each platform.

CPU usage is influenced by the CPU architecture of each platform. The cores and threads for each CPU architecture influence the percentage usage on each platform. All platforms have multi core, multi-threading CPUs, which the library for the ML algorithm uses to accelerate the creation of the model. All the platforms use some of their capabilities, but the library is not able to use all the resources available. The Orange Pi 5 is the one with the highest usage of its resources for training the algorithm.

f. Device identification conclusions: All the platforms provide diverse capabilities for different use cases. An essential application is IoT identification, which can be significantly enhanced through the usage of ML techniques, given their adaptability in construction classification models. The conducted test showed that the Orange Pi 5 has an edge over the other platforms. The device demonstrated its ability to create robust models, serving as an effective IoT gateway for recognizing connected devices. The second option is the Raspberry Pi 5, which offers almost similar performance, but with the lacking capability for future proof in AI usage, Orange Pi 5 having the advantage with the onboard NPU that can be leveraged at a later time to create even more effective models for device identification.




Figure 5. CPU usage [%]

5.2 Intrusion detection results

The model trained for intrusion detection is created by using the eXtreme Gradient Boost algorithm. Due to the huge amount of data used for creating the model, the Coral Dev Board couldn't execute the creation script successfully. The low amount of RAM of only 1 GB is not enough to successfully create the model. This information is useful as for a future development, a custom dataset will be used and the 1 GB of RAM will not be enough to create the model.

a. Data reading times: Figure 6. contains reading times of both training and testing data for each platform.




Fig. 6. Dataset reading times [s].

The dataset reading times are as expected as the more powerful the hardware is, the faster it will read the dataset. The faster CPU and RAM components on the Orange Pi 5 resulted in faster reading times of the dataset.

b. Training time XGB: Figure 7. contains training times for the XGB model for each platform.




Fig. 7. Training time XGB [s].

The resulting times for each platform are as expected. The Orange Pi 5 has more powerful hardware, so the execution time of the model training is more 1.5 times faster than the Raspberry Pi 5. This means that the Orange Pi 5 can create models with huge amounts of data in a timely manner. Even though the Raspberry Pi 5 creates the model slower, it still represents a viable option for other machine learning use cases.

c. CPU times: Figure 8. contains CPU times for the XGB model training for each platform.




Fig. 8. CPU times [s].

The CPU processing times represent how much time the CPU was used when executing the training script. The Orange Pi 5 has slightly lower CPU time due to its more powerful CPU that needs less time to execute the script than the Raspberry Pi 5. Both hardware platforms are suitable for this type of ML technique as both platforms have powerful CPUs embedded for edge computing. The edge here is for the Orange Pi 5 as it has a dedicated NPU that can be used to further accelerate the use of ML models in a later iteration.

d. Average RAM usage: Figure 9. contains the average RAM usage for the XGB model training for each platform.




Fig. 9. RAM usage [MB].

RAM usage has a notable difference between the two platforms. The Orange Pi 5 used less RAM with almost 1.5 times than the Raspberry Pi 5. Using the same dataset, the RAM component on the Orange Pi 5 is more efficient for bigger datasets than the component on the Raspberry Pi 5 which used almost the entire memory available. The Orange Pi 5 can be used to create machine learning models with even bigger datasets that will have increased accuracy when used.

e. Average CPU usage: Figure 10. contains the average CPU usage for the XGB model training for each platform.

CPU usage is influenced by the CPU architecture of each platform. The Orange Pi 5 platform has higher CPU usage as the machine learning library makes use of more resources than the Raspberry Pi 5. The XGB algorithm takes advantage of the multi-threading architecture of both CPUs, and it uses the most processing power that the library is able to pull out from the available resources.




Fig. 10. CPU usage [%].

The more usage on the CPU of the Orange Pi 5 resulted in faster times to create the machine learning model.

f. Intrusion detection conclusions: All platforms offer different levels of processing power for ML algorithms. The Orange Pi 5 is the best choice for intrusion detection based on machine learning models as it can easily create models, with a complex dataset, with low processing times. The Raspberry Pi 5 represents a good competitor for the Orange Pi 5, but the latest has an edge over the Raspberry Pi 5 due to its NPU component that can be used at a later time to create even more efficient models. The Coral Dev Board couldn't be used for the intrusion detection part as the dataset was too large for the 1 GB of memory that the board has, however, in the future, it can be used with different algorithms to remake the tests as it opens opportunities for further investigations.

