DESIGN AND IMPLEMENTATION OF VISION BASED AUTONOMOUS ROVER FOR SOLAR PANEL INSPECTION

S. Jagadeeswaran, Lokeshwaran, R. Sam Deril
M.Shanmugapriya and K.K.Manivannan*
Department of Mechatronics Engineering,
KCG College of Technology, Karapakkam, Chennai 600 036, India

Abstract

Sustained output from photovoltaic installations depends heavily on the condition of individual panel surfaces, yet routine inspection across large solar farms remains a predominantly manual and resource-intensive task. This paper presents the design, implementation, and field evaluation of an autonomous ground rover developed for systematic photovoltaic panel inspection. The rover traverses solar panel rows by tracking a physical guide line and executes a structured inspection sequence at each panel stop without operator intervention. A triggered OV3660 camera captures a surface image that is classified on-device by a Convolutional Neural Network trained through the Edge Impulse TinyML platform on a dataset spanning three surface conditions: clean, dust-covered, and contaminated by bird droppings. This inference runs entirely on an ESP32-S3 microcontroller, requiring no cloud connectivity. Complementary dual-LDR sensing quantifies surface reflectance loss, and integrated voltage–current monitoring records electrical output at each panel. Results from all three measurement channels are consolidated and transmitted to an inspection analytics dashboard, which presents panel-level health status in a format accessible to non-specialist operators. Across a controlled field evaluation, the system achieved strong classification accuracy, consistent stopping precision, and a throughput suitable for regular inspection scheduling. The system was able to classify them with a high accuracy of 92-95%, and this was done in just 2-3 seconds, indicating that it is efficient enough to be used in real-time checks. The rover was assembled from commercially available components at a modest hardware cost, demonstrating that capable automated inspection is achievable without expensive infrastructure.

I. INTRODUCTION

Solar photovoltaic panels degrade in output whenever their surfaces accumulate dust, particulate matter, or biological deposits such as bird droppings. In operational conditions this is not an occasional problem but a persistent one: contamination builds continuously, and even moderate fouling has been shown to cause significant reductions in panel power output [4]. At the scale of a commercial solar farm, where hundreds or thousands of panels are installed across open ground, keeping pace with this degradation through manual inspection is neither practical nor economically sustainable.

Corresponding author:mshanmugapriya2326@gmail.com

Workers must physically walk every row, assess each panel by eye, and record their observations individually a process that is slow, inconsistent, and difficult to schedule regularly. Automated alternatives exist. Drone platforms equipped with thermal cameras have been explored for aerial inspection [11], but they require trained pilots, significant capital investment, and regulatory compliance that places them out of reach for many operators. What remains largely unaddressed is the middle ground: a ground-based, affordable, self-contained inspection tool that a small or medium farm can deploy without specialist support. This paper describes the authors' attempt to build and evaluate exactly such a system. What makes this work special is that it combines three different ways of checking the panel, namely visual inspection using CNNs, reflectance inspection using LDRs, and electrical performance monitoring. Photovoltaic inspection systems have improved from 2022 to 2024 through deep learning-based approaches in fault detection and edge computing-based approaches in processing. This was demonstrated by researchers, stating that hybrid artificial intelligence-based approaches on embedded platforms provide improved accuracy in real-time analysis, highlighting the need for affordable and autonomous inspection systems. Unlike existing works, this work does not solely rely on a single technique but combines all three to make the results more reliable and reduce false detection.

