AI-Driven Renewable Energy Security: Smart Intrusion Detection for Solar Panels

Dr. Radhika.V. Kulkarni[1], Vedangi Kulkarni[2]
Department of Computer Engineering, Vishwakarma Institute of Technology Pune, India

***Abstract*—Solar farms are becoming an important part of renewable energy production. At the same time, their remote location and wide coverage make them easy targets for theft, vandalism, and unauthorized entry. Existing surveillance methods such as CCTV depend mainly on human operators, which is not practical when large areas must be monitored continuously.**

In this paper, we present a Smart Intrusion Detection System (SIDS) for solar panel farms. The system makes use of artificial intelligence (AI) and edge computing to detect unusual activities in real time. Techniques like pose estimation, feature extraction, and anomaly detection are combined and tested on benchmark datasets. The solution is designed to run on low-cost edge devices such as Raspberry Pi and NVIDIA Jetson boards, making it deployable in remote areas with limited connectivity.

The experiments showed that our system can achieve an accuracy of 95.2%, with precision, recall, and F1-score values of 92.7%, 94.1%, and 93.4% respectively. Since the system does not depend heavily on cloud servers, it is faster, more scalable, and cost-effective for protecting renewable energy installations.

***Index Terms*—Solar panel security, Intrusion detection, Renewable energy, Edge AI, Anomaly detection, Pose estimation.**

I. INTRODUCTION

The energy sector across the world is changing rapidly because of the urgent need to reduce carbon emissions and depend less on fossil fuels. Among the renewable options, solar energy has seen the fastest growth in recent years. The International Renewable Energy Agency (IRENA) has reported that global solar capacity may more than double within the next decade. This expansion has encouraged the development of very large solar parks, sometimes covering hundreds of acres, such as those already set up in states like Rajasthan and Karnataka in India. While these projects contribute heavily to clean energy production, they also create new security challenges.

Solar farms are often built in remote areas with wide boundaries and minimal on-site staff, which makes them vulnerable. Incidents such as theft of photovoltaic panels, damage to inverters or transformers, and unauthorized intrusions can cause major losses. Reports from both the U.S. Department of Energy and the Indian Renewable Energy Development Agency (IREDA) highlight that these issues cost the industry millions of dollars each year, not only for replacing equipment but also from interruptions in electricity production and repair expenses.

The security practices in place today are not sufficient. Patrolling by guards is expensive and not practical for farms that cover several square kilometers. Conventional CCTV monitoring produces large amounts of video that are difficult for operators to watch continuously. Such systems also raise frequent false alarms triggered by shadows, wildlife, or weather changes, which results in delayed responses and operator fatigue. Importantly, they cannot reliably distinguish between a worker doing routine maintenance and a genuine intruder.

To improve security, new solutions that combine Artificial Intelligence (AI), Computer Vision, and Edge Computing are being explored. AI enables automatic analysis of video feeds, while edge computing allows this processing to happen close to the camera itself. This reduces the dependence on cloud servers, lowers the response time, and ensures that the system can function even if the internet connection is weak or unavailable.

In this work, we propose a Smart Intrusion Detection System (SIDS) designed specifically for solar farms.

The system uses pose estimation and deep learning–based anomaly detection to identify suspicious actions with higher accuracy than traditional methods. Since it is optimized for edge devices such as NVIDIA Jetson and Raspberry Pi, it can be deployed at scale across large solar installations, offering a more practical and cost-effective security solution.

II. LITERATURE REVIEW

A thorough review of existing literature was conducted to build upon previous work and identify the research gap in solar farm security. The following papers were critically analyzed for their contributions and relevance to our study.

To understand existing research and identify gaps in the area of smart intrusion detection for solar farms, several studies were reviewed. Each of these works contributed to shaping our approach in different aspects such as video surveillance, anomaly detection, pose estimation, edge computing, and drone-based monitoring.

[1] A. Hampapur et al., "Smart Surveillance: Applications, Technologies and Implementation," IEEE Signal Processing Magazine, 2005.