5.3 Platform conclusions

In the conducted benchmark, the comparison between the three platforms, namely Raspberry Pi 5, Orange Pi 5, and Coral Dev Board, resulted with the Orange Pi 5 as the superior choice across different performance metrics. The benchmark was conducted on two critical tests: the first involved device identification utilizing the GB algorithm model, while the second focused on intrusion detection utilizing an XGB algorithm model. The Orange Pi 5 outperformed the other platforms in almost all aspects, including dataset reading times, CPU and RAM utilization and overall CPU time. The resulting performance places the Orange Pi 5 as an ideal candidate for serving as a secure IoT gateway for devices communicating using the TCP/IP protocol. Furthermore, the Orange Pi 5 has an NPU component embedded that can accelerate AI operations which can enhance the efficiency in processing the network data. As a result, the Orange Pi 5 can be used as a reliable solution for building secure IoT infrastructures.

6. FUTURE DEVELOPMENT

Leveraging the capabilities of NPUs and TPUs for secure IoT gateway solutions significantly enhances scalability, enabling the system to handle increased data loads and more complex threat landscapes effectively. By integrating AI accelerators for real-time data processing and inference, the solution can rapidly adapt to the growing number of connected devices and diverse security requirements. This approach not only ensures robust and scalable intrusion detection and device

identification but also influences future hardware development by emphasizing the necessity for AI accelerator compatibility. Consequently, it guides the selection of hardware platforms that prioritize advanced computational efficiency and real-time security, essential for the evolving and expanding IoT ecosystem.

Looking ahead, the future development of the secure IoT gateway solution based on the mentioned platforms creates a need to explore the full capabilities of NPUs and TPUs. Given that the NPUs specialize in accelerating neural network computations (Tan and Cao, 2021) and TPUs excel in processing deep learning models (Shahid and Mushtaq, 2020), a common approach is essential to ensure compatibility across diverse hardware platforms. The development roadmap should prioritize the creation of a model that integrates with both NPUs and TPUs, offering an intrusion detection and device identification solution that can be applied across different devices. This approach facilitates the utilization of AI accelerators, maximizing the potential for real time data processing and inference.

In the evolving IoT domain, where security is one of the most important factors for IoT ecosystems, leveraging the full potential of NPUs and TPUs becomes mandatory. While CPUs of the mentioned hardware platforms demonstrate good performance, the incorporation of AI accelerators ensures not only enhanced computational efficiency, but also real time threat detection and response. By using these accelerators for instantaneous processing of data, the secure IoT gateway can rapidly identify and mitigate potential threats, providing robust defense mechanisms. The forward strategy is to minimize vulnerabilities and to create a safe environment for IoT devices that can operate securely, contributing to the advancement of IoT.

7. CONCLUSIONS

The current paper presented a comprehensive benchmark analysis of three different hardware platforms, namely Raspberry Pi 5, Orange Pi 5, and Coral Dev Board, in the context of IoT security applications. The research focused on evaluating the performance of these platforms through two main tests: intrusion detection using the XGB algorithm and device identification using the GB algorithm. Through experimentation and analysis, the benchmark results revealed that the Orange Pi 5 was the best performing platform based on various metrics, showcasing its advantages in addressing the demanding requirements of AI driven IoT security applications for real time protection.

The current benchmark not only contributes to the comparative performance of popular hardware platforms but also puts an emphasis on the importance of selecting the right hardware for effective implementation of security algorithms for IoT ecosystems. As the demand for robust and scalable IoT solutions continues to rise, the identified strengths of the Orange Pi 5 make it compelling for practitioners and researchers to use the device for creating powerful ML models for IoT security.

