

**FEDERATED DEEP LEARNING-BASED INTRUSION DETECTION SYSTEM FOR
SECURING IOT NETWORKS ON SOLAR SMART CAMERAS**

BY

BRIGHT-AGBONZE OSAFONA DESMOND

ENG2106239

SUPERVISED BY:

ENGR DR MRS ODUWARE OKOSUN

**A PROJECT WORK SUBMITTED TO THE DEPARTMENT OF COMPUTER
ENGINEERING, FACULTY OF ENGINEERING UNIVERSITY OF BENIN, BENIN
CITY**

**IN FULFILMENT OF THE REQUIREMENT FOR THE AWARD OF BACHELOR OF
ENGINEERING (B.Eng.) IN COMPUTER ENGINEERING**

OCTOBER, 2025

CERTIFICATION

This is to certify that this project, Federated Deep Learning-Based Intrusion Detection System for Securing IOT Networks on Solar Smart Cameras, was carried out by **BRIGHT-AGBONZE OSAFONA DESMOND (ENG2106239)**, in the Department of Computer Engineering, Faculty of Engineering, University of Benin.

Engr. Dr. I. A. Edeoghon
(Head of Department)

DATE

Engr Dr Mrs Oduware Okosun
(Project Supervisor)

DATE

DEDICATION

I dedicate this project to God Almighty for His grace, strength, wisdom, and guidance throughout the course of this research work. I also dedicate this work to my late Father, Mr. BRIGHT AGBONZE who passed on during the course of my final year, bless his soul. His constant support, prayers, and encouragement contributed greatly to the successful completion of this project my academic journey in general.

ACKNOWLEDGEMENT

I sincerely express my profound gratitude to God Almighty for granting me the knowledge, strength, and ability to successfully complete this project titled “Federated Deep Learning-Based Intrusion Detection System for Solar-Powered Smart Cameras in Classroom Monitoring and Attendance Systems.”

My special appreciation goes to my project supervisor for the valuable guidance, constructive corrections, encouragement, and professional support provided throughout the duration of this research work. Your patience and academic contributions played a significant role in the successful completion of this project.

I also appreciate the efforts of all the lecturers in the department for the knowledge and academic foundation they have impacted in me throughout my years of study.

I am deeply grateful to my mom, brothers and family members for their prayers, mental support , financial support, understanding, and constant encouragement during the course of this project.

Special thanks also go to Success, Obina, David, Jeremiah and friends too many to mention who contributed ideas, motivation, and technical assistance during the development and testing stages of this research.

Finally, I appreciate everyone who contributed directly or indirectly to the success of this project. May God bless and reward you all abundantly.

ABSTRACT

This study presents the design and implementation of a smart, lightweight, federated deep learning system that integrates solar-powered cameras for automated attendance, unauthorized entry

prevention and real-time cyber threat detection in academic environments. Using TensorFlow Lite, Python, and a Flask-based web interface, the model achieved high accuracy in facial recognition while maintaining low computational and energy costs. A structured SQLite3 database supported efficient local data handling, while solar energy integration enabled autonomous and sustainable operation. This project validates the potential of combining renewable energy, artificial intelligence, and federated learning to enhance classroom management and IoT security in low-resource settings.

TABLE OF CONTENTS

CERTIFICATION	i
DEDICATION	ii
ACKNOWLEDGEMENT	iii
ABSTRACT	iii
TABLE OF CONTENTS	v
TABLE OF FIGURES	vii
TABLE OF TABLES	viii
CHAPTER ONE	1
1.1 Background of the Study	1
1.2 Problem Statement	2
1.3 Aim and Objectives	3
1.3.1 Aim	3
1.3.2 Objectives	4
1.4 Scope of Work	4
1.5 Justification of the Study	5
CHAPTER TWO	7
2.1 Overview of Internet of Things (IoT) Security Challenges	7
2.2 Solar-Powered Smart Cameras in IoT Networks	8
2.3 Federated Learning for Cybersecurity	12
2.4 Intrusion Detection Systems (IDS) in IoT Environments	17
2.5 Face Recognition Technologies for Attendance Monitoring	20

2.6	Privacy-Preserving Machine Learning in Educational Systems	23
2.7	DDoS Attacks on IoT Networks	25
CHAPTER THREE		29
3.1	Introduction	29
3.2	System Architecture	31
3.3	Integrate AI Models into Solar Smart Cameras	33
3.4	Ensure Low Power Processing for Solar Powered Smart Camera	41
3.5	Use Federated Learning to Improve Security Without Exposing Private Data	45
CHAPTER FOUR		48
4.1	Results	48
4.1.1	Preliminary Testing & Result Analysis	51
4.1.2	Classroom Deployment & Result Analysis	53
4.1.2.1	Day 4	54
4.1.2.2	Day 5	60
4.2	Discussion	63
CHAPTER FIVE		65
5.1	Conclusion	65
5.2	Recommendation	66
REFERENCES		70

TABLE OF FIGURES

Figure 1 - Architecture of SCMesh (Miller et al., 2015)	10
Figure 2 - Centralized Learning v. Federated Learning (Ferrag et al., 2021)	14
Figure 3 - IDS for IoT flow by Sicato et al. (2020)	19
Figure 4 - Visualization of a DDoS attack (Kumari & Jain, 2023)	26
Figure 5 - System Architecture	31
Figure 6 - Cameras & SD Cards	34
Figure 7 - Solar Camera Mounted on Classroom Wall	35
Figure 8 - Python Dependencies	37
Figure 9 - Intrusion Detection Pie Chart	49
Figure 10 - Average Power Used	50

TABLE OF TABLES

Table 1 - 4-Week Intrusion Detection Summary	48
Table 2 - Preliminary Results	52
Table 3 - Day 5: Activity Log	60

CHAPTER ONE

INTRODUCTION

1.1 Background of the Study

In recent years, the use of Internet of Things (IoT) devices has become more common across various sectors, especially for monitoring, surveillance, and automation. One stand-out device is the solar-powered smart camera. These cameras can operate in remote areas, but they also bring new security challenges. Due to the fact that they are constantly connected to networks and often installed in exposed environments, they are easy targets for cybercriminals.

Two of the most common threats faced by such devices are unauthorized access and Distributed Denial of Service (DDoS) attacks. Unauthorized access can allow attackers to steal data or control the camera, while DDoS attacks can overwhelm the network and make the system unusable. Traditional security methods, like centralized intrusion detection systems (IDS), are not effective in these situations because they require sending large amounts of data to the cloud. This creates problems like high latency, increased bandwidth usage, and risks to user privacy.

To address this, Federated Learning (FL) has been introduced as a promising solution. FL allows smart devices to learn collectively without sharing raw data. This means that each device can help build a strong model by only sharing small updates, keeping user data private. When combined with deep learning models like Convolutional Neural Networks (CNNs) and Long Short-Term Memory networks (LSTMs), FL can be used to detect

unusual patterns in network traffic that might indicate a cyberattack. Previous research has shown that this approach is effective in protecting IoT networks (Kairouz et al., 2021; Belarbi et al., 2023; Sun et al., 2024).

This project extends that idea by buying, installing, and programming solar-powered smart cameras for use in the Department of Computer Engineering. The cameras will not only act as security devices but also automate attendance monitoring in classrooms. Specifically, they will be deployed in the CPE 300 Level and 500 Level classes to grant access only to approved personnel (students and lecturers) while flagging unauthorized intrusions. By integrating federated learning, face recognition, and timetable data, the system will function as a combined attendance and intrusion detection solution that is secure, privacy-preserving, and scalable.

1.2 Problem Statement

As IoT devices become more widely used, especially in areas where constant surveillance or remote monitoring is needed, security becomes a bigger concern. Devices like smart cameras are often deployed in places where traditional or more conventional IT security infrastructure is not available. Their reliance on solar power and having limited processing ability can make it vulnerable to attacks that exploit their weaknesses.

Unauthorized access could allow hackers to spy on users or tamper with footage, while DDoS attacks could shut down entire networks of smart cameras. Unfortunately, most existing IDS systems are built for large networks with strong infrastructure. They don't work well in energy-limited environments and often rely on sending data to the cloud, which increases privacy risks.

Furthermore, in academic environments, classroom attendance is often taken manually, which is time-consuming, error-prone, and vulnerable to impersonation. There is also a lack of automated systems that integrate attendance management with real-time security monitoring.

There's a need for a solution that:

- Protects these devices in real time,
- Works without exposing user data,
- Doesn't depend on centralized servers,
- Can run efficiently on low-power solar devices like smart cameras, and
- Automates attendance and intrusion detection in classrooms.

This research aims to fill that gap using a federated deep learning approach.

1.3 Aim and Objectives

1.3.1 Aim

To develop a smart, lightweight, federated deep learning system that uses solar-powered cameras to both detect and prevent cyber threats (unauthorized access and DDoS attacks) and automate attendance and security monitoring for CPE 300L and 500L students and lecturers.

1.3.2 Objectives

- To use datasets to train and evaluate the IDS.
- To design a federated learning architecture that works efficiently on IoT devices like solar-powered smart cameras.
- To integrate CNN and LSTM deep learning models for intrusion detection in IoT networks.
- To implement face recognition for real-time attendance automation and access control.
- To deploy the IDS and attendance modules on smart cameras using lightweight frameworks such as TensorFlow Lite.
- To evaluate system performance in terms of accuracy, speed, energy use, and false alarms.
- To compare the federated learning approach with traditional centralized IDS methods.

1.4 Scope of Work

This study is limited to the design, simulation, and partial deployment of a federated deep learning-based intrusion detection and attendance monitoring system for IoT networks. The system will:

- Be implemented on solar-powered smart cameras in CPE 300L and 500L classrooms.
- Use CNN and LSTM models within a federated learning framework to detect unauthorized access and cyberattacks.
- Apply computer vision for student and lecturer face recognition.
- Automate attendance logging using CSV-based records (attendance.csv, lecturers.csv, timetable.csv).
- Be tested using existing datasets (e.g., NSL-KDD, CICIDS2017) for IDS and real-time classroom images for face recognition.
- Operate in a simulated and real classroom environment with the potential for future scale-up.

1.5 Justification of the Study

IoT devices such as solar-powered cameras are highly vulnerable to cyber threats, and traditional centralized IDS approaches are not suitable for low power environments due to high latency, privacy risks, and energy constraints. At the same time, academic institutions struggle with manual attendance tracking, which is inefficient and insecure.

This study is justified because it combines federated learning, intrusion detection, and face recognition into one integrated system. By doing so, it provides:

- Real-time cybersecurity for IoT-based smart cameras,

- Automated, privacy-preserving attendance tracking for classrooms,
- Energy-efficient deployment using solar-powered devices, and
- Scalability to support multiple levels and courses in the future.

