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

Dr. Radhika.V. Kulkarni[1], Vedangi Kulkarni[2], Ketaki Dharmraj Raut[3], Harshal Prashant Runwal[4]
[1,2,3,4]*Department of Computer Engineering, Vishwakarma Institute of Technology, Pune, India*

***Abstract*—Solar farms are being deployed at large scales to meet renewable energy targets, especially in geographically remote locations. However, the physical isolation of these sites increases the risk of unauthorized entry, damage, and theft of installed components. Traditional CCTV-based monitoring systems rely extensively on manual observation, which becomes impractical when coverage spans several acres and continuous monitoring is needed.**

In this study, a Smart Intrusion Detection System (SIDS) is proposed, specifically designed for solar farm environments. The system integrates lightweight AI-based video analysis with edge-level deployment to identify suspicious activities in real time. Key approaches such as human-pose estimation, feature extraction and anomaly scoring using learned normal patterns were implemented and evaluated using publicly available datasets.

The developed approach demonstrated promising performance, achieving 95.2% accuracy and strong precision–recall balance. As the computation happens on low-power embedded devices, the system remains operational even with limited network access, making it a viable security solution for remote and large-scale solar installations.

***Index Terms*—Solar farms; Edge processing; Intrusion detection; Pose-based analysis; Renewable energy security.**

I. INTRODUCTION

The global energy landscape is undergoing continuous transformation as governments and industries work toward reducing environmental impact and improving energy security. Solar energy remains central in this shift due to its scalability, declining installation cost, and the ability to deploy plants in vast open lands. International renewable energy agencies have reported steady growth in installed solar capacity, and many countries continue to expand large solar farm projects at both government and private levels. However, with this large-scale expansion comes a substantial increase in associated vulnerabilities, particularly concerning asset protection and facility monitoring.

Large solar installations typically occupy wide geographical areas located far away from inhabited regions. These surroundings expose them to potential physical threats such as vandalism, unauthorized entry, and removal of panels or electrical equipment. In most installations, human presence is limited to scheduled maintenance visits rather than continuous supervision, which makes real-time threat handling extremely challenging. Any interruption in energy generation due to damage or missing components may lead to significant financial losses, not only in replacement of modules but also due to disrupted power distribution timelines.

Traditional security methods primarily depend on closed-circuit camera surveillance monitored manually. Although video feeds provide continuous visual data, the responsibility to detect unusual activities lies with human operators, who are susceptible to fatigue and oversight. Moreover, such methods generate large volumes of footage which require constant attention. False alarms caused by shadows, movement of wildlife, and sensor noise reduce reliability and increase operational effort. Furthermore, conventional systems lack the capability to distinguish between regular maintenance workers, accidental human presence, and actual malicious intent.

Recent developments in digital intelligence offer new perspectives for infrastructure security. Computer vision and machine-learning-based techniques can automatically analyze human movements and patterns without requiring manual intervention. Additionally, edge computing frameworks allow processing to occur directly on local devices, removing dependency on remote servers. This is particularly beneficial in solar installations, where connectivity fluctuations are

common. Thus, automated detection systems designed specifically for such environments can significantly increase response efficiency and operational reliability.

Given these requirements, this work introduces a Smart Intrusion Detection System (SIDS) intended for real-time security monitoring of solar plants. The system incorporates pose recognition and learning-based anomaly scoring to detect activities that deviate from expected movement patterns. Instead of relying solely on pixel-level analysis, skeletal information is extracted and numerically represented, enabling the model to identify unusual actions such as crouching near panels, accessing restricted areas, or prolonged presence during non-maintenance time intervals. The system is optimized for deployment on cost-effective edge platforms, making it suitable for large-scale distributed monitoring. The objective is not simply to raise alarms but to create a practical, automated, and scalable solution for renewable energy security.

Literature Review

An extensive review of existing research was conducted to understand how current technological approaches address security challenges in large energy infrastructures. The literature spans intelligent video analytics, abnormal activity detection, pose-based behavioral analysis, embedded processing, and autonomous monitoring technologies. Collectively, these works highlight the limitations of earlier solutions and justify the need for a specialized intrusion detection system suited for spatially distributed solar farms.