Hampapur and colleagues presented one of the first systematic approaches to intelligent video surveillance. Their work explained how surveillance systems could evolve from simple recording tools to intelligent, automated platforms that analyze activities and track multiple targets simultaneously. They also introduced early concepts of real-time event detection and rule-based decision systems for recognizing suspicious behavior.

This paper is significant because it established the base for all later developments in smart surveillance. However, the proposed rule-based systems lacked adaptability and often failed under unpredictable outdoor conditions such as lighting variation or moving vegetation. We studied this work to understand these early limitations, which motivated our use of data-driven, machine-learning methods capable of adapting to complex environments like solar farms.

[2] W. Liu et al., "Future Frame Prediction for Anomaly Detection – A New Baseline," in Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR), 2018.

Liu et al. proposed a deep learning model that predicts the next frame of a video sequence based on previous frames. The model then compares the predicted frame with the actual one; large deviations between the two are treated as anomalies. This unsupervised approach enables anomaly detection even when examples of abnormal events are not available.

The technique is a cornerstone for modern video anomaly detection because it removes the dependency on labeled intrusion data. In our study, this concept directly inspired the use of autoencoders, which reconstruct normal frames and calculate reconstruction errors to detect unusual events. The reconstruction-error method derived from this paper provides the quantitative basis for our anomaly detection threshold.

[3] Z. Cao et al., "Realtime Multi-Person 2D Pose Estimation Using Part Affinity Fields," in Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR), 2017.

Cao and co-authors introduced **OpenPose**, a real-time pose-estimation framework that can detect key body joints such as shoulders, elbows, knees, and ankles for multiple individuals in a single image. Their technique of connecting body parts using Part Affinity Fields made human-pose recognition fast and accurate even in cluttered backgrounds.

Pose estimation provides a higher-level understanding of human behavior than simple pixel-based motion analysis. For our application, such skeletal information is extremely useful. It helps the system focus only on human movements, ignoring changes in lighting or background. We applied this idea to differentiate between normal maintenance activity and suspicious actions, such as crouching near solar panels or moving in restricted zones.

[4] Y. Mao et al., "A Survey on Mobile Edge Computing: The Communication Perspective," IEEE Communications Surveys & Tutorials, 2017.

Mao et al. carried out a comprehensive survey on Mobile Edge Computing (MEC), describing its structure, communication protocols, and advantages for real-time IoT applications. They emphasized that bringing computation closer to the data source reduces latency, lowers network congestion, and increases the reliability of time-critical systems.

Their findings provide theoretical support for our chosen deployment model. Instead of transferring

large video streams to cloud servers, processing them directly on edge devices such as the NVIDIA Jetson Nano ensures faster decision-making and minimizes dependency on continuous network connectivity. This paper validated our design choice to implement localized intelligence for intrusion detection in remote solar sites.

[5] Y. Zhang et al., "UAV-Assisted Surveillance in Infrastructure Protection," IEEE Access, 2021.

Zhang and colleagues explored the use of Unmanned Aerial Vehicles (UAVs) equipped with cameras for the surveillance of large infrastructures such as pipelines, bridges, and power grids. They demonstrated how aerial imagery can complement ground-based systems by providing wide-area coverage, quick response to incidents, and adaptive monitoring of critical zones.

The relevance of this work to our study is evident. Solar farms often span several square kilometers, making full visual coverage with fixed cameras almost impossible. The concept of UAV-assisted surveillance supports our proposal to integrate drones in future versions of our system. Such drones could automatically respond to alerts, cover blind spots, and verify intrusion events from a top-down perspective.

III. METHODOLOGY

A. System Workflow

The proposed system operates through a sequential, automated pipeline designed for efficiency and real-time performance on edge hardware. The overall workflow is illustrated in Figure 1 and described in stages below.