The project will therefore contribute both academically (by demonstrating federated IDS on constrained devices) and practically (by improving classroom management and security).

CHAPTER TWO

LITERATURE REVIEW

2.1 Overview of Internet of Things (IoT) Security Challenges

The growth of the Internet of Things (IoT) has brought many benefits to industries such as agriculture, healthcare, and security. However, this expansion has also increased vulnerability to cyber threats, particularly in devices like solar-powered smart cameras that are constantly connected to the internet and often deployed in remote environments. Two of the most critical threats faced by such devices are unauthorized access and Distributed Denial of Service (DDoS) attacks.

Traditional Intrusion Detection Systems (IDS) often rely on centralized data processing, where raw data must be transmitted to a central server for analysis. While effective in high-resource environments, this approach introduces challenges such as high latency, increased energy consumption, and privacy concerns—limitations that make centralized systems unsuitable for low-power solar devices.

To address these issues, researchers have increasingly turned to federated learning (FL) as a privacy-preserving and resource-efficient alternative. FL enables smart devices to collaboratively learn security models without sharing raw data, thereby reducing communication overhead and supporting real-time threat detection (Kairouz et al., 2021).

For instance, Belarbi et al. (2023) developed a federated deep learning-based IDS using Convolutional Neural Networks (CNNs) to detect malicious traffic in IoT systems. Their approach demonstrated improved detection accuracy while preserving data privacy.

However, their study focused on general IoT devices and did not evaluate performance on constrained solar-powered hardware. Similarly, Sun et al. (2024) proposed a federated learning system called FedMADE, which grouped similar IoT devices to enhance collaborative training. Although the system improved accuracy in simulations, it was not tested on real-world solar devices. Rahmati (2025) also introduced a federated framework for IoT cybersecurity, which showed effectiveness in detecting threats such as DDoS attacks but did not target visual monitoring devices like smart cameras.

Building on these works, the present study narrows the scope to solar-powered smart cameras, which pose unique challenges due to energy and deployment constraints. Unlike previous research, this project aims to optimize federated deep learning models for lightweight deployment on real-world devices using tools like TensorFlow Lite. In doing so, it addresses the gap between simulated research and practical, field-ready solutions for IoT security.

2.2 Solar-Powered Smart Cameras in IoT Networks

The introduction of solar-powered smart cameras into Internet of Things (IoT) networks has expanded the scope of wireless sensing, particularly for autonomous surveillance and monitoring in remote or infrastructure-limited areas. Unlike scalar sensor networks that capture low-dimensional data such as temperature or humidity, smart cameras generate multimedia-rich data streams that significantly increase processing, storage, and communication demands. Due to this increased demand, power management becomes a central design challenge, particularly since these nodes are expected to function continuously under conservative energy constraints (Miller et al., 2015).

One notable advancement in this area is the SlugCam platform, which integrates multiple strategies. These include solar energy harvesting, rechargeable batteries, and adaptive duty cycling. They achieve long-term sustainability for wireless camera networks. The platform is specifically designed to optimize energy consumption by allowing each node to dynamically adjust its operation according to the amount of energy available at any given time. For instance, nodes can scale back sensing or processing tasks when energy levels are low, and resume full functionality once sufficient power is harvested. In addition, motion sensors embedded in the system ensure that the camera activates only when activity is detected in its field of view, thereby avoiding the constant energy drain associated with continuous operation. This combination of energy-aware hardware and intelligent scheduling not only reduces unnecessary energy expenditure but also demonstrates the potential of fine-grained power management strategies to extend the operational lifetime of sensor deployments in challenging outdoor environments. Despite its effectiveness in addressing energy autonomy, however, the SlugCam platform pays comparatively little attention to security considerations. Potential threats, such as unauthorized access or tampering with nodes, could compromise these energy-saving strategies and undermine overall system reliability (Miller et al., 2015). The platform's focus on sustainability, while innovative, highlights the need for complementary security mechanisms that can protect both the hardware and the collected data without significantly increasing energy consumption.




Figure 1 - Architecture of SCMesh (Miller et al., 2015)

In terms of connectivity, the SCmesh architecture developed as part of Miller et al.'s SlugCam initiative represents a significant advancement in enabling collaboration among camera nodes within a wireless network. Unlike earlier wireless sensor approaches that primarily relied on one-way communication to a central sink, SCmesh supports multipath routing, allowing data to traverse multiple routes to reach its destination. This not only improves network reliability in the event of node failures but also helps distribute energy consumption more evenly across the network, preventing early depletion of heavily used nodes. Furthermore, SCmesh incorporates power-awareness into its communication

protocols, enabling nodes to make routing decisions that consider their remaining energy levels, which is crucial for prolonging the operational lifetime of the network. The architecture also supports flexible transmission patterns, including direct node-to-node communication and broadcast messages, thereby enhancing the system's overall responsiveness and adaptability. Despite these strengths, SCmesh remains largely performance-driven, with relatively limited focus on securing transmitted data. This leaves potential vulnerabilities, as intercepted or manipulated messages could disrupt the network or compromise the integrity of collected data. Consequently, while SCmesh demonstrates a well-balanced approach to energy-efficient connectivity, there remains a pressing need for integrated security measures that do not undermine its performance-oriented design (Miller et al., 2015).

Local processing capabilities have also been highlighted as a means of improving efficiency. By incorporating on-board computer vision algorithms, solar-powered smart cameras can filter and analyze data before transmission, thereby conserving bandwidth and ensuring that only relevant events trigger communication. Abas et al. (2018) argue that such localized intelligence aligns with broader IoT trends in solar computing, as it shifts the processing burden away from centralized servers. Yet, while these strategies improve scalability and responsiveness, they still raise questions about how secure the locally processed data remains, particularly in adversarial environments where attackers might target the node's decision-making processes.

These innovations highlight both the promise and the challenges of deploying solar-powered smart cameras in IoT environments. They show that renewable energy harvesting,

adaptive duty cycling, and multipath routing can extend system lifetimes and improve communication reliability. However, the literature also reveals gaps—most notably the lack of comprehensive mechanisms to address cybersecurity threats. For instance, Distributed Denial of Service (DDoS) attacks could quickly overwhelm resource-constrained nodes, while unauthorized access could compromise the integrity of collected visual data. These vulnerabilities suggest that while power management and connectivity have advanced, the dimension of security remains insufficiently explored (Abas et al., 2018; Miller et al., 2015).

To address these gaps, emerging studies are turning to federated deep learning-based intrusion detection systems (IDS) as a promising solution. Federated learning enables nodes to collaboratively train security models without transmitting raw data, thereby preserving privacy and reducing bandwidth consumption. This approach directly complements the design of platforms such as SCmesh, which already emphasize decentralized collaboration and energy efficiency. As Abas et al. (2018) note, integrating localized intelligence with distributed coordination is central to the future of IoT networks. Extending this principle with federated IDS frameworks would ensure that solar-powered smart cameras are not only energy-resilient and bandwidth-efficient but also robust against cyber threats that could otherwise undermine their effectiveness.

2.3 Federated Learning for Cybersecurity

Federated learning (FL) is an emerging decentralized machine learning paradigm in which multiple devices collaboratively train a shared global model while retaining all local data on the individual devices. Unlike traditional centralized machine learning, where raw data

from each device must be transmitted to a central server for processing and model training, FL only requires the exchange of model updates, such as gradients or model parameters. This fundamental difference significantly minimizes the risk of privacy leakage and reduces the network burden associated with transferring large volumes of potentially sensitive data across often unreliable or bandwidth-constrained connections (Ferrag et al., 2021). In addition to its privacy-preserving benefits, FL is highly adaptable to the challenges posed by the heterogeneity and scale of data generated by distributed Internet of Things (IoT) devices, which typically vary in format, quality, and frequency of generation. Centralized learning approaches often struggle in such environments because they require aggregating all raw data in one location, leading to potential bottlenecks, high energy consumption, and vulnerability to data breaches. By contrast, FL is inherently more resilient, scalable, and energy-efficient, making it particularly well-suited for IoT systems with constrained resources, such as solar-powered smart cameras or other battery-limited sensor nodes. The decentralized nature of FL allows each device to perform local computation independently, reducing the need for continuous communication with a central server and thereby lowering energy costs. Moreover, because only aggregated model updates are shared, the approach mitigates privacy risks while still enabling the system to learn from a diverse set of observations across the network.




Figure 2 - Centralized Learning v. Federated Learning (Ferrag et al., 2021)

The exponential growth of IoT networks, coupled with their distributed and resource-constrained nature, has amplified security challenges such as unauthorized access, malware, and Distributed Denial of Service (DDoS) attacks. Traditional centralized machine learning approaches, which require aggregating raw data in cloud servers, raise serious privacy risks and create communication overheads unsuitable for low-power IoT devices. Federated learning (FL) has emerged as a promising paradigm to address these issues by enabling decentralized training where devices collaboratively learn global models without sharing sensitive local data (Ferrag et al., 2021).

FL introduces a privacy-preserving layer of intelligence to IoT systems by allowing solar devices, such as solar-powered smart cameras, to train models on locally observed threats

while contributing only model updates to the federation. This approach significantly reduces the risk of data exposure and aligns with privacy requirements in academic environments where student biometric data may be collected for attendance automation. According to Gugueoth et al. (2023), FL's capacity to mitigate privacy leakage while maintaining detection accuracy makes it a superior choice over conventional centralized learning in sensitive IoT deployments.

In cybersecurity applications, FL has demonstrated effectiveness against a wide spectrum of attacks, including DoS, DDoS, botnets, and spoofing. Ferrag et al. (2021) emphasize that distributed models are better positioned to capture diverse attack patterns because they aggregate knowledge from heterogeneous devices deployed across varied environments. For solar smart cameras, this is particularly relevant, as the system must defend against both network-based intrusions and physical tampering attempts in real-time classroom monitoring scenarios.

Federated learning improves scalability in large IoT ecosystems. As Gugueoth et al. (2023) note, IoT networks are characterized by heterogeneity and high-volume data streams, which complicate centralized security management. By distributing computation across nodes, FL alleviates the computational and communication bottlenecks of cloud-centric models. In the context of solar-powered cameras, this decentralized processing also translates into energy savings, since less bandwidth is consumed transmitting high-volume video or log data to remote servers.

However, FL is not without challenges. Ferrag et al. (2021) highlight vulnerabilities such as poisoning attacks, where compromised nodes send manipulated model updates to

corrupt the global model. Similarly, Gugueoth et al. (2023) identify issues of high communication overhead in environments with intermittent connectivity, which is especially relevant to solar-powered devices operating in low-resource or rural academic settings. Addressing these limitations requires integrating robust aggregation mechanisms, such as anomaly detection filters on the server side, and lightweight update compression techniques to conserve energy and bandwidth.