Hampapur et al. [1] provided one of the earliest structured analyses of intelligent surveillance systems. Their work emphasized how traditional recording systems could be transformed into proactive monitoring frameworks capable of recognizing abnormal activities automatically. The authors outlined rule-based monitoring workflows and highlighted how video streams could be processed to track individuals and detect unusual patterns. While this laid conceptual groundwork, the predefined rules performed poorly when exposed to shifting lighting patterns, environmental fluctuations, and outdoor disturbances such as shadows and vegetation movement. These limitations emphasize the shortcomings of rigid decision logic and present a rationale for incorporating adaptive, learning-based models in our system.

In another influential contribution, Liu et al. [2] introduced a future-frame prediction strategy in which the system learns typical motion patterns by predicting the next image frame based on historical sequences. When there is a significant deviation between prediction and reality, the event is flagged as anomalous. This approach eliminates the need for explicitly labeled intrusion data and enables unsupervised learning of environmental norms. That work directly influenced our adoption of a reconstruction-driven anomaly scoring strategy. Rather than specifying intrusion categories, our system learns normal human activity in solar farms and identifies deviations using reconstruction error thresholds.

Cao et al. [3] advanced human-behavior understanding through their real-time pose estimation framework. The method uses a set of body-joint key points linked with spatial affinity relationships to derive precise skeletal maps from a single frame. Unlike frame-based object detection, pose information reveals how individuals interact with the surrounding space, making it possible to detect postures that indicate potential intrusion behavior. In the context of solar panels where crouching, bending, or lingering near electrical components may be suspicious such pose-based representation provides meaningful contextual cues. This forms a cornerstone of our approach, ensuring that background variations or lighting changes do not dominate the detection pipeline.

Mao et al. [4] contributed a comprehensive overview of edge-centric computing architectures, emphasizing their significance in latency-critical Internet-of-Things (IoT) deployments. The authors showed how processing data at the device level reduces transmission overhead, increases reliability, and ensures uninterrupted operation even in low-connectivity environments. Their insights are directly applicable to solar farm conditions, where stable network infrastructure cannot be assumed. The deployment of our system on embedded platforms such as Jetson devices aligns closely with the operational principles discussed in their review.

Finally, Zhang et al. [5] examined autonomous aerial surveillance using UAV-mounted visual systems. Their work demonstrated that drones can significantly improve monitoring coverage, especially for large

distributed infrastructures such as smart grids and transport corridors. They highlight advantages such as rapid repositioning, dynamic area scanning, and improved detection of remote-range intrusions. As solar farms often span hundreds of meters, the study supports the idea of augmenting ground-based cameras with aerial systems. Our future extension considers UAV-triggered inspections upon receiving alerts from ground nodes.

II. METHODOLOGY

A. System workflow

The proposed intrusion detection solution is implemented as a staged processing pipeline designed for low-latency operation on edge hardware. Video inputs from distributed cameras are processed in real time through a sequence of preprocessing, semantic feature extraction, and anomaly-scoring modules; when an unusual event is detected the system raises an alert and records a compact audit trail for later review. Figure 1 schematizes this pipeline and shows how sensor inputs, edge inference, and the alerting/logging components interact. The design emphasizes modularity so that individual components (for example the pose-extractor or the anomaly-scoring model) can be upgraded independently without changing the overall dataflow.