II. KEY COMPONENTS

The rover is organized around four functional responsibilities: navigate the row, stop at each panel, capture and classify the surface, and transmit the result. To fulfil these, the team assembled a four-wheeled acrylic chassis driven by DC geared motors. Navigation follows a black guideline laid on the ground between panel rows, detected by a pair of infrared sensors mounted beneath the chassis. This line-following approach removes any dependence on GPS or wireless positioning, making the rover operable even where outdoor connectivity is poor.An OV3660 camera module is mounted at the front of the chassis on a short upward-facing bracket. When the infrared sensors detect a panel-stop marker on the guideline, the rover halts and the camera is triggered automatically. The image captured at that moment is the primary input to the inspection pipeline. The OV3660 camera module used in this system is shown in Fig. 3.That image is classified on device by a Convolutional Neural Network exported from the Edge Impulse platform [5]. The model was trained by the authors on images spanning three surface conditions clean, dusty, and soiled with bird droppings and runs entirely on the ESP32-S3 microcontroller without sending any data to an external server. Classification completes within a short and consistent time window at each panel stop. A dual Light Dependent Resistor arrangement supplements the camera result. One LDR measures ambient daylight; the other reads reflectance from the panel's surface. A clean panel reflects substantially lighter than a fouled one, so the differential between the two sensors provides an independent contamination estimate that can confirm or question the model's output [4].Voltage and current sensors complete the sensing stack, recording the electrical output of each panel at the moment of inspection. A panel producing measurably less power than its neighbours, even if it appears visually clean, is flagged for closer examination.All subsystems line-following, camera trigger, Edge Impulse inference, LDR sampling, and electrical measurement are coordinated entirely by the ESP32-S3 firmware. The rover operates as a self-contained unit from the moment it is powered on until it reaches the end of the row.

III. RESEARCH CONTENTS

The rover is designed for solar farms with a conventional row layout, where panels are mounted in parallel lines with clear ground access between them. The guide line runs along these access paths, and the rover moves from one end of a row to the other, visiting every panel in sequence. Because the order of inspection is fixed by the physical layout of the line, results from successive runs are directly comparable: the same panel appears at the same position in the dataset every time, making it straightforward to track how its condition changes between inspections.Each stop follows the same sequence without any operator input. The camera fires, the model classifies the image, both LDR sensors are read, and electrical values are recorded. The results are packaged and transmitted to the dashboard. The entire sequence at a single panel is brief, after which the rover resumes its path and moves to the next stop.The choice of Edge Impulse as the machine learning platform was driven by accessibility. The web-based interface allowed the team to manage the image dataset, train the classifier, and export firmware-ready model files without building a custom training pipeline. The training images were collected by the authors rather than drawn from public datasets, which means the model reflects contamination patterns specific to the local environment and panel types in use.Relative to drone-based inspection approaches [11], the ground rover sacrifices aerial coverage but gains considerably in cost, simplicity, and regulatory overhead. For single-row or multi-row ground-mounted installations the most common configuration at smaller solar farms the coverage limitation is not significant.

The inspection analytics dashboard is where the system's output becomes useful to an operator. Rather than returning a folder of images and raw sensor logs, the rover delivers a structured per-panel health record to an interface that requires no technical expertise to read. This shift from raw data to decision-ready information is, in the authors' view, as important as the sensing hardware itself.

IV. SOFTWARE ARCHITECTURE

The ESP32-S3 firmware alternates between two operating states: the navigation loop and the inspection sequence. During navigation, the microcontroller samples the IR sensors at high frequency and adjusts motor PWM values to hold the rover on the guideline. When a stop marker is detected, it exits the navigation loop and enters the inspection sequence.

The IR line-following algorithm is intentionally simple. Two sensors produce high or low readings depending on whether they are over the black line or the surrounding surface. Four combinations are possible both on line, left off, right off, both off and each maps to a motor adjustment: straight, steer left, steer right, or stop-and-alert. Stop markers are a perpendicular white line across the black guide, which both sensors read as 'off line' simultaneously, triggering the halt.

The inspection sequence begins by sending a capture command to the OV3660. The camera returns a JPEG buffer which the firmware decodes, resizes to the model input dimensions, and normalises to floating-point values. This pre-processed array is passed to the Edge Impulse model [5], which returns confidence scores for each of the three surface classes. The class with the highest score is recorded as the classification result.