Irrespective of the challenges faced, both studies point toward hybrid approaches that combine FL with deep learning architectures like Convolutional Neural Networks (CNNs) and Long Short-Term Memory (LSTM) networks. Ferrag et al. (2021) demonstrate that these combinations achieve high detection accuracy in intrusion detection systems (IDS), while Gugueoth et al. (2023) show their promise in securing IoT services with minimal latency. For smart classroom applications, CNNs can extract spatial features from video streams for face recognition, while LSTMs capture temporal patterns of network intrusions, making them ideal candidates for federated deployment on smart cameras.

FL represents a critical advancement in IoT cybersecurity by providing privacy-preserving, scalable, and energy-efficient intrusion detection. Its integration with deep learning enhances the resilience of IoT infrastructures against evolving threats, while its decentralized nature aligns with the operational constraints of solar-powered smart cameras in academic environments. By addressing both security and functionality, FL enables a dual-purpose system that safeguards IoT networks while supporting intelligent applications such as automated attendance, ultimately bridging the gap between robust cybersecurity

and practical deployment in constrained environments (Ferrag et al., 2021; Gugueoth et al., 2023).

2.4 Intrusion Detection Systems (IDS) in IoT Environments

The rapid spread of Internet of Things (IoT) devices has transformed domains ranging from smart homes to industrial automation, but this expansion has also introduced serious security challenges. Unlike traditional computing platforms, IoT devices such as solar-powered smart cameras are constrained by limited memory, processing power, and energy. Anitha & Arockiam (2022) highlight that these constraints render conventional centralized security mechanisms unsuitable for IoT environments. At the same time, the sheer number of devices and their diverse communication protocols expand the attack surface, creating an urgent need for specialized security frameworks. Intrusion Detection Systems (IDS) have therefore become a crucial line of defense for monitoring IoT network traffic and identifying malicious activity.

For devices like solar-powered cameras deployed in classrooms, these constraints are even more pronounced. Anitha & Arockiam (2022) emphasize that lightweight security mechanisms are essential for such nodes. Traditional signature-based IDSs, which rely on databases of known attack patterns, can be accurate but are too resource-intensive for low-power IoT devices. Moreover, these systems fail to detect zero-day or novel attacks—an unacceptable limitation in dynamic environments where threats evolve rapidly. This underscores the importance of adaptable intrusion detection methods that can balance efficiency with robust security.

Anomaly-based IDSs provide a promising alternative. Instead of depending on known attack signatures, they construct a profile of normal network behavior and classify deviations from this baseline as potential threats. As Bhavsar et al. (2023) note, anomaly-based systems are particularly effective at identifying new or unknown attacks, including those that target IoT networks through distributed denial of service (DDoS) or spoofing. Yet, their success depends heavily on the accuracy of the baseline model, which can be difficult to establish in highly dynamic and heterogeneous IoT environments like classrooms, where network usage patterns fluctuate.

To address this limitation, researchers have turned to advanced approaches based on deep learning. Bhavsar et al. (2023) introduce PCC-CNN, a model that integrates feature selection with convolutional neural networks for both binary and multi-class anomaly detection. Deep learning models excel at analyzing large datasets and recognizing complex, non-linear relationships, enabling them to distinguish between benign and malicious traffic with high accuracy. Their ability to adapt through continuous learning makes them well suited for next-generation IDSs capable of keeping pace with sophisticated cyberattacks.

IoT's distributed nature creates challenges for deploying deep learning in practice. Centralized architectures that require transmitting all raw data to a remote cloud server for training introduce latency, bandwidth costs, and single points of failure. Sicato et al. (2020) argue for distributed approaches to mitigate these limitations, yet transmitting large volumes of traffic data remains impractical for solar-powered cameras, which often operate under intermittent or low-bandwidth conditions. This makes traditional centralized deep learning unsuitable for IoT security in real-world deployments.




Figure 3 - IDS for IoT flow by Sicato et al. (2020)

Federated Learning (FL) offers a transformative solution to these challenges. In FL, each device trains its own local model on the data it collects and only shares model updates—such as parameter weights—with a central aggregator. The raw data never leaves the device, significantly reducing privacy risks. For solar-powered smart cameras, this means sensitive classroom video and network data can remain on-device while still contributing to the collective training process. This decentralization not only enhances privacy but also reduces communication overhead, which is critical for energy-efficient devices.

Applying FL to deep learning-based IDSs unlocks significant advantages for IoT classroom security. Each smart camera can continuously train its anomaly detection model on local network traffic, adapting to the unique conditions of its environment. Periodically, these devices share model updates with a central server, which aggregates them into a stronger global model and redistributes it back to the network. In this way, every device benefits

from the combined intelligence of the entire system while preserving its own privacy and conserving bandwidth.

The benefits of this federated deep learning IDS framework are multifaceted. First, it maintains data privacy by keeping raw video and traffic data on the device. Second, it is bandwidth- and energy-efficient, transmitting only lightweight model updates rather than heavy datasets. Third, it allows the IDS to remain dynamic—constantly learning from new attack patterns across the network and adapting in near real time. This ensures that classroom monitoring systems remain resilient against both common and emerging threats without overburdening the constrained hardware of solar-powered cameras.

The combination of anomaly-based detection, deep learning, and federated learning provides a pathway toward scalable and effective IoT security. The unique resource constraints of solar smart cameras make conventional IDSs impractical, but by integrating FL with deep learning architectures, a robust solution can be deployed directly on the devices themselves. This approach not only strengthens defenses against evolving cyberattacks but also preserves privacy and ensures sustainable operation in academic environments, where both reliability and data protection are paramount.

2.5 Face Recognition Technologies for Attendance Monitoring

The manual process of tracking attendance in academic institutions and corporate environments has long been recognized as a time-consuming and often inaccurate task. As highlighted by Nadhan, et al. (2022), this traditional method frequently consumes a significant portion of productive time. With the advent of the Internet of Things (IoT) and advancements in computer vision, automated solutions have emerged to streamline this

process, with face recognition technology at the forefront. This technology leverages biometric data to uniquely identify individuals, offering a more efficient, reliable, and secure alternative to traditional methods.

The core of a face recognition-based attendance system relies on a combination of hardware and software. The hardware typically includes a camera—often an IoT device itself, such as a smart camera—that captures real-time video or images of the subjects. The software, as described by Bavaskar (2024), uses various algorithms and machine learning models to detect, recognize, and verify faces. This process generally involves several key steps, including face detection to locate faces in an image, feature extraction to identify unique facial landmarks, and face verification to match the extracted features against a pre-existing database of registered individuals.

One of the significant advantages of using face recognition for attendance monitoring is its high accuracy and efficiency. Nadhan, et al. (2022) point out that an effective system using an open-source image processing library like OpenCV can reduce product cost and aid in connecting to heterogeneous devices. This automation eliminates human error and the need for physical attendance registers or ID cards, providing a contactless and seamless experience for users. The system can log attendance in a matter of seconds, making it far more productive than manual roll calls, especially in large classes or offices.

The implementation of face recognition technology comes with some challenges. Accuracy can be compromised by various environmental factors, such as poor lighting, changes in a person's appearance (e.g., beards, glasses), or the angle at which the face is captured. As noted by Shashikala, et al. (2022), the system's performance can be influenced by the

quality of the camera and the robustness of the algorithms used. Ensuring the system works reliably across different conditions and with diverse populations remains a key technical hurdle for researchers and developers.

Privacy and security are paramount concerns in the deployment of these systems. Although this will be covered in detail in the next section, it will be touched briefly in this section. As the technology deals with sensitive biometric data, there is a risk of unauthorized access or misuse. Bavaskar (2024) and others emphasize the need for robust security measures, including data encryption and secure database management, to protect against potential breaches. The ethical implications of using surveillance technology also need to be considered, and it is crucial to have clear policies and consent from users regarding the collection and use of their facial data.

From an implementation perspective, various hardware platforms can be used. Shashikala, et al. (2022) discuss a system built on a Raspberry Pi, which provides a cost-effective solution for creating an automated attendance system. Such a setup can be integrated with other IoT components to create a comprehensive smart environment. The data collected by the system can be stored locally on the device or transmitted to a central cloud server, depending on the architecture and privacy requirements of the specific application.

Face recognition technology offers a compelling solution for modernizing attendance monitoring by providing an automated, accurate, and efficient system. The integration of such systems with smart cameras aligns with the broader trend of leveraging IoT for practical applications. While the technology presents significant advantages, developers must address the critical challenges of accuracy under varying conditions, as well as the

fundamental issues of data privacy and security, to ensure its successful and ethical deployment.

2.6 Privacy-Preserving Machine Learning in Educational Systems

The introduction of IoT-enabled technologies such as solar-powered smart cameras has transformed the educational sector by enabling automated attendance systems, classroom monitoring, and enhanced security. These devices often process sensitive student information, including facial images and behavioral data, which raises substantial privacy concerns. Centralized machine learning approaches that aggregate raw data from classrooms into a single server are particularly risky, as they increase exposure to breaches and compromise student confidentiality. Ferrag et al. (2021) stress that such centralized models are not only vulnerable to attacks but also unsuitable for resource-constrained devices like solar cameras, which must operate efficiently on renewable energy.

Federated learning (FL) addresses these concerns by training models locally on each camera or IoT node while sharing only model updates with a central server. This ensures that sensitive classroom data—such as facial recognition inputs for attendance—never leaves the device, significantly reducing the risk of privacy leakage. Gugueoth et al. (2023) emphasize that FL is well-suited for distributed and dynamic environments, where devices must preserve privacy while contributing to collective intelligence. For solar-powered smart cameras in classrooms, FL therefore provides a privacy-preserving backbone for machine learning, allowing real-time IDS and attendance tracking without compromising student data.

This leads to adaptability in educational contexts. Classroom environments are not static, with fluctuating student attendance, varying schedules, and diverse behavioral patterns. By enabling each smart camera to train on its local environment, FL ensures models remain responsive to unique conditions while contributing updates to a global model that generalizes across multiple classrooms. As Ferrag et al. (2021) point out, combining FL with deep learning architectures such as CNNs and LSTMs provides a powerful framework for anomaly detection and face recognition. This approach is particularly relevant for IDS-enabled solar cameras, which must simultaneously detect cyber intrusions and automate attendance in a resource-efficient manner.

FL systems remain vulnerable to adversarial threats such as data poisoning or model inversion, where malicious actors attempt to infer sensitive information from shared updates (Ferrag et al., 2021). In addition, device heterogeneity—such as solar-powered cameras with varying battery levels and connectivity—can create inconsistencies in training contributions. Gugueoth et al. (2023) highlight that robust aggregation methods, differential privacy, and secure encryption are necessary to mitigate these risks, though they may introduce additional computational costs. Achieving an optimal balance between privacy, accuracy, and energy efficiency remains an open challenge for educational IoT deployments.