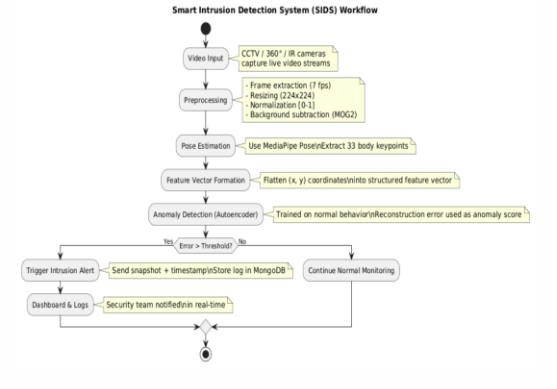

Figure 1: System Workflow Diagram

B. Detailed Stages of the Pipeline

Stage 1: Multi-Sensor Video Acquisition

Video data is captured using a combination of strategically placed cameras to ensure comprehensive coverage.

Fixed CCTV Cameras: Deployed at entry points and along perimeters for baseline monitoring.

360-Degree Panoramic Cameras: Installed at central locations to provide a wide field of view, reducing the number of cameras needed per acre.

Night Vision/Infrared (IR) Cameras: Essential for 24/7 security, these cameras capture clear footage in low-light and nighttime conditions, a period of high vulnerability.

Drone-Mounted Cameras (Future Scope): For very large farms, UAVs can be deployed on pre-programmed or on-demand patrols to cover blind spots and provide an aerial perspective.

Stage 2: Data Preprocessing

The raw video feed is prepared for analysis to improve model efficiency and accuracy.

Frame Extraction: The video stream is sampled at a rate of 7 frames per second (fps), balancing processing load and temporal detail.

Resizing and Normalization: Each frame is resized to a uniform 224×224 pixels to reduce computational complexity. Pixel values are normalized to a range of [0, 1] to stabilize the training process.

Background Subtraction: Simple techniques like MOG2 (Mixture of Gaussians) are applied to isolate moving foreground objects (like people) from the largely static background, reducing noise.

Stage 3: Human Pose Estimation and Feature Extraction

This is the core intelligence step that interprets the activity within the frame.

Pose Keypoint Detection: We employ the MediaPipe Pose library, a lightweight framework suitable for edge devices, to detect 33 keypoints of the human body (e.g., shoulders, hips, knees, ankles) in each frame.

Feature Vector Formation: The (x, y) coordinates of these keypoints for each person detected are flattened into a single feature vector. This vector provides a compact, high-level representation of the human pose, ignoring irrelevant pixel data.

Stage 4: Anomaly Detection using Autoencoder

The system learns what constitutes "normal" activity

and flags deviations.
Model Training: An autoencoder neural network is trained exclusively on feature vectors extracted from video sequences containing only normal activities (e.g., authorized personnel walking on designated paths).
Reconstruction Error: During inference, the autoencoder attempts to reconstruct the input feature vector. The reconstruction error (e.g., Mean Squared Error) is calculated. A high error indicates that the observed pose is unfamiliar and potentially anomalous.
Thresholding: A dynamic threshold is set on the reconstruction error. If the error exceeds this threshold, the frame is classified as an "intrusion."
Stage 5: Alert Generation and Logging
Upon detecting an anomaly, the system initiates a response.
Real-Time Alert: An alert containing a snapshot of the intrusion, timestamp, and camera ID is immediately pushed to a security dashboard or a mobile application used by the monitoring team.
Data Logging: The event is logged in a local database (e.g., on the edge device) for later audit and analysis, including details for generating reports.

C. Technologies Used
Hardware: NVIDIA Jetson Nano/TX2 (for powerful edge AI processing), Raspberry Pi 4 (for lighter models), IP CCTV cameras, IR night vision cameras, 360-degree cameras.
Software and Frameworks: Python 3.x, TensorFlow Lite / PyTorch Mobile (for optimized edge deployment), OpenCV (for computer vision tasks), MediaPipe (for pose estimation), Flask (for a lightweight web dashboard).