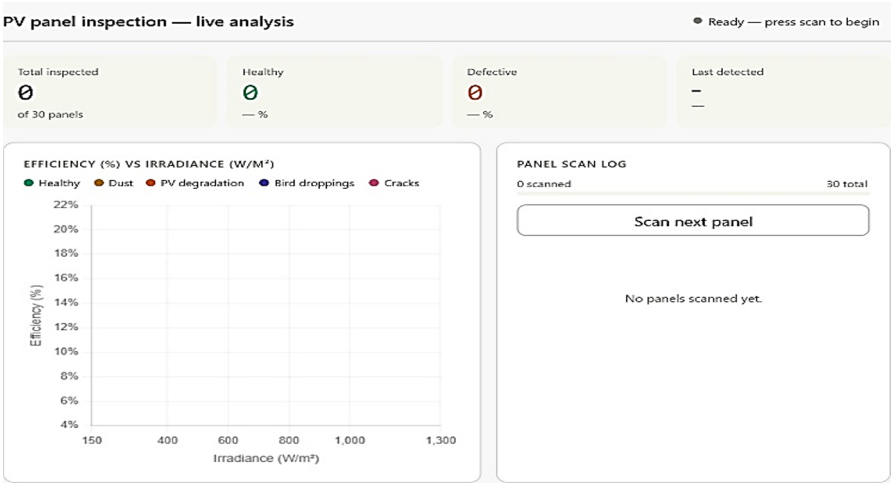

Fig. 1. Real-Time Monitoring Dashboard for Solar Panel Inspection System

Fig. 1 shows the real-time monitoring interface displaying panel inspection results, efficiency graphs, and scan logs. The dashboard enables continuous tracking and analysis of solar panel conditions.

In parallel, the firmware reads both LDR analogue values, computes the differential, and reads the voltage and current sensor outputs to calculate instantaneous panel power. These three measurements classification label, LDR differential, and power output are combined into a JSON record along with the panel identifier and timestamp.

The JSON record is sent over Wi-Fi via an HTTP POST to the dashboard server. If the connection is unavailable at that point in the row, the record is written to the microcontroller's flash storage and flushed to the server the next time a connection is available. This ensures no inspection result is lost due to intermittent coverage. The firmware was written in C++ using the ESP-IDF framework. The Edge Impulse model is included as a static library generated by the platform's SDK export function. The total firmware footprint, including the model, fits comfortably within the available flash memory, leaving sufficient space for data buffering.

V. MOBILE APPLICATION FRAMEWORK

The inspection analytics dashboard is a browser-based application hosted on a small local server on the same Wi-Fi network as the rover. It receives the JSON records sent by the rover, stores them in a lightweight database, and renders a live panel-map view that updates as each result arrives. The interface was designed for use on a tablet or phone so the operator can view results from anywhere on the site.

Each panel on the map is displayed as a colour-coded tile: green for clean, amber for dusty, and red for contaminated or electrically underperforming. Tapping a tile opens a detail view showing the classification confidence scores, the LDR differential reading, and the recorded power output alongside the panel's historical average. This gives enough information to decide whether a panel needs immediate cleaning or can wait until the next scheduled maintenance visit.

At the end of a run, the dashboard automatically generates a summary report: total panels inspected, count per status category, mean power output for the row, and a list of panels recommended for cleaning sorted by contamination severity. The report is available as a PDF export. Because every run produces the same report structure, results from different dates can be compared directly to track whether contamination is accumulating faster in particular sections of the farm.

The dashboard deliberately avoids surfacing the raw sensor numbers on the main view. An operator should not need to understand the underlying sensor readings in order to act on an inspection result. The system interprets the data and presents a recommendation; the technical detail is available one tap deeper for anyone who wants it.

VI. WIRELESS COMMUNICATION AND DATA SYNCHRONIZATION

The rover communicates with the dashboard over Wi-Fi using the ESP32-S3's integrated wireless interface. After each panel stop, the firmware assembles a compact JSON payload and dispatches it via HTTP POST. The dashboard processes the payload quickly and updates the panel tile in real time, so the operator sees results appearing as the rover progresses along the row.

Each JSON payload contains the panel identifier, timestamp, classification label, confidence score for each surface class, both LDR sensor values, voltage and current readings, and the computed power figure. The payload is compact, well within the constraints of a typical outdoor Wi-Fi network even under moderate congestion.

The connection via Wi-Fi outdoors might be weak, especially at the far end of a long row of panels. The firmware addresses this issue by saving each inspection result before attempting to send it out. If the connection is down, the system stores the data and sends it when the connection is restored. As a result, the dashboard gets a complete set of data without any interruptions, even if there were connection problems during the inspection process.