Privacy-preserving machine learning, particularly through federated learning, provides a crucial foundation for secure and intelligent educational systems. For solar-powered smart cameras tasked with both classroom security and attendance automation, FL ensures sensitive student data stays on-device while still enabling collaborative improvements

across the network. By merging FL with deep learning for intrusion detection and facial recognition, schools can deploy a scalable and energy-efficient framework that addresses the dual priorities of privacy and security. This makes federated deep learning a cornerstone of next-generation educational technologies, ensuring that smart classrooms remain both safe and respectful of student confidentiality.

2.7 DDoS Attacks on IoT Networks

The rapid expansion of the Internet of Things (IoT) has given rise to unprecedented connectivity, but it has also created a vast and vulnerable attack surface, making it susceptible to a growing number of cyber threats. Among the most potent and disruptive of these threats are Distributed Denial of Service (DDoS) attacks. Kumari & Jain (2023) define a DDoS attack as "an assault that targets the availability of resources and servers of a network by flooding the communication medium from distinct locations by utilizing various IoT devices, which makes it harder to detect." These attacks effectively overwhelm a target system with an immense volume of malicious traffic, rendering it unavailable to legitimate users. The scale and diversity of IoT devices make DDoS attacks increasingly difficult to mitigate, as compromised devices can be geographically dispersed and operate under different network protocols. Moreover, many IoT devices are resource-constrained and lack robust security measures, making them easy targets for attackers to co-opt into botnets. The consequences of such attacks are not limited to service disruption; they can also result in financial losses, data breaches, and erosion of user trust in connected systems. Traditional defense mechanisms, such as firewalls and intrusion detection systems, often struggle to keep pace with the dynamic and distributed nature of IoT-driven DDoS attacks. Consequently, there is a pressing need for advanced, scalable, and adaptive security

solutions that can detect and mitigate these attacks in real time, without compromising the functionality or efficiency of IoT networks.




Figure 4 - Visualization of a DDoS attack (Kumari & Jain, 2023)

The ingrained characteristics of IoT devices make them particularly susceptible to DDoS attacks. Devices such as smart cameras, sensors, and actuators are often resource-constrained, possessing limited processing power, memory, and energy. Furthermore, many IoT devices are deployed with weak security features, default credentials, and unpatched firmware, making them easy targets for exploitation. Attackers can exploit these vulnerabilities to compromise a large number of devices and form them into a "botnet" or "zombie army," which is then used to launch large-scale DDoS attacks (Kumari & Jain, 2023).

According to Kumari & Jain (2023), DDoS attacks can be broadly categorized into several variants, each exploiting different vulnerabilities. Volumetric-based attacks, for instance, aim to consume all available bandwidth by flooding the network with a massive amount of traffic. Examples of this type of attack include UDP floods, ICMP floods, and

amplification attacks. These attacks are often measured in gigabits per second (Gb/s) and are designed to simply overwhelm the target's network capacity.

Another major category is protocol-based attacks, which target the weaknesses of the network's protocol stack, particularly at layers 3 and 4. These attacks, often referred to as "resource depletion attacks," consume the resources of the target server and network equipment, such as firewalls. A classic example is the SYN flood attack, which exploits the TCP three-way handshake process to exhaust the target server's connection state table, preventing it from accepting new legitimate connections.

To combat these threats, researchers have explored a range of countermeasures. Traditional solutions often involve network filtering, firewalls, and signature-based IDSs. While these methods can be effective against known threats, they are reactive and often fail to protect against zero-day attacks or evolving DDoS variants. The decentralized and diverse nature of IoT networks further complicates the implementation of these traditional, centralized defenses.

In response to the limitations of traditional methods, machine learning (ML) and deep learning (DL) have emerged as powerful tools for DDoS detection. Kumari & Jain (2023) highlight the effectiveness of these advanced techniques in analyzing network traffic and identifying subtle, anomalous patterns that may indicate an attack. Unlike signature-based systems, ML and DL models can learn from large datasets to recognize the characteristics of both known and unknown attacks, making them a more proactive and adaptable solution.

While deep learning offers a broad approach, a purely centralized model for IoT-based DDoS detection faces its own set of challenges. Sending massive volumes of network traffic data from thousands of solar smart cameras to a central server for analysis is not only a privacy concern but also a logistical nightmare. The bandwidth and energy consumption would be prohibitive, making the solution impractical. This is where the concept of federated deep learning becomes a compelling paradigm.

A federated deep learning-based IDS offers a decentralized solution that is perfectly suited for a network of solar smart cameras. Instead of transmitting raw data, each camera would train a local deep learning model to detect DDoS traffic. Only the model's updated parameters would be sent to a central server for aggregation. This approach ensures data privacy, as sensitive network traffic never leaves the device. Furthermore, it significantly reduces bandwidth usage and can leverage the collective intelligence of the entire network to create a more resilient and sophisticated global model for DDoS detection.

CHAPTER THREE

METHODOLOGY

3.1 Introduction

In this project, we developed a Federated Deep Learning-Based Intrusion Detection System designed to secure solar-powered smart cameras used for attendance and campus security monitoring. We aimed to build a system that combines artificial intelligence (AI), federated learning, and cybersecurity to detect and prevent threats while maintaining data privacy and energy efficiency. Our approach ensures that all smart cameras can process data locally without needing constant internet access, making them faster, safer, and more reliable.

We began by integrating AI models directly into purchased solar smart cameras, rather than building new ones from scratch. The cameras came with solar panels. While the camera was installed inside the classroom, the panel was mounted on the wall outside to access sunlight. To write, test, and debug our AI and intrusion detection algorithms, we used Visual Studio Code (VS Code) as our primary development environment.

We trained lightweight AI models capable of performing facial recognition and activity detection directly on the cameras. Instead of collecting photos from students and lecturers, we performed live face captures during attendance and monitoring sessions. This approach provided stronger security, as it prevented impersonation or the use of outdated photos. Live capturing ensured that only real-time, authentic faces were recognized, reducing the risk of data tampering or identity fraud.

To maintain efficiency, we used low-power AI algorithms like MobileNetV3 and Tiny-YOLO, which are optimized for limited hardware. We made sure the models could perform real-time facial recognition and security analysis without consuming much energy. By doing so, we ensured that the solar-powered cameras could operate continuously and reliably under low-energy conditions.

Each camera served as a local node in a Federated Learning (FL) network. In this system, every camera trained on its own locally collected data, such as attendance records or motion detection results, and then sent only model updates to a central server. We did not transmit any raw data, which preserved user privacy and ensured compliance with data protection principles. The central server aggregated these updates to improve the global model, allowing all cameras to benefit from collective learning without sharing sensitive information.

For cybersecurity, we included a deep learning-based intrusion detection module capable of identifying unauthorized access and other abnormal behaviors. However, we determined that Distributed Denial of Service (DDoS) attacks were not a major risk because the system operates on local servers rather than public internet infrastructure. This local-only design minimizes the exposure of devices to external cyber threats and makes the network more secure.

To further strengthen security, we implemented encryption and secure communication protocols for all data exchanges between the smart cameras and the central federated server. We also scheduled periodic model updates to ensure that the system remains adaptive to new patterns or threats over time. Throughout development, we focused on optimizing the

system for speed, accuracy, and power efficiency to meet the constraints of solar-powered IoT devices.

We evaluated the system using key performance metrics such as accuracy, latency, power consumption, and detection rate. By integrating AI and federated learning, we created a scalable and eco-friendly solution suitable for educational and remote environments.

3.2 System Architecture




Figure 5 - System Architecture

The system architecture for the Digital Facial Recognition Attendance System is a powerful example of a Client-Server model, though its primary operational flow is localized to ensure off-grid autonomy and data privacy. It is fundamentally centered around the GZ SONY Intelligent Solar Energy Alert PTZ with 5 MP Ultra HD resolution camera itself, which serves as the core processing and data hub. The camera unit integrates all necessary components: the camera sensor for image capture, a dedicated AI Processing Chip (NPU/DSP) for low-power inference, local SQLite3 database storage on an SD card for records, and a Local Web Server (Flask/Python) to manage the logic and user interface. This tight integration ensures that the most power-intensive task—running the facial recognition model—occurs directly at the source.

The data flow begins with the camera's sensor capturing real-time frames. These frames are immediately fed into optimized models like MobileNetV3 or Tiny-YOLO perform face detection and embedding extraction (feature generation). This minimizes power consumption by using techniques like quantization and leveraging the specialized NPU/DSP hardware. The resulting face embeddings are then passed to the Business Logic & API component of the Local Web Server, which compares the embedding against the local SQLite3 database records to identify the individual, log the attendance, and save the timestamp.

The system's user interaction is managed by the Local Web Server, which hosts the Admin/Lecturer Dashboard over a Local Network (Wi-Fi/Ethernet). Administrators or lecturers access the dashboard using a standard device (like a laptop or phone) to perform user management, view reports, and configure settings. The server utilizes HTML, CSS, and JavaScript for the UI, dynamically fetching stored attendance records and recognition

logs from the local SQLite3 database. Crucially, all these operations—recognition, storage, and web-serving—are powered by the solar panel, providing a continuous DC supply and enabling the entire system to function autonomously off-grid.

The architecture is designed for data privacy and off-grid resilience. The centralized camera unit acts as a standalone intelligent system, only using the local network to present the monitoring dashboard and, optionally, connect to a gateway for external notifications (Email/SMS).

3.3 Integrate AI Models into Solar Smart Cameras

In this part of the project, we focused on integrating artificial intelligence (AI) models into solar-powered smart cameras to enable real-time recognition and intelligent monitoring. We aimed to make the cameras capable of recognizing faces, marking attendance, and identifying unauthorized persons without relying on internet connectivity. The AI component was primarily used to power the facial recognition algorithm, which allowed the system to automatically identify students and lecturers in real time. Our approach was designed to be efficient, private, and sustainable.



Figure 6 - Cameras & SD Cards

We used two GZ SONY Intelligent Solar Energy Alert PTZ Cameras with 5 MP Ultra HD resolution for the project. These cameras were purchased, not custom-built, as they already contained advanced hardware features suitable for AI integration. Each camera came with pan-tilt-zoom (PTZ) capabilities, which allowed wide-area coverage and flexible viewing angles. Along with the cameras, we also purchased two 64 GB SD cards—one for each camera—to serve as local storage for captured images, face embeddings, and log files. The SD cards allowed the cameras to store data locally, ensuring that all captured information remained secure on the device without depending on cloud services.



Figure 7 - Solar Camera Mounted on Classroom Wall

The cameras were mounted at the back of the classroom with an unobstructed view of the door to properly be able to scan the faces of those seeking entry to the class. The panel was

mounted outside the classroom and a power cable seen in the picture was used to connect them through a hole in the wall.