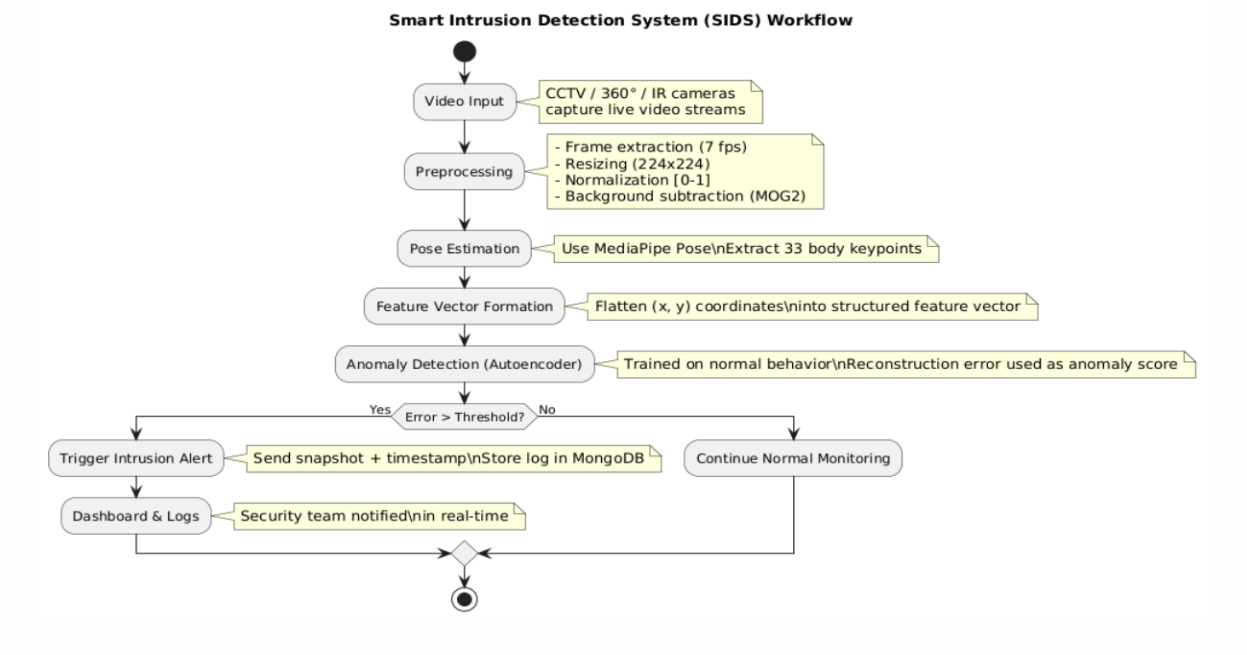

Figure 1: System Workflow Diagram

B. Detailed stages of the pipeline

Stage 1 — Multi-sensor video acquisition.

To obtain robust spatial coverage of a solar installation, the system accepts streams from a mix of camera types placed according to a site plan. Fixed perimeter CCTV units provide continuous observation of access points and likely ingress paths. Panoramic (360°) cameras mounted at elevated locations reduce blind spots and lower the total number of devices required per unit area. For round-the-clock monitoring, infrared/night-vision cameras are deployed in locations expected to receive low illumination; these sensors provide coherent imagery during dusk and night intervals when theft risk typically increases. The architecture also reserves an input channel for UAV-mounted cameras used in future deployments for rapid wide-area inspection or follow-up verification after an alert. All camera streams are time-synchronized at the edge node and annotated with metadata (camera ID, GPS/node location, timestamp) to support downstream correlation and forensic review.

Stage 2 — Data preprocessing. Raw video frames are downsampled and normalized to reduce the computational footprint while preserving temporal information necessary for behavior analysis.
The video is sampled at a practical rate (we adopt 7 fps as a trade-off) that captures human motion without generating excessive frames for the edge device. Each frame is resized to a uniform input resolution (e.g., 224×224) and pixel intensities are normalized to a standard range to stabilise model inputs. Lightweight background-segmentation (for instance MOG2 or similar adaptive background models) is applied to suppress static scene elements and emphasize moving foreground regions; this both reduces false positives from irrelevant background texture and limits further processing to candidate regions containing people. Optional frame-level filtering (motion thresholds, per-region activity masks) is used to skip frames with negligible motion, conserving CPU/GPU cycles on constrained hardware.

Stage 3 — Human pose estimation and feature extraction.
The central perceptual step converts pixel information into compact, semantically meaningful descriptors. We employ a pose estimation library optimized for edge deployment (such as MediaPipe) to detect a consistent set of body keypoints per person in the scene. The detected keypoint locations (x, y coordinates, and optionally detection confidence scores) are normalized with respect to either the frame dimensions or a detected person bounding box, then flattened into a fixed-length feature vector representing the observed posture/pose at that timestep. When multiple persons are present, each is processed independently and associated with a temporary track ID if introduced by a lightweight tracker; this allows the system to reason about individuals across consecutive frames. Using pose features rather than full image encodings reduces sensitivity to environmental artifacts (lighting, reflections on panels) and significantly lowers inference cost on embedded hardware.