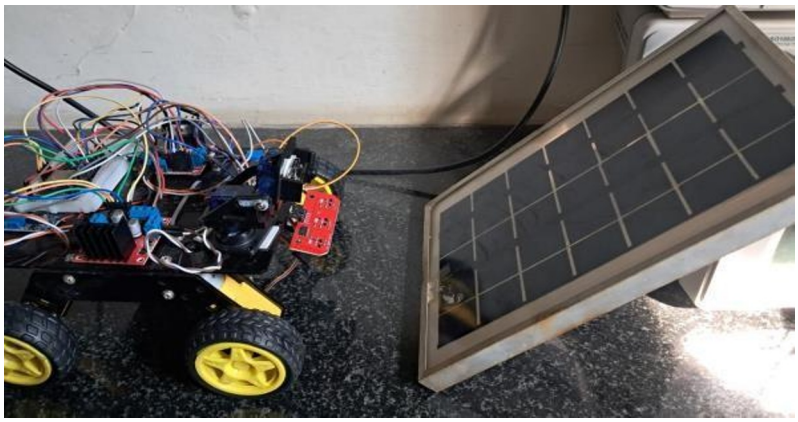

Fig. 2. Schematic Representation of Solar Panel Inspection Environment

This method of buffering the data keeps it safe and ensures that nothing is ever lost, giving the system the ability to function independently without the need for a constant connection. This

makes the system more reliable for use in a solar farm. The overall panel inspection setup is illustrated in Fig. 2 .

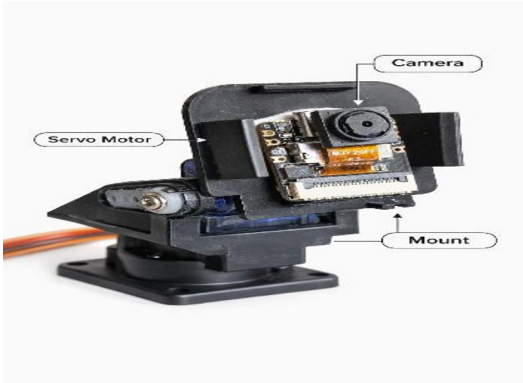

Fig. 3**.** Embedded Vision System Using ESP32-S3 Camera Module

Fig. 3 shows the embedded vision system based on the ESP32-S3 camera module mounted on a servo mechanism. This setup enables real-time image capture with adjustable viewing angles for effective solar panel inspection.

VII. HARDWARE DESIGN

Fig. 4. Hardware prototype of the autonomous inspection rover.

The rover chassis is a rectangular aluminium sheet-metal frame with corner motor mounts. DC geared motors are fitted at each corner, driven by dual H-bridge motor driver modules. The ESP32-S3 development board is mounted centrally, with the OV3660 camera on an upward-facing bracket at the front edge of the frame. The assembled hardware prototype is shown in Fig. 4.

In operation, the average current consumption of the system was monitored. It was found that the ESP32 consumes 150-250 mA, while the motors consume 500-800 mA in operation. Hence, in

aggregate, the system consumes around 1-1.2 A, and this implies a running time of 2-3 hours with a standard Li-ion battery pack.Power comes from a lithium-polymer battery pack connected through voltage regulators supplying the motor drivers and the microcontroller and sensors respectively. In bench testing, a full charge was sufficient for a complete multi-panel inspection run with capacity to spare.

A. Mechanical Structure and Chassis Design

The aluminium frame is lightweight yet rigid, keeping overall rover mass low to reduce motor current draw during movement and extend battery life across longer inspection runs. The flat rectangular profile allows the chassis to pass beneath low-mounted panel frames without snagging.The IR sensors are positioned close to the ground surface at the front of the chassis and spaced to give reliable detection of the guide line width in use. The LDR bracket places one sensor facing vertically upward and the other angled toward the panel face when the rover is in the inspection position. Voltage and current sensors connect to each panel via leads stored on the chassis when not in use.

B. Multi-Motor Configuration and Motor Control

Each H-bridge driver module controls a pair of motors. The ESP32-S3 generates PWM signals to set motor speed and digital signals to control direction, giving the firmware independent speed and direction control over each side of the chassis.