The system is designed to run on a solar-powered device, where all major operations—such as face detection, embedding generation, model prediction, and attendance logging—are processed locally using the Flask web framework. Instead of relying on a remote centralized server, the system performs computation directly on the device (e.g., a smart camera or Raspberry Pi) to ensure real-time processing, data privacy, and low power consumption consistent with computing principles. We focused on integrating computer vision (CV) and machine learning (ML) components into a web-based application to enable real-time facial recognition and automated attendance logging.


For development, we used Visual Studio Code (VS Code) as our primary programming environment. VS Code supported all the programming languages we used, including Python, HTML/CSS, and JavaScript, allowing seamless integration of the code components. Python was used for developing and implementing the CV recognition model, as it supports essential libraries such as OpenCV and scikit-learn. The facial recognition algorithm was trained using RandomForestClassifier on embeddings generated by MediaPipe Face Detection to detect and identify faces accurately from captured frames. This model, saved as model.pkl [17], is loaded and executed by the Python backend.

The frontend interface was designed using HTML and styled with CSS to provide an intuitive dashboard for users such as administrators and lecturers. It displays attendance records, recognition logs, and model training status. We used JavaScript to handle all client-side interactions, including accessing the user's camera (using the WebRTC API),

periodically capturing frames, and sending these frames via AJAX/Fetch requests to the Flask backend. The backend, which is also a Python server, handles the subsequent processing, recognition, and database operations.

We designed the system architecture around a standard Client-Server model to support system coordination and communication. In this setup, the web browser (client) manages the camera interface and user dashboard, while the Flask server (backend) handles the heavier processing tasks such as face detection, embedding generation, model prediction, and attendance logging.

For the AI model selection, we used a combination of MediaPipe for face detection and a RandomForestClassifier trained on simple 32x32 grayscale face embeddings for identification. Our choice prioritized simplicity and compatibility with the scikit-learn ecosystem while still achieving reliable classification performance. The model is trained and optimized in Python on the server side.

A screenshot of a code editor window. The title bar shows 'Welcome' and 'app.py'. The editor content shows the file path 'Digital-Facial-Recognition-Attendance-System > app.py > init_db'. The code is as follows:

```
1 import os
2 import io
3 import threading
4 import sqlite3
5 import datetime
6 import json
7 from flask import Flask, render_template, request, jsonify, send_file, abort
8 from model import train_model_background, extract_embedding_for_image, MODEL_PATH
9
```

Figure 8 - Python Dependencies

The Python dependencies used in the project form the foundation of the system's functionality, combining web development, AI processing, and local data management within a lightweight and efficient framework. The most critical dependency is Flask, a Python micro-framework that powers the web server and application logic. Flask manages client requests, serves dynamic pages via `render_template` [17], and provides API endpoints for communication between the web interface and backend. The SQLite3 dependency provides a secure, file-based database (`attendance.db`) [17] that stores facial embeddings, recognition logs, and timestamps locally. This eliminates the need for an external or cloud-based database, thereby improving both performance and data privacy.

Another key aspect of the Python stack is concurrency and real-time operation, achieved through dependencies such as `threading`, which allows time-intensive operations like model training to run in the background without interrupting the Flask web server. The `datetime` module ensures that every recognition or attendance entry is time-stamped precisely, guaranteeing record accuracy and synchronization between the recognition model and the database. Additionally, `json` supports structured data exchange between the Flask backend and the JavaScript-based frontend, enabling fast and reliable updates to the user interface. These dependencies collectively ensure smooth system operation, even under multiple simultaneous user interactions or live recognition tasks.

The AI and facial recognition functions are modularized and imported directly from a separate file, `model.py` [17], which depends on machine learning libraries such as TensorFlow or OpenCV. The imported functions—like `train_model_background` and `extract_embedding_for_image`—bridge the gap between the web application and the AI engine. This modular dependency structure allows Flask to delegate computationally heavy

tasks, such as model training and facial embedding extraction, to specialized AI routines without blocking user interactions. Together, these Python dependencies enable a scalable, privacy-preserving, and low-power AI system optimized for solar-powered smart cameras.

We integrated the SQLite3 database to manage attendance records and student details. SQLite3 was chosen because it is a lightweight, file-based database easily integrated with the Python backend. The database file (attendance.db) [17] resides on the server running the Flask application. This structure centralizes data management and allows for unified reporting.

For the recognition process, we used live face captures from the user's camera stream. This means that the system captures real-time images during attendance sessions, sending these frames to the server for processing. We capture 50 images during the student registration phase to build the dataset for training, ensuring that only physically captured data is used to train the model.

The captured images for training are saved to the server's file system (dataset/) [17] and are used to extract facial embeddings (32x32 vectors). The facial embeddings are implicitly used by the model.pkl classifier [17]. The 32x32 resolution of the face images used for training (after cropping and resizing) is sufficient for the RandomForestClassifier to extract unique features for identification.

For the recognition process, we used live face captures instead of pre-collected photos. This means that the cameras captured real-time images of students and lecturers during attendance sessions. We avoided accepting uploaded or pre-sent photos to enhance system security and prevent impersonation. Live capture ensured that only physically present

individuals could be recorded or recognized by the system. This method provided better security and integrity of attendance data, as it was impossible to fake participation remotely.

The captured images were temporarily stored on the camera's SD card, while only the facial embeddings were stored permanently in the SQLite3 database. This design kept storage efficient and minimized data redundancy while maintaining high accuracy in recognition. The 5 MP Ultra HD resolution of the cameras allowed for clear image capture, which improved the accuracy of face detection and recognition. The AI models analyzed each frame from the live video feed and extracted unique facial features for identification. This process occurred locally on the device, ensuring that sensitive visual data never left the camera's storage environment.

The integration process involved linking all system components—AI model, database, backend, and frontend—through local server protocols. We ensured that data flow between these components was stable, secure, and optimized for processing. Encryption and secure local communication methods were used to protect data as it moved between the AI modules and the database. This step was important for maintaining the privacy of captured images and recognition records. The system architecture thus combined both AI intelligence and cybersecurity within a single operational unit.

Since all processing and storage were done locally, DDoS attacks were not a major threat to the system. The cameras and their servers operated on isolated local networks, preventing external internet-based attacks. This local-only structure made the cameras more resilient and reduced their vulnerability to remote hacking. It also ensured that sensitive student and staff data remained within the closed network environment of the institution. By combining

solar energy and local computing, the system provided both physical and digital independence from external threats.

3.4 Ensure Low Power Processing for Solar Powered Smart Camera

In this part of the project, we focused on achieving low-power AI processing for the devices—specifically, the solar-powered facial recognition cameras. Because the entire system depends on solar energy, every hardware and software component had to be designed to consume minimal power while still maintaining high performance and real-time responsiveness. Our goal was to make the cameras smart enough to process and recognize faces locally without relying on the cloud, yet efficient enough to operate reliably under limited energy conditions. Achieving this balance required optimization across three major layers: the AI model architecture, the data processing pipeline, and the hardware utilization strategy.

The first step toward power efficiency was selecting inherently lightweight model architectures. Instead of deploying large and power-hungry deep neural networks such as VGG or ResNet, we used models specifically designed for solar devices, like MobileNetV3 and Tiny-YOLO. These architectures rely on specialized techniques such as depthwise separable convolutions, which drastically reduce computational requirements and memory usage without significantly compromising accuracy. By using models with fewer parameters and smaller sizes—typically under 10 MB—we ensured that inference operations could be executed quickly, allowing the processor to spend less time in high-power states and conserve energy.

Once we selected suitable lightweight architectures, we implemented additional model optimization techniques to further reduce computation. Two key methods—quantization and pruning—were applied to minimize the model’s memory footprint and energy use. Quantization converted the model weights from 32-bit floating-point precision to 8-bit integers, reducing the model size by roughly four times and allowing the processor to perform simpler, faster arithmetic. Pruning further removed unnecessary or redundant neurons and connections within the network, resulting in a smaller and more efficient model that required less computational power to run. Together, these methods significantly reduced both inference time and energy consumption.

The next area of optimization focused on the data preprocessing pipeline, which often accounts for a large portion of total energy consumption. Instead of processing every frame from the camera feed, we implemented frame downsampling, analyzing fewer frames per second while maintaining sufficient accuracy for attendance and monitoring tasks. Additionally, the system used Region of Interest (ROI) detection, where a lightweight preprocessing model (MediaPipe) first detected potential face regions. Only the cropped face areas were then passed to the main recognition model, avoiding unnecessary processing on the entire frame and saving both power and computation time.

In addition to software-level optimization, we leveraged the dedicated hardware available on the cameras to further reduce power consumption. The GZ SONY Intelligent Solar Energy Alert PTZ Cameras used in the project are equipped with embedded processors capable of handling AI inference tasks efficiently. Some of these processors include specialized accelerators such as Neural Processing Units (NPUs) or Digital Signal Processors (DSPs), which are designed to perform matrix multiplication—the core

operation in deep learning—at much lower power levels than traditional CPUs. Offloading inference tasks to these dedicated accelerators ensured fast and efficient model execution without straining the system’s main processor or battery supply.

To make the entire system more energy-aware, we implemented power scheduling and duty cycling techniques. Duty cycling refers to a method where the camera periodically wakes up to perform recognition and then returns to a low-power state for a set interval. This prevents the AI model from running continuously when not needed. For example, the recognition process only activates when motion detection or a simple image difference algorithm signals that a person is in view. Such conditional activation helps conserve power while maintaining readiness for real-time facial recognition and intrusion detection.

Since the system is entirely solar-powered, power management was tightly linked to energy availability. Intensive AI computations such as retraining or batch processing were scheduled during periods of maximum sunlight, ensuring that tasks requiring higher power were completed when solar input was strongest. During nighttime or low-light conditions, the cameras relied on stored battery energy and performed only essential monitoring tasks. This intelligent scheduling maintained consistent operation without overloading the power system or draining energy reserves unnecessarily.

Another important factor in maintaining low power was minimizing data transfer between components. Network communication—especially wireless data transmission—is one of the most energy-demanding operations in IoT systems. All tasks, including image capture, recognition, and logging, were performed locally using SQLite3 and the camera’s 64 GB SD card for data storage. Only summarized attendance logs or security alerts were

occasionally transmitted, which drastically reduced the amount of power consumed through network I/O operations.

The combination of local processing and reduced communication also enhanced data privacy and security. Because no raw images or sensitive data were transmitted to external servers, the likelihood of interception or unauthorized access was minimized. This design not only supported energy efficiency but also aligned with the project's goal of maintaining privacy through decentralized computation.

We conducted detailed profiling of CPU and memory usage during runtime. This allowed us to identify and eliminate performance bottlenecks, ensuring that every computation step used the least amount of power necessary. The software was optimized to handle variable power inputs, dynamically adjusting computational intensity based on available solar energy. This adaptability made the system robust under different environmental conditions, ensuring continuous operation even during cloudy or rainy days.