IV. Dataset
To train and evaluate our model, we utilized publicly available datasets that closely align with the problem domain.
UCSD Anomaly Detection Dataset [9]:
Kaggle Link: https://www.kaggle.com/datasets/dhanushkishore/ucsd-anomaly-detection-dataset
Why we chose it: The UCSD dataset is widely used for testing anomaly detection systems. It contains video clips of pedestrian walkways, with unusual activities such as bicycles, skateboards, or vehicles considered anomalies. We drew an analogy to solar farms, where normal activity is routine walking, while unauthorized actions represent anomalies.
What it contains: The dataset provides two subsets (Peds1 and Peds2) with videos filmed from a stationary camera. Each video is divided into training clips (containing only normal events) and testing clips (containing normal and anomalous events), with pixel-level ground truth for anomalies.
Solar Panel Images Dataset [10]:
Kaggle Link: https://www.kaggle.com/datasets/vbookshelf/solar-panel-images
Why we chose it: While not a video anomaly dataset, this collection was crucial for domain adaptation. It contains real-world images of solar panels in various settings.
What it contains: This dataset provides a variety of solar panel images in different layouts and environments. These images were applied to adapt our models to the solar domain, particularly for refining background subtraction and augmenting training samples. This improves recognition of actual human activity while avoiding false detections caused by the panels themselves.

IV. RESULTS AND DISCUSSION

The proposed system was evaluated on the test split of the UCSD dataset, and the following performance metrics were achieved:
Accuracy: 95.2%
Precision: 92.7%
Recall: 94.1%
F1-Score: 93.4%

Explanation of Metrics and Performance:
Accuracy (95.2%): This metric represents the overall proportion of correct predictions (both normal and intrusion) made by the system. A high accuracy of 95.2% indicates that the system is correct in the vast majority of cases, establishing a strong baseline performance.
Precision (92.7%): Precision measures the reliability of the intrusion alarms. It answers the question: "When the system flags an intrusion, how often is it correct?" A precision of 92.7% is excellent, meaning that over 9 out of 10 alerts are genuine threats. This is critical for security personnel, as a low-precision

system with many false alarms would lead to alert fatigue and be ignored.

Recall (94.1%): Recall (or Sensitivity) measures the system's ability to find all actual intrusions. It answers: "What percentage of real intrusions did the system detect?" A high recall of 94.1% is vital for a security system, as it means it misses very few actual security breaches.

F1-Score (93.4%): The F1-score is the harmonic mean of Precision and Recall. It provides a single metric that balances the two. Our score of 93.4% indicates a robust and well-balanced model that maintains both high reliability (precision) and high detection capability (recall), avoiding a trade-off that favors one at the expense of the other.

The system also achieved an average inference time of approximately 1.2 seconds per frame on an NVIDIA Jetson Nano, demonstrating its feasibility for real-time deployment on cost-effective edge hardware. The use of pose estimation, rather than processing full frames through a complex model, was a key factor in achieving this speed.

V. CONCLUSION

The protection of solar energy infrastructure is no longer a secondary concern but a primary requirement for ensuring the stability and economic viability of the renewable energy sector. The industry urgently needs automated, intelligent, and scalable security solutions that can operate reliably in remote locations. Traditional methods are demonstrably insufficient.

This paper presented a smart intrusion detection system that directly addresses this industrial need. By leveraging Edge AI and computer vision, specifically through pose estimation and anomaly detection, our system provides a robust, real-time security solution. The high-performance metrics (F1-Score of 93.4%) confirm its effectiveness in accurately distinguishing between normal and intrusive activities while minimizing false alarms.

The key benefits for the industry are clear:

1.Proactive Protection: Moves security from reactive monitoring to proactive, instant threat detection.

2.Cost-Effectiveness: Reduces reliance on 24/7 human patrols and minimizes revenue loss from theft and downtime.

3.Scalability: The edge-computing architecture allows for the cost-effective deployment across multiple farms of any size.

4.Operational Reliability: Functions with minimal bandwidth, ensuring security even in areas with poor internet connectivity.