REFERENCES

- Ahmed, A., & Alheeti, K.M. (2022). Intelligent Identification of Unauthorized Internet of Things Devices. 2022 Fifth International Conference of Women in Data Science at Prince Sultan University (WiDS PSU), 37-42.
- Al-Mutawa, R., & Eassa, F. (2020). A Smart Home System based on Internet of Things. ArXiv, abs/2009.05328.
- Ariza, J.Á., & Báez, H. (2021). Understanding the role of single-board computers in engineering and computer science education: A systematic literature review. *Computer Applications in Engineering Education*, 30, 304 - 329.
- Banaamah, A.M., & Ahmad, I. (2022). Intrusion Detection in IoT Using Deep Learning. *Sensors (Basel, Switzerland)*, 22.
- Baral, P., Yang, N., & Weng, N. (2023). IoT Device Identification Using Device Fingerprint and Deep Learning. *IntechOpen*. doi: 10.5772/intechopen.111554.
- Bouazza, A., Debbi, H., & Lakhlef, H. (2022). Machine Learning-based Intrusion Detection System Against Routing Attacks in the Internet of Things. *Tunisian-Algerian Joint Conference on Applied Computing*.
- Chen, T., & Guestrin, C. (2016). XGBoost: A Scalable Tree Boosting System. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*.
- Craciun, R., Nicolae Pietraru, R., & Alexandru Moisesescu, M. (2023). IoT Device Identification: A Machine Learning Assessment. 2023 13th International Symposium on Advanced Topics in Electrical Engineering (ATEE), 1-6.
- Craciun, R., Pietraru, R.N., Mocanu, S., & Moisesescu, M.A. (2023). IoT Intrusion Detection: A Machine Learning Assessment. 2023 24th International Conference on Control Systems and Computer Science (CSCS), 307-312.
- Farhan, L., Kharel, R., Kaiwartya, O., Quiroz-Castellanos, M., Alissa, A., & Abdulsalam, M. (2018). A Concise Review on Internet of Things (IoT) -Problems, Challenges and Opportunities. 2018 11th International Symposium on Communication Systems, Networks & Digital Signal Processing (CSNDSP), 1-6.
- He, M., Vijaykumar, T.N., & Thottethodi, M. (2020). Booster: An Accelerator for Gradient Boosting Decision Trees Training and Inference. 2022 IEEE International Parallel and Distributed Processing Symposium (IPDPS), 1051-1062.
- Javed, A., Awais, M., Shoaib, M., Khurshid, K.S., & Othman, M. (2023). Machine learning and deep learning approaches in IoT. *PeerJ Computer Science*, 9.
- Kiran, K.S., Devisetty, R.N., Kalyan, N.P., Mukundini, K., & Karthi, R. (2020). Building a Intrusion Detection System for IoT Environment using Machine Learning Techniques. *Procedia Computer Science*, 171, 2372-2379.
- Koball C, Rimal BP, Wang Y, Salmen T, Ford C. IoT Device Identification Using Unsupervised Machine Learning. *Information*. 2023; 14(6):320. <https://doi.org/10.3390/info14060320>

- Lee, K.J. (2021). Chapter Seven - Architecture of neural processing unit for deep neural networks. *Adv. Comput.*, 122, 217-245.
- Meidan, Y., Bohadana, M., Shabtai, A., Guarnizo, J.D., Ochoa, M., Tippenhauer, N., & Elovici, Y. (2017). ProfilIoT: a machine learning approach for IoT device identification based on network traffic analysis. *Proceedings of the Symposium on Applied Computing*.
- Nguyen, G.L., Dumba, B., Ngo, Q., Le, H., & Nguyen, T.N. (2021). A collaborative approach to early detection of IoT Botnet. *Comput. Electr. Eng.*, 97, 107525.
- Shahid, A., & Mushtaq, M. (2020). A Survey Comparing Specialized Hardware And Evolution In TPUs For Neural Networks. 2020 IEEE 23rd International Multitopic Conference (INMIC), 1-6.
- Tan, T., & Cao, G. (2021). Deep Learning on Mobile Devices Through Neural Processing Units and Edge Computing. *IEEE INFOCOM 2022 - IEEE Conference on Computer Communications*, 1209-1218.
- Ushakov, D., Dudukalov, E.V., Kozlova, E., & Shatila, K. (2022). The Internet of Things impact on smart public transportation. *Transportation Research Procedia*.
- Vishali Priya, O., & Sudha, R. (2021). Impact of Internet of Things (IoT) in Smart Agriculture. *Recent Trends in Intensive Computing*.
- Yang, H., Kumara, S.R., Bukkapatnam, S.T., & Tsung, F. (2019). The internet of things for smart manufacturing: A review. *IJSE Transactions*, 51, 1190 - 1216.
- Yugha, R., et al. (2022) "Privacy Protected IoT-Blockchain using ZKP for Healthcare application." *CONTROL ENGINEERING AND APPLIED INFORMATICS* 24.4 76-87.