This section of the project demonstrates how low-power AI processing can be achieved on solar-powered devices through a combination of model optimization, efficient coding, hardware acceleration, and intelligent power management. By using lightweight architectures, quantization techniques, duty cycling, and localized processing with SQLite3 storage, the system maintained a careful balance between computational performance and energy sustainability. This approach proves that even under tight energy constraints, AI-driven applications like facial recognition and intrusion detection can operate effectively and autonomously on solar-powered devices.

3.5 Use Federated Learning to Improve Security Without Exposing Private Data

In this section of the project, we focused on integrating federated learning into the system to improve security, intelligence, and privacy. Federated learning is a distributed machine learning approach where multiple devices or nodes train models locally using their own data and share only the trained model parameters, such as weights and biases, instead of the raw data. This allows each node to learn independently while still contributing to a shared global model without exposing sensitive information. In this setup, each solar-powered camera processed its own facial recognition and intrusion detection data locally while periodically sharing learning updates with a central server for aggregation. This structure ensured that the cameras became smarter over time without sacrificing data privacy or increasing network vulnerability.

Each camera acted as an independent learning node, performing local computations like facial feature extraction, anomaly detection, and unauthorized access recognition. Communication between nodes occurred through data packets, which are small, structured units of data transmitted over a network. To measure network reliability and stability, the system monitored the Packet Delivery Ratio (PDR), calculated using the formula:

$$\text{PDR} = (\text{Number of packets received} / \text{Number of packets sent}) \times 100.$$

A high PDR value close to 100% means that most data packets were successfully delivered, indicating smooth and secure communication. For example, if 98 out of 100 packets were received correctly, the PDR would be 98%, showing strong network performance. However, if only 75 out of 100 packets arrived, the PDR would drop to 75%, signaling possible network interference, congestion, or malicious packet loss. Monitoring the PDR

helped detect when the network became unstable, which could indicate either a technical fault or an ongoing intrusion attempt affecting data transmission.

In intrusion detection systems, threat weighting is used to assess the severity of potential risks or anomalies. Weight values represent how the machine learning model classifies and prioritizes each event. Low-weight events, such as normal user recognition or stable packet flow, are assigned around 0.1 and are treated as safe. Medium weights, around 0.7, may represent new or unrecognized faces, irregular data transfers, or unexpected patterns that require observation but are not immediately dangerous. High weights, such as 0.9 or above, are linked to repeated failed entries, abnormal network packets, or attempts to spoof identities. For example, a camera detecting the same unrecognized person multiple times or receiving irregular packets from an unfamiliar device would assign a higher weight, prompting immediate alert or review by the system.

Integrating PDR analysis with threat weighting provides a layered defense mechanism that enhances accuracy in both physical and network-based intrusion detection. For instance, if the PDR remains stable but facial detections fluctuate at medium threat levels, the anomaly is likely physical, such as new individuals entering the frame. However, if the PDR suddenly drops alongside high-weight alerts, it may signal a network-level attack or data interference. These two data types—packet reliability and behavioral weight—work together to give the system context-aware intelligence, allowing it to differentiate between harmless irregularities and genuine security threats.

By combining packet monitoring, PDR calculation, and threat weighting within a federated learning structure, the system becomes more adaptive and resilient. Each node continuously

learns from local activity while contributing anonymized insights to the shared global model. This decentralized learning process strengthens the entire network's detection capabilities without centralizing data, ensuring that security remains high even under low power or offline conditions. The combination of network reliability tracking and weighted threat analysis creates a balanced, intelligent defense mechanism suitable for long-term use in privacy-sensitive IoT environments.

CHAPTER FOUR

RESULTS & DISCUSSION

4.1 Results

The results in this project will be discussed under two sections: Preliminary Testing & Classroom Deployment. The analysis phase will focus on evaluating the performance of the system during both preliminary testing and classroom implementation. Emphasis will be placed on examining detection accuracy, response time, and system resilience under varying network and environmental conditions. Data collected from these phases will be analyzed to determine how the federated learning architecture influences communication efficiency, power consumption, and data privacy during distributed model updates.

Table 1 - 4-Week Intrusion Detection Summary

Weight	Week1	Week 2	Week 3	Week 4	Average (%)
0.1(normal)	50(60.24)	45(54.9)	15(37.5)	25(49.02)	50.42
0.2(low)	15(18.07)	20(24.4)	10(25)	12(23.5)	22.74
0.3(moderate)	10(12.04)	9(10.98)	5(12.5)	7(13.7)	12.3
0.4(high)	5(6.02)	7(8.54)	7(17.5)	2(3.9)	8.99
0.5(critical)	3(3.61)	1(1.22)	3(7.5)	5(9.8)	5.53
	83	82	40	51	

The four-week intrusion-detection summary table above offers a detailed breakdown of how the smart camera classified events across the five weighted categories. By showing both raw counts and percentages, the table makes it easy to track changes in activity patterns from week to week. The averages in the final column give a clear long-term view of system behaviour, revealing that the 0.1 “normal” category dominated at slightly above

50%. This confirms that most observed activities were routine and did not trigger unnecessary alerts, indicating strong baseline accuracy.

A closer look at the weekly figures shows consistent distribution patterns across the four weeks, with only small fluctuations between categories. Low-level disturbances (0.2) and moderate events (0.3) maintained similar proportions across the dataset, suggesting that the system responded predictably even as classroom conditions varied. This uniformity is important because it shows that the intrusion-detection model remained stable over time rather than being influenced by temporary spikes in activity or inconsistent sensor performance.




Figure 9 - Intrusion Detection Pie Chart

The first pie chart visually reinforces the numerical trends from the table. The dominant 50% slice for the 0.1 category makes the prevalence of normal behaviour immediately clear. Meanwhile, the 0.2 (23%) and 0.3 (12%) slices show that minor disturbances accounted for

a meaningful but controlled portion of the detected events. In contrast, the relatively small slices corresponding to 0.4 (9%) and 0.5 (6%) highlight that high-risk and critical incidents were uncommon. Together, these proportions show that the weighting mechanism successfully mapped real classroom behaviour to appropriate risk categories.

The graphical interpretation also helps confirm that the system was neither overly sensitive nor prone to false alarms. If the algorithm had been too aggressive, the higher-weight categories would have taken up a larger portion of the chart. Instead, the low percentages for 0.4 and 0.5 indicate that severe events were detected only when necessary. This balance between sensitivity and restraint is essential for practical intrusion detection, ensuring that administrators receive alerts only for meaningful deviations instead of being overwhelmed by unnecessary notifications.

0.1:(100%-90%) : Excellent.

0.2: (89% - 80%): Good

0.3: (79% - 50%) : Moderate

0.4: (49% - 30%): Low

0.5: (29% - 0%): Critical




Figure 10 - Average Power Used

The second pie chart introduces a qualitative interpretive scale—Excellent, Good, Moderate, Low, and Critical—that maps directly to the weight categories. This translation

from numbers to qualitative labels helps contextualize the system’s performance and overall operational health. When these labels are viewed alongside the pie chart on the right, it becomes evident that the system operated predominantly in the “Excellent” and “Good” ranges. This means that the camera’s behaviour classification aligned strongly with expected conditions and maintained operational efficiency.

Taken together, the table, charts, and interpretive scale all highlight the strength of the weighting strategy used by the smart camera system. By assigning increasing weights to more severe behaviours, the model could effectively distinguish normal events from genuine threats. The dominance of low-weight categories, combined with the small percentage of high-risk detections, demonstrates that the system remained accurate, energy-efficient, and reliable throughout the monitoring period. Overall, the results confirm that the detection method is robust enough for real-world deployment, offering both practical security value and consistent performance.

4.1.1 Preliminary Testing & Result Analysis

Before the full classroom deployment, a preliminary testing phase was conducted to verify the functionality of the facial recognition and database logging code. This initial test was performed on a standard laptop using its built-in webcam rather than the solar-powered smart camera. The objective was to confirm that the recognition pipeline—from image capture to database storage—worked as intended under controlled conditions. During this phase, several facial samples were collected to test the system’s response time, accuracy, and logging reliability. The SQLite3 database was used to store the embeddings, timestamps, and associated metadata for

each recognition event. Once the preliminary verification confirmed the stability and accuracy of the code, the system was subsequently transferred to the solar-powered smart camera for the main classroom implementation.

Table 2 - Preliminary Results

STUDENT_ID	NAME	TIMESTAMP	INITIAL_SCAN_TIME (s)	AVERAGE_MATCH_TIME (s)
1	HAMZAT	2025-09-25T12:18:00.395441	9.8	3.0
2	MIMI	2025-09-25T12:18:10.282810	10.4	2.9
3	JUNE	2025-09-25T12:18:20.413571	8.7	3.4
4	ADELEKE	2025-09-25T12:18:32.644771	10.9	2.6
5	ADEWALE	2025-09-25T12:18:48.353592	9.3	3.2

The table above represents a sample of the preliminary testing records obtained during webcam-based evaluation. The initial scan time indicates the period required for the system to process and register a new facial embedding into the SQLite3 database, averaging around 10 seconds per individual. Subsequent recognitions of the same faces took approximately 3 seconds on average, demonstrating fast and consistent matching performance. The recognition algorithm maintained an average accuracy of about 90% with no false matches reported throughout the testing. It also performed effectively under various lighting conditions, including bright natural light, dim indoor illumination, and partially shaded environments. Additionally, the model adapted well to minor facial variations such as the use of eyeglasses, hairstyle changes, and slight differences in facial expression between scans.

4.1.2 Classroom Deployment & Result Analysis

The classroom deployment phase marked the transition of the system from a controlled laptop-based environment to a real-world academic setting using the solar-powered smart cameras. Prior to deployment, facial scans were collected from 40 students and 10 lecturers, totaling 50 unique identities. This number was deliberately selected to align with the upper limit programmed into the system's code for the maximum number of stored embeddings. The limit of 50 scans was considered sufficient for the project's objectives, ensuring manageable data processing while maintaining high recognition accuracy. The smart cameras were positioned strategically at classroom entrances, continuously capturing video streams and automatically performing facial scans for all individuals entering the room.

Records were maintained continuously over five school days, with each day dedicated to testing specific aspects of the system and addressing emerging technical challenges. On Day 1, the team observed that recognition latency was slightly higher than expected, primarily due to network synchronization delays between the local database and the solar camera. This issue was resolved by optimizing the caching process and improving the real-time communication buffer. Day 2 presented a lighting-related challenge as glare from nearby windows affected facial clarity during morning sessions; to address this, the camera angles were adjusted, and exposure settings were recalibrated. On Day 3, minor mismatches occurred when individuals wore face masks or hats that partially obscured facial

features; this was mitigated by updating the preprocessing algorithm to better handle partial occlusions and recalibrating the recognition threshold.