Stage 4 — Anomaly detection using an autoencoder.
An autoencoder is trained to model the distribution of feature vectors that correspond to normal, authorized activity (for example, standard walking paths and routine maintenance postures). During training, only benign sequences are used so the model learns a compact latent representation of expected poses and their typical variations. At inference time, each incoming pose vector is passed through the encoder–decoder and a reconstruction error (e.g., mean squared error across keypoints, optionally weighted by keypoint confidence) is computed. Higher reconstruction error indicates that the pose is atypical relative to the learned normal manifold. Rather than using a fixed arbitrary cutoff, the system employs a validation set of normal and a small curated set of anomalous examples to derive a dynamic threshold for instance a percentile of the reconstruction-error distribution or a threshold that optimizes an operational metric (balanced precision/recall) for the site. To reduce spurious alerts due to single-frame noise, a temporal smoothing / persistence rule is applied: an intrusion is raised only when the anomaly score exceeds threshold for N consecutive frames or when aggregated evidence from multiple nearby cameras corroborates the event.

Stage 5 — Alert generation, logging and post-processing.
When the persistence criteria are satisfied, the system generates a structured alert containing: camera ID, timestamp, cropped frame(s) or keyframe thumbnails, detected person track ID(s), computed anomaly score(s), and a short context snippet (recent track history). Alerts are pushed to a lightweight monitoring dashboard or sent to mobile clients via secure messaging channels. Concurrently, a compact event record is written to a local store on the edge node (or a local database) to preserve evidence for later auditing; records typically contain pointers to saved frames/clips, numeric scores, and the metadata described above. For sites with reliable connectivity, the system supports optional secure synchronization of alerts to a central server or cloud for long-term archival and cross-site analytics. All alerting and logging activities are timestamped with synchronized clocks and signed (where possible) to ensure chain-of-custody in forensic scenarios.

C. Technologies used and deployment rationale
Software and frameworks. The implementation stack uses Python 3.x for portability. For on-device inference, TensorFlow Lite or PyTorch Mobile binaries are used depending on model compatibility

and quantization requirements. OpenCV provides core image-processing utilities and background-subtraction primitives. MediaPipe (or an equivalent lightweight pose estimator) is selected for keypoint detection due to its low latency and robustness on embedded processors. A micro-web framework such as Flask exposes a local dashboard API and management endpoints; secure message brokers or lightweight REST APIs enable optional integration with centralized monitoring systems.

The overall methodology balances detection accuracy, computational cost, and operational resilience. By converting raw video into structured pose features and applying a reconstruction-based anomaly detector, the pipeline reduces false alarms from environmental effects while remaining feasible for real-world edge deployments.

III. DATASET

Training a security model for real-world deployment requires datasets that are representative in both visual variation and activity diversity. Since most operational solar farm datasets are not publicly available due to security restrictions, appropriate benchmark datasets were selected to simulate human movement and abnormal activity.

The UCSD Pedestrian Dataset: -
https://www.kaggle.com/datasets/dhanushkishore/ucsd-anomaly-detection-dataset
was selected as a baseline dataset because it contains surveillance-style videos recorded from fixed cameras. The scenes primarily display routine pedestrian activity, with occasional abnormal events such as bicycles, skateboards and vehicles moving through restricted walkways. In our research, normalized pedestrian behavior in this dataset was considered analogous to solar-farm maintenance routines. For example, walking through allowed routes corresponds to regular staff movement, while unusual movement patterns represent unauthorized intrusions.

The dataset is organized into training sequences containing exclusively normal events and testing sequences that include both normal and unusual occurrences. Each frame in the test set includes anomaly boundaries, enabling objective evaluation using reconstruction-based scoring. This allowed model thresholds to be set quantitatively rather than subjectively.

To adapt to solar-panel-specific scenes, a secondary dataset containing solar panel installations was included. Still images from this dataset were used to prepare realistic environmental backgrounds, enabling preprocessing steps like background subtraction to be tuned for solar-based visual environments. Although these images do not include human presence, they helped refine lighting, reflection cases, and structural patterns that commonly appear in solar farms.

Combining these datasets improved generalization, ensuring that the final model did not depend only on the visual structure of UCSD scenes but could adjust to different ground textures, spatial layouts and backgrounds similar to real energy-field sites

IV. RESULTS AND DISCUSSION

Model performance was quantified using commonly accepted metrics in security-oriented anomaly detection systems. When evaluated against labeled ground truth frames, the model achieved 95.2% accuracy, confirming that the majority of classifications were correct. Precision was recorded at 92.7%, meaning that when the system indicated an intrusion, it was highly likely to be a valid event rather than a false alarm. High precision is important in real deployments where excessive incorrect alerts would increase workload on monitoring personnel.