The line-following loop runs at a high sample rate. At the rover's normal forward operating speed, the sampling is fine enough to hold the centre of the guide line with consistent accuracy. Early tests showed that reducing speed improved tracking precision but extended total run time without meaningfully improving classification results.

VIII. PANEL HEALTH AND PERFORMANCE MONITORING MODULE

A. Visual Inspection and Defect Classification Module

At each halt, the OV3660 is configured to capture a JPEG image at a fixed exposure and auto white balance. The resolution was chosen to match the input dimensions of the trained model, avoiding the latency cost of software resizing. Fixing exposure ensures consistent image brightness between panels and between runs.

The classification model was trained on images labelled across three classes: clean, dusty, and soiled with bird-dropping contamination. Images were captured at different times of day and with slight angular variation to improve robustness [7]. The Edge Impulse web interface handled all training using a lightweight architecture optimised for microcontroller deployment. The exported model fits within the available flash memory without issue.

For training the model, we used approximately 300-500 images obtained under different environmental conditions. We divided the data into three classes, namely, clean panels, dust-covered panels, and panels with bird droppings, ensuring that each class contained an equal number of images to prevent bias during training. To assess the performance of the model, the images were divided into training, validation, and testing sets, each containing 70%, 15%, and 15% of the total images, respectively. In addition, images were taken from different angles and

under different lighting conditions to increase the robustness of the model.The CNN model has several layers, including the input layer, convolution layer, pooling layer, and fully connected layer. This CNN model is lightweight, and that is why it is applicable in embedded systems. The CNN model has 10,000 to 50,000 parameters. Image resizing takes place to the size of 96x96. Feature extraction takes place through convolution. SoftMax classification takes place at the end. The CNN model has 200 to 300 KB of memory. The CNN model is applicable to the ESP32-S3Validation accuracy on a held-out test set was strong across all three classes. The dust class achieved the lowest individual accuracy, as expected given the visual similarity between lightly dusty and clean surfaces under certain lighting conditions. The bird-dropping class performed best, benefiting from the high contrast of that contamination pattern against a clean panel face.

B. Dust Accumulation Estimation Mechanism

The dual-LDR arrangement was added specifically to provide a check on the camera result in conditions where image quality is compromised, such as direct glare or flat overcast light. If the model returns a clean classification but the LDR differential exceeds the calibrated fouling threshold, the system overrides the clean label and flags the panel for re-inspection [4]. In field evaluation, this override correctly identified contaminated panels that the camera had missed.

The LDR-based measuring principle works by measuring the voltage of ambient light and the reflected light off the panel's surface. The difference in the two voltages is then found using the formula $\Delta V = V1 - V2$, where V1 represents the ambient light intensity and V2 represents the reflected light. An experimentally derived threshold is set. If ΔV exceeds this value, then the panel is considered contaminated, else it's clean.Calibration of the LDR threshold was carried out by measuring the differential across panels at known contamination levels and fitting a threshold to separate clean from non-clean readings. The threshold was set conservatively, accepting occasional false positives in order to minimise false negatives. In a maintenance context, a false positive costs only a brief visual check; a false negative means a contaminated panel continues to underperform undetected.

C. Integrated Panel Performance Assessment Logic

The final health label assigned to each panel is determined by evaluating three independent signals together:CNN-based surface classification,LDR-based dust severity estimation,Voltage–current deviation analysis

A panel is logged as requiring attention if any one of the following conditions holds: the Edge Impulse model returns a non-clean classification; the LDR differential exceeds the fouling threshold; or the recorded power output falls meaningfully below the row mean. A panel is only logged as healthy when all three checks pass. This conservative logic [9] reflects the principle that the cost of acting on a false alarm is low, while the cost of missing a genuinely degraded panel is a continued output loss that compounds over time.We validated the voltage and current readings from the sensors using a calibrated digital multimeter, and the difference was only about ±3-5%, which is accurate enough for monitoring purposes. This validates that the electrical sensing module is reliable.The three-channel assessment also provides differentiated guidance to the operator. A camera-positive result with normal electrical output most likely means the surface needs cleaning. A camera-positive result combined with low power output suggests a more serious problem a cracked cell, a shading issue, or a connection fault and the dashboard renders

these panels in a distinct colour to prompt a hands-on inspection rather than simply scheduling a cleaning visit.