By Days 4 and 5, the system had stabilized, and structured data collection was conducted for performance analysis. The final phase focused on recording attendance data, match confidence levels, and system responsiveness under normal classroom conditions. During these two days, the system’s recognition performance, data integrity, and power efficiency were closely monitored. The outcomes of this phase, including quantitative and qualitative results from the classroom deployment, are presented and discussed in the next sections.

4.1.2.1	Day	4
Activity		Log
<i>7:55 AM – System Boot</i>		
<i>Solar-powered smart camera activates automatically.</i>		
<i>Dashboard message:</i>		
<i>“System Online – Monitoring Active – Power Source: Solar.”</i>		
<i>Battery Level: 100%</i>		
<i>Status: Secure</i>		
<i>8:03 AM – Dr. Isi Arrives</i>		
<i>Camera detects motion.</i>		
<i>Recognition process: “Matching with Database...”</i>		

Result: "Match Found – Confidence: 95%."

System logs attendance.

Dashboard light turns green.

Message: "Authorized Entry Recorded – 8:03 AM."

8:42 AM – Intruder #1 Detected

Unrecognized face appears at the entrance.

Camera alert:

"Unknown Face – Possible Intrusion."

Automatic actions:

- *Image saved as "Intruder_0842AM.jpg"*
- *Silent alert sent to admin dashboard*
- *Event logged with timestamp*

The intruder leaves after standing briefly by the door.

9:16 AM – Dr. Mrs. Okosun Arrives

Camera identifies face after two seconds.

*Display: "Authorized **Personnel** – Confidence: 92%."*

Attendance recorded at 9:16 AM.

Dashboard status: Secure.

9:58 AM – Intruder #2 Detected (Repeated Entry)

Same unrecognized person returns wearing a cap.

System compares with stored intrusion images.

Message: “Match with Previous Unknown Entry – Intrusion Flagged.”

Admin dashboard flashes red.

Security notification sent automatically.

10:11 AM – Engr. Sly Arrives

Camera matches face successfully.

“Authorized Personnel – Confidence: 90%.”

Attendance logged.

System note: “Lecturer recognized after intrusion event.”

Dashboard returns to normal.

11:07 AM – Prof. Apeh Arrives (Late)

Camera picks face after he pauses by the entrance.

“Face Detected – Processing...”

*Result: “Authorized **Personnel** – Confidence: 97%.”*

Attendance recorded at 11:07 AM.

Battery Level: 93% (Stable).

11:42 AM – Intruder #3 Detected

Two unrecognized students appear at the door.

Camera logs both faces.

Display:

“Unknown Faces – Dual Intrusion – Logged Automatically.”

System saves two frames, assigns labels Intruder_1142A and Intruder_1142B.

Admin alerted silently.

12:26 PM – Engr. Omisgho Arrives

Camera captures face on entry.

“Match Found – Confidence: 98%.”

Attendance entry saved.

Dashboard shows six total detections (four authorized, two intruders).

1:09 PM – Engr. Solomon Arrives

System detects movement near the board area.

“Face Matched – Confidence: 94%.”

Attendance recorded.

Dashboard message: “All Recognized Entries Authorized.”

1:48 PM – Intruder #4 Detected (Unknown Student)

A face not in the database is detected sitting at the back.

Alert displayed:

“Unrecognized Individual in Frame – Possible Unauthorized Presence.”

Camera locks frame, logs timestamp, saves image.

System remains stable and continues recording the class.

2:00 PM – End of Session Summary

Automatic report generated:

- *Authorized Entries:*
- *Dr. Isi – 8:03 AM*
- *Dr. Mrs. Okosun – 9:16 AM*
- *Engr. Sly – 10:11 AM*
- *Prof. Apeh – 11:07 AM*
- *Engr. Omisgho – 12:26 PM*
- *Engr. Solomon – 1:09 PM*
- *Intrusions Detected: 4*
- *Intruder_0842AM*
- *Intruder_0958AM*
- *Intruder_1142A*
- *Intruder_1142B*
- *System Power: Solar Stable (Battery 81%)*
- *Data Logs: Saved Locally (attendance.db)*
- *Network: Private and Secure*

Dashboard final message:

“Monitoring Complete – Attendance and Intrusion Logs Stored Successfully.”

The daily activity log shows that this particular day was set aside to test the facial recognition system specifically on lecturers, not on students or

for general attendance taking. The goal was to see how well the camera could identify authorized teaching staff under normal classroom conditions. The system operated from 8 AM to 2 PM, powered completely by solar energy, with the battery remaining stable above 80%. All six lecturers were successfully recognized with accuracy levels between 90% and 98%, and their detection times were automatically recorded in the database. Each recognition took only a few seconds, showing that the camera and the facial recognition software worked smoothly together during real-time operation.

During the same period, the camera detected four intruders—individuals whose faces were not registered in the database. One of these intruders attempted to re-enter later, and the system correctly matched the face with a previous record, proving that it could remember and cross-check unrecognized faces. Each intrusion event triggered a silent alert, saved a photo, and logged the exact time it occurred. This meant that while the focus of the test was lecturer recognition, the intrusion detection feature was also being verified in parallel to ensure that unauthorized entries could be identified without disrupting normal classroom activities.

The facial recognition system performed well across different lighting conditions, from bright morning sunlight near the door to softer indoor light in the afternoon. It handled multiple face detections efficiently without lag or errors. The solar-powered operation stayed consistent, and the system continued recording without needing an external power source. This day's test confirmed that the facial recognition model worked effectively for identifying lecturers and that the system could maintain both accuracy and stability during continuous operation.

Table 3 - Day 5: Activity Log

TIME	EVENT DESCRIPTION	SYSTEM OBSERVATION	SYSTEM ACTION	DETECTION LEVEL	REMARKS
8:00 AM	System starts	Solar camera powered on	Initializes monitoring on	0.1	No intrusion
8:15 AM	Students begin to arrive	Faces matched with database	Attendance recorded for 20 students	0.1	Normal operation
8:25 AM	Dr. Isi arrives	Face recognized as authorized	Attendance recorded	0.1	Stable operation
9:10 AM	Unknown person walked in	Face not recognized	Silent alert sent to the admin	0.3	Possible intrusion; person left immediately
9:30 AM	More students joined the class	Recognized faces	Attendance updated	0.1	Stable operation
10:00 AM	Dr. Isi leaves	Recognizes faces	Updated attendance	0.1	Normal operation
10:15 AM	Dr. Mrs. Okosun enters	Authorized face recognized	Attendance updated	0.1	Stable operation
10:45 AM	Another intruder enters	Unrecognized face	Silent alert sent to the admin	0.3	Possible intrusion; left immediately
11:00 AM	Class in session	Recognized faces	Attendance updated	0.1	System stable
11:10 AM	Dr. Mrs. Okosun leaves	Recognized faces	Updated attendance	0.1	Stable operation
11:20 AM	Engr. Sly arrives	Authorized face recognized	Attendance updated	0.1	Normal operation
12:00 PM	Same two intruders returned	Unrecognized faces	Silent alert sent to the admin	0.4	Intrusion flagged
12:30 PM	Engr. Sly leaves	Recognizes faces	Updated attendance	0.1	Stable operation
12:35 PM	Students on break	Normal movement detected	Updated attendance	0.1	Normal operation
1:05 PM	Prof. Apeh arrives	Face successfully recognized	Updated attendance	0.1	Stable operation
2:15 PM	Unknown student spotted	Unrecognized face seen at the back	Silent alert sent to the admin	0.4	Intrusion recorded
3:00 PM	Prof. Apeh leaves	Face recognized	Updated attendance	0.1	Stable operation
3:15 PM	Engr. Omisgho arrives	Authorized face recognized	Attendance updated	0.1	Normal operation
3:45 PM	Engr. Omisgho leaves	All recognized and verified	Updated attendance	0.1	Stable operation
3:55 PM	Students leave the class	25 students were recorded	Attendance finalized	0.1	Normal operation
4:00 PM	System shuts down	Solar camera powered off	Monitoring ended	—	No operation

The day's activity log shows a full operational cycle running on the solar-powered smart camera from 8:00 AM to 4:00 PM in the Computer Engineering Department. Throughout the day, the system performed dual functions — facial recognition for attendance and real-time intrusion detection — while maintaining stable solar power performance. The log recorded several types of events, including authorized staff recognition, student attendance, and unrecognized face detections that were automatically flagged as potential security threats. The consistent timestamps and recorded confidence levels indicate that the system was running continuously, collecting and analyzing data without interruptions or power loss.

In this system, threat levels are represented as weights in the machine learning model, allowing the camera to classify events based on their perceived risk. Normal attendance and recognized users were assigned a weight of 0.1, indicating a low-risk or “safe” classification. When the system detected a new or unrecognized face, it generated a higher weight of up to 0.7, classifying the event as a moderate-level threat that required administrative attention. In cases of repeated failed entries, duplicate intrusion attempts, or unrecognized network packets, the weight increased to 0.9, marking the event as a high-level threat. This weighting system helped the federated learning model to distinguish between normal, suspicious, and malicious behavior patterns, allowing the system to adapt and respond accordingly over time.

The log shows that the system effectively identified multiple threat levels throughout the day. Early events such as students arriving and lecturers entering were processed under the 0.1 threshold, which signified normal activity. However, intrusion events at 9:10 AM, 10:45 AM, 12:00 PM, and 2:15 PM showed higher detection levels ranging from 0.3 to 0.4, corresponding to unrecognized faces. These readings reflected how the model scaled its detection weights in real time, learning to adjust its classification confidence based on environmental and contextual factors. When repeated intrusions occurred (as seen at 12:00 PM), the detection level rose, signaling that the system had detected a pattern consistent with repeated failed access attempts.

In addition to facial recognition, the system also monitored network packets, which are small units of data transmitted between connected IoT devices. Each packet carries information necessary for device communication, and by monitoring them, the system could detect irregular data flows or unauthorized access attempts. The Packet Delivery Ratio (PDR) — the percentage of successfully delivered packets compared to those sent — served as a reliability metric. A stable PDR near 100% indicated that communication between the smart camera, database, and dashboard was functioning smoothly. Any sudden drop in PDR would suggest network interference or a possible intrusion attempt at the data level, prompting the system to assign a higher threat weight for further inspection.

By the end of the day, the system had accurately logged attendance for 25 students and multiple lecturers, while also detecting four intrusion attempts without any false alarms. Each alert triggered a silent notification and automatic image logging, ensuring traceability of all events. The continuous power supply from the solar source and the stable communication performance confirmed that the system's design could sustain extended monitoring without external support. Overall, the data from this session highlights the system's ability to intelligently categorize threats, maintain secure communication, and ensure consistent operation in a live classroom environment using machine learning-based adaptive weighting and packet monitoring.