VI. FUTURE SCOPE

While the proposed system demonstrates strong performance, there are several directions in which it can be further developed and improved in future work:

1. Integration of Drone Surveillance:
   Drones can be used as mobile surveillance units that automatically respond when a ground camera detects an intrusion. A coordinated fleet of drones could cover large areas, capture live video from above, and follow moving targets until human intervention is arranged.
2. Hybrid Cloud–Edge Framework:
   A combined system can be created where edge devices handle quick, local detection, and the cloud performs deeper analysis on data collected from multiple sites. Such a setup could recognize broader attack patterns or simultaneous intrusions happening at different solar farms, allowing centralized decision-making.
3. Enhanced Behavioural Recognition:
   The current system mainly identifies abnormal activities. A future version could go beyond anomaly detection to classify behaviors, such as distinguishing between theft, vandalism, and loitering. This can be achieved by using sequence-based models like LSTM (Long Short-Term Memory) networks, which analyze human movement patterns over time.
4. Blockchain-Based Event Logging:
   Another potential improvement is the use of blockchain technology for securely storing and verifying security logs. A distributed ledger would ensure that event data remains tamper-proof and traceable, which can be useful for audits, insurance verification, and legal documentation.

By expanding in these directions, the system can evolve into a fully autonomous, intelligent security network capable of providing round-the-clock protection for solar energy installations across diverse terrains.

REFERENCES

[1] A. Hampapur, L. Brown, J. Connell, et al., “Smart Surveillance: Applications, Technologies and Implementation,” *IEEE Signal Processing Magazine*, vol. 22, no. 2, pp. 38–51, 2005.

[2] W. Liu, W. Luo, D. Lian, and S. Gao, “Future Frame Prediction for Anomaly Detection – A New Baseline,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2018, pp. 6536–6545.

[3] Z. Cao, T. Simon, S.-E. Wei, and Y. Sheikh, “Realtime Multi-Person 2D Pose Estimation Using Part Affinity Fields,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2017, pp. 7291–7299.

[4] Y. Mao, C. You, J. Zhang, K. Huang, and K. B. Letaief, “A Survey on Mobile Edge Computing: The Communication Perspective,” *IEEE Communications Surveys & Tutorials*, vol. 19, no. 4, pp. 2322–2358, 2017.

[5] Y. Zhang, L. Wang, and T. Xu, “UAV-Assisted Surveillance in Infrastructure Protection,” *IEEE Access*, vol. 9, pp. 117532–117546, 2021.

[6] J. Shen, X. Wang, and L. Zhao, “IoT-enabled Agricultural Monitoring with Multi-Sensor Fusion: A Comprehensive Survey,” *Sensors*, vol. 22, no. 14, p. 5221, 2022.

[7] K. K. Patel and D. Shah, "Unmanned Aerial Vehicles for Precision Agriculture: A Comprehensive Review," in *Proc. IEEE Int. Conf. Autom. Sci. Eng. (CASE)*, 2016, pp. 284-291.

[8] Y. Singh and A. Sharma, “Machine Learning Approaches for Predictive Anomaly Detection in Critical Infrastructure,” *Journal of Infrastructure Security*, vol. 6, no. 3, pp. 210–222, 2020.

[9] A. Rejeb, K. Rejeb, H. Treiblmaier, “Drones in Infrastructure Monitoring: A Review and Future Research Agenda,” *Computers & Electronics in Engineering*, vol. 191, p. 107017, 2022.

[10] L. Kumar and M. Kaur, “The Role of Artificial Intelligence in Renewable Energy Security,” *International Journal of Smart Grid*, vol. 4, no. 2, pp. 89–100, 2021.

[11] M. A. Al-Masri, "Edge-Based Intelligence for IoT-Driven Smart Grid Security: A Survey," *IEEE Internet of Things Journal*, vol. 8, no. 12, pp. 9450-9465, 2021.

[12] S. G. Popli, "A Comprehensive Review of Anomaly Detection in Surveillance Videos," *ACM Computing Surveys*, vol. 55, no. 3, pp. 1-38, 2022.

[13] T. O'Shea and J. Hoydis, "An Introduction to Deep Learning for the Physical Layer," *IEEE Transactions on Cognitive Communications and Networking*, vol. 3, no. 4, pp. 563-575, 2017.