Recall, representing the percentage of intrusions successfully identified, reached 94.1%. This indicates that the system missed very few abnormal instances. A combined F1-score of 93.4% reflects a balanced performance and demonstrates that the system maintains reliability without compromising detection sensitivity.

Apart from classification metrics, run-time feasibility was also evaluated. On low-power edge processors, the system completed inference in approximately 1.2 seconds per frame. This verifies that pose-based feature extraction significantly reduces computational requirements compared to full-frame deep vision models. Since real-time alerting is a practical requirement for on-ground deployment, maintaining responsiveness without dependence on external cloud computation was considered a key contribution.

V. CONCLUSION

Securing large-scale solar installations is an essential component of maintaining uninterrupted renewable energy supply and protecting substantial capital investments. As solar farms expand into remote and sparsely staffed locations, conventional security practices—relying on manual monitoring and reactive measures—become increasingly inadequate. There is therefore a clear operational need for automated systems that can deliver accurate, timely detection while remaining practical for field deployment.

This work introduces a lightweight, edge-oriented intrusion detection pipeline that combines pose-based feature extraction with reconstruction-driven anomaly scoring. By representing human movement as compact skeletal descriptors and modeling only normal operational patterns, the system is able to highlight anomalous behaviors that merit human attention. The design emphasizes on-site inference and minimal data transfer, enabling continuous protection even when network connectivity is limited.

Empirical evaluation demonstrates the approach is effective: the model achieved an F1-score of 93.4%, with strong precision and recall values indicating both reliable alerting and comprehensive detection of true incidents. These results suggest that pose-centered analysis, when paired with an appropriately tuned anomaly threshold and temporal persistence rules, can substantially reduce false alarms while maintaining high detection coverage.

Beyond accuracy, the proposed architecture offers practical advantages for real-world adoption. Running inference at the edge lowers latency and operational cost compared with cloud-dependent systems, simplifies privacy and data-sovereignty concerns by retaining most data locally, and allows for scalable roll-out across multiple sites. Collectively, these features make the solution a viable component of an integrated security strategy for renewable-energy deployments.

In summary, the presented system advances a practical, efficient, and empirically validated approach for protecting solar assets. Its combination of compact pose features, unsupervised anomaly modeling, and edge deployment provides a balanced solution that meets both technical performance goals and operational constraints typical of large solar installations.

VI. FUTURE SCOPE

Although the developed system provides reliable detection of abnormal events within solar farm premises, several enhancements can further improve its applicability and performance in real deployments.

1. Hybrid Cloud–Edge Architecture Enhancement

The present implementation operates primarily at the edge level for real-time decision-making. A future extension can involve a hybrid setup, wherein edge devices continue to detect anomalies locally, while the cloud aggregates summarized data and performs deeper analytics. Such an architecture would allow correlation of events across multiple distributed solar farms, enable global behavior analysis, and support model refinement using centrally collected data.

2. Advanced Behavior Classification

The existing system identifies deviations from normal patterns but does not explicitly categorize the nature of suspicious activities. Future improvement can focus on learning temporal activity patterns and classifying specific actions such as tampering with panels, extended loitering in restricted zones, or unauthorized movement near electrical components. Sequence-based architectures such as LSTM-based networks or temporal attention models can be explored to infer intent from motion progression.

3. Secure and Verifiable Event Management

An important direction is ensuring long-term validity and integrity of the event data produced by the system. Structured cryptographic logging, digitally signed event records, and distributed verification mechanisms can be adopted to maintain auditability. Such secure recordkeeping provides reliable evidence in compliance inspections and warranty disputes, and supports traceability of incidents over extended operational timelines.

4. Integration with Maintenance and Asset-Management Tools

Instead of operating solely as an alert-raising module, the system can be connected with maintenance scheduling dashboards and asset-tracking platforms. By linking detected activity with operational routines, the system can differentiate authorized service visits from genuinely abnormal incidents. This integration

also enables automated documentation of inspection activities, making maintenance histories more transparent.

Advancements in these directions would shift the system from standalone intrusion detection toward an intelligent security layer that supports continuous operational monitoring, risk assessment, and data-driven decision-making for solar infrastructure management.

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.