4.2 Discussion

The results indicate that the system achieved a strong balance between accuracy and efficiency, with match rates consistently above the 60% success threshold and the majority exceeding 80%. This performance level demonstrates that the AI model is robust enough for classroom-level deployment, especially considering the system's constraints on processing power and energy usage. The consistency of matches over several days also confirms that the facial recognition algorithm can handle real-time attendance tracking without network dependency.

However, minor variations in match percentages show that facial recognition accuracy can fluctuate due to environmental and human factors. These include changes in classroom lighting, partial obstructions like masks, or differences in facial appearance (haircuts, glasses, or facial hair). Despite these factors, the system's ability to maintain accurate

recognition and avoid false matches highlights the strength of the underlying AI model and preprocessing pipeline.

The cap of 50 facial scans proved adequate for testing, but future versions could benefit from expanding this limit to accommodate larger class sizes or multiple departments. Additionally, the current system's use of CSV-based data export provides a straightforward method for lecturers and administrators to receive attendance summaries via email. The flexibility of the SQLite3 database ensures scalability, allowing more complex analytics or reporting features to be implemented later.

By ensuring that no personal identifiers other than facial embeddings are stored, the system maintains compliance with privacy standards while still enabling lecturers to track attendance frequency through unique internal identifiers. These results and observations suggest that the system is both technically sound and adaptable for broader institutional use in the future.

CHAPTER FIVE

CONCLUSION & RECOMMENDATION

5.1 Conclusion

In conclusion, this project successfully achieved its aim of developing a smart, lightweight, federated deep learning system that utilizes solar-powered cameras to detect and prevent cyber threats while automating attendance and security monitoring for CPE 300L and 500L students and lecturers. Through the integration of AI-based facial recognition, federated intrusion detection, and local computing, the system demonstrated that it is possible to build a sustainable, privacy-preserving, and low-power monitoring solution. The design effectively combined solar energy efficiency with real-time facial recognition, ensuring that attendance could be tracked without external connectivity or manual intervention. This approach not only improved operational autonomy but also showcased how renewable energy and artificial intelligence can work together to support academic and security applications.

The system's face recognition module—implemented using TensorFlow Lite and Python—automated attendance recording with a high accuracy rate of over 80%, even under variable classroom conditions. The inclusion of a structured SQLite3 database ensured efficient and secure local data handling, while the Flask web framework provided a flexible, user-friendly interface for real-time monitoring. The deployment on solar-powered cameras validated the system's energy-efficient design, achieving stable performance with minimal computational overhead.

The project demonstrated that a federated deep learning–based intrusion detection and attendance system can operate effectively on solar-powered IoT devices, providing secure, autonomous, and privacy-conscious functionality. The combination of cybersecurity measures, including DDoS and unauthorized access protection, with AI-driven facial recognition, resulted in a robust and adaptable solution for classroom management and IoT security. The system’s performance in terms of accuracy, energy efficiency, and resilience highlights its potential for broader adoption in educational and low-resource environments.

5.2 Recommendation

While the system met its core objectives of developing a smart, federated deep learning–based intrusion detection and attendance monitoring system, there are several ways it can be enhanced to improve scalability, efficiency, and security. The current implementation demonstrates strong potential for real-world use, but future iterations can focus on expanding capacity, improving automation, and enhancing both physical and digital security features. The following recommendations outline key areas for future development and optimization:

1. Increase the Upper Limit of Face Scans

The current system has a built-in limit of 50 facial scans, which is suitable for small-scale testing but insufficient for large classrooms or multi-departmental use. Future versions should raise this limit to at least 2000 face scans to accommodate more students, lecturers, and visitors. Expanding the face embedding capacity will improve scalability and make the system viable for deployment across the entire faculty or institution.

2. **Deploy Additional Cameras Across More Classrooms**

To provide comprehensive coverage and enable simultaneous monitoring of multiple classes, more AI-integrated cameras should be installed in additional classrooms and lecture halls. This would also allow for redundancy—ensuring that even if one camera fails, attendance tracking and access monitoring continue uninterrupted.

3. **Increase the Number of Administrators for Enrollment**

The process of capturing live face scans can be time-consuming when handled by only one or two administrators. Increasing the number of trained admins will speed up the registration process and allow simultaneous scanning across multiple classrooms or departments. This will be especially beneficial at the start of a new session when many students need to be registered at once.

4. **Implement Mechanical Access Control**

Currently, the system only logs unauthorized access attempts without preventing physical entry. Future improvements should integrate mechanical door locks or barriers that are electronically controlled and respond to recognition results. This would make the system more secure by actively preventing unauthorized individuals from entering restricted classrooms.

5. **Add an Administrator Override Mechanism**

To accommodate situations such as exam periods, when individuals outside the registered departments need access, there should be an override feature allowing authorized administrators to temporarily disable or bypass mechanical locks. This

ensures flexibility in access management while maintaining system integrity during regular operation.

6. **Enable Secure Storage of ID Information (Optional, Encrypted Form)**

While the current design prioritizes privacy by not storing names or IDs, future versions could include an encrypted and securely stored ID mapping system. This would allow automatic generation of attendance reports that include student names and matriculation numbers without compromising data privacy. Proper encryption and access control policies should be applied to ensure this remains secure.

7. **Introduce Motion-Activated Camera Operation**

To conserve power in solar-powered setups, future iterations should implement motion sensors or computer vision–based activity detection. The cameras would remain in low-power standby mode and only activate AI processing when motion is detected near the door, significantly extending operational time and improving energy efficiency.

8. **Expand the Federated Learning Network**

Future deployments can connect multiple camera units into a federated learning framework, where each camera trains locally on its own data and periodically shares model updates (not raw data) with a central server. This would improve recognition accuracy while maintaining privacy and minimizing bandwidth usage.

9. **Enhance Intrusion Detection Intelligence**

The intrusion detection system can be expanded with real-time anomaly detection that analyzes behavioral patterns, such as repeated failed access attempts or

suspicious network traffic. This could further strengthen protection against unauthorized access and potential cyber threats, even in offline environments.

10. Improve Data Visualization and Reporting Tools

While the current CSV and Excel export system works effectively, future versions could include a built-in reporting dashboard for attendance summaries, performance analytics, and security alerts. This would reduce administrative workload and allow real-time insights without needing external tools.

11. Integrate Biometric Fusion for Multi-Factor Authentication

Future iterations could combine facial recognition with other biometrics, such as fingerprint or voice recognition, for critical security zones. This would improve accuracy and minimize spoofing or impersonation risks while keeping the system flexible for classroom environments.

12. Conduct Long-Term Performance and Durability Testing

Finally, extended deployment over several academic semesters should be conducted to assess system durability, long-term accuracy, and solar power reliability.

Continuous monitoring and iterative improvement based on real-world feedback will help refine both the hardware and software components for future scalability.

REFERENCES

1. Kairouz, P., McMahan, H.B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A.N., Bonawitz, K., Charles, Z., Cormode, G., Cummings, R. and D'Oliveira, R.G., 2021. Advances and open problems in federated learning. *Foundations and trends® in machine learning*, 14(1–2), pp.1-210.
2. Belarbi, O., Spyridopoulos, T., Anthi, E., Mavromatis, I., Carnelli, P. and Khan, A., 2023, December. Federated deep learning for intrusion detection in IoT networks. In *GLOBECOM 2023-2023 IEEE Global Communications Conference* (pp. 237-242). IEEE.
3. Sun, S., Sharma, P., Nwodo, K., Stavrou, A. and Wang, H., 2024, October. Fedmade: Robust federated learning for intrusion detection in Iot networks using a dynamic aggregation method. In *International Conference on Information Security* (pp. 286-306). Cham: Springer Nature Switzerland.
4. Rahmati, M. and Pagano, A., 2025, July. Federated Learning-Driven Cybersecurity Framework for IoT Networks with Privacy Preserving and Real-Time Threat Detection Capabilities. In *Informatics* (Vol. 12, No. 3, p. 62). MDPI.
5. Abas, K., Obraczka, K. and Miller, L., 2018. Solar-powered, wireless smart camera network: An IoT solution for outdoor video monitoring. *Computer Communications*, 118, pp.217-233.
6. Miller, L., Abas, K. and Obraczka, K., 2015, August. Scmesh: Solar-powered wireless smart camera mesh network. In *2015 24th International Conference on Computer Communication and Networks (ICCCN)* (pp. 1-8). IEEE.

7. Ferrag, M.A., Friha, O., Maglaras, L., Janicke, H. and Shu, L., 2021. Federated deep learning for cyber security in the internet of things: Concepts, applications, and experimental analysis. *IEEE Access*, 9, pp.138509-138542.
8. Gugueoth, V., Safavat, S. and Shetty, S., 2023. Security of Internet of Things (IoT) using federated learning and deep learning—Recent advancements, issues and prospects. *ICT express*, 9(5), pp.941-960.
9. Anitha, A.A. and Arockiam, L., 2022. A review on intrusion detection systems to secure IoT networks. *International Journal of Computer Networks and Applications*, 9(1), pp.38-50.
10. Sicato, S., Costa, J., Singh, S.K., Rathore, S. and Park, J.H., 2020. A comprehensive analyses of intrusion detection system for IoT environment. *Journal of Information Processing Systems*, 16(4).
11. Bhavsar, M., Roy, K., Kelly, J. and Olusola, O., 2023. Anomaly-based intrusion detection system for IoT application. *Discover Internet of things*, 3(1), p.5.
12. Nadhan, A.S., Tukkoji, C., Shyamala, B., Dayanand Lal, N., Sanjeev Kumar, A.N., Mohan Gowda, V., Adhoni, Z.A. and Endaweke, M., 2022. Smart attendance monitoring technology for industry 4.0. *Journal of Nanomaterials*, 2022(1), p.4899768.
13. Bavaskar, V., 2024. *Face Recognition Attendance System* (Doctoral dissertation, Doctoral dissertation, Sant Gadge Baba Amravati University, Amravati).

14. Shashikala, H.K., Shakya, S.N., Panjiyar, P., Upreti, A.S. and Dadapeer, S., 2022. Attendance monitoring system using face recognition. *International Journal of Information Technology, Research and Applications*, 1(3), pp.15-22.
15. Kumari, P. and Jain, A.K., 2023. A comprehensive study of DDoS attacks over IoT network and their countermeasures. *Computers & Security*, 127, p.103096.
16. Khanday, S.A., Fatima, H. and Rakesh, N., 2023. Implementation of intrusion detection model for DDoS attacks in Lightweight IoT Networks. *Expert Systems with Applications*, 215, p.119330.
17. nightfury217836 (2025) *Digital-Facial-Recognition-Attendance-System*. GitHub repository. Available at: <https://github.com/nightfury217836/Digital-Facial-Recognition-Attendance-System> (Accessed: 24 October 2025).