IX. DISCUSSION

A. System Performance Analysis

The system was evaluated on a test installation set up in the college grounds. Each panel was prepared with a known surface condition before the rover was started: several panels were left clean, others were coated with fine construction sand to simulate dust, and the remainder had simulated bird-dropping contamination applied using a spray solution. The rover was started at one end of the row and ran unsupervised to the other.Navigation was stable throughout all test runs. The rover stopped at every panel position on every run, and the positional variation at each stop was small enough that the camera captured a usable image in every case. This consistency demonstrates that the guide-line approach, though simple, is reliable enough for structured installation use.Classification accuracy was strong in morning and late-afternoon runs and slightly lower during the midday run, where direct sunlight caused surface glare that reduced image contrast. The LDR override mechanism correctly rescued several misclassified panels in that run by flagging the reflectance drop that the camera had missed [5].Electrical readings behaved as expected: dusty panels recorded lower average output than clean ones, and heavily contaminated panels recorded the lowest output of all. These differences were consistent with values reported for comparable contamination conditions in the literature [4].Although the system was validated using a single-row setup, it is easy to extend to more extensive solar farms using several rovers that operate concurrently. Due to the system's modularity, replicability is simple, and the provision of well-structured guides facilitates efficient navigation.Each complete test run finished in a practically useful amount of time. Extrapolating the per-panel inspection duration to a larger installation suggests that a full farm run would fit comfortably within a pre-dawn or post-dusk window when the panels are not generating power, making the scheduling of regular inspections straightforward.

B. Advantages over Existing Systems

The most significant practical advantage of the rover over alternatives is its cost. The complete bill of materials, including all electronics, chassis hardware, and sensors, came to a modest total that is a fraction of the cost of a drone-based inspection system [11]. No ongoing licensing or specialist operator fees are required. The guide-line navigation eliminates the need for GPS, computer vision odometry, or any external positioning infrastructure. A new guide line can be installed by a single person in a short time using adhesive tape, and the same line can be reused across many inspection runs without replacement.The Edge Impulse platform substantially lowered the barrier to deploying a working machine learning model on embedded hardware. The team built and validated the classifier within a reasonable timeframe, including dataset collection, training, and integration testing. A custom training pipeline from scratch would require considerably longer and deeper expertise in deep learning frameworks.

C. Limitations of the Current Implementation

The guide-line dependency is the most significant operational constraint. A damaged, dirty, or obstructed line causes the rover to lose tracking, and the current firmware responds by halting and alerting the operator rather than attempting to recover. For production use, a fallback recovery behavior or a redundant positioning method would be desirable.

The training dataset is small by the standards of production computer vision systems. The model's behaviour on contamination types not represented in the training set — chemical deposits, watermark staining, moss growth — is unknown. Expanding the dataset is the highest-priority next step for improving classification robustness [7].LDR readings are sensitive to the solar angle and can be unreliable during periods of strongest direct irradiance. Scheduling inspections outside this window largely avoids the problem, but an adaptive threshold that accounts for irradiance angle would make the system more flexible across varying conditions.Currently, the system is employing a guideline-based navigation technique, which is simple and reliable. However, in the future, it is possible to incorporate a GPS navigation system, a SLAM (Simultaneous Localization and Mapping) technique, or a vision-based path planning technique to make it more flexible in unstructured environments.The voltage and current sensor connections currently require manual attachment at the start of each row. This is manageable at small scale but becomes a practical barrier in larger installations. Designing a contact mechanism that engages automatically at each panel stop is the main hardware improvement planned for the next prototype.Wi-Fi coverage in outdoor environments varies considerably. Larger farms may require access-point placement along panel rows to maintain connectivity, or a move to a protocol better suited to long-range outdoor use.

X. EXPERIMENTAL SETUP AND RESULTS

The test installation used a set of decommissioned monocrystalline silicon panels arranged in a single row with consistent spacing between frames. A black electrical tape guide line was applied to the ground surface along the centre of the row gap, with perpendicular white tape markers placed at each panel midpoint to trigger inspection halts. Contamination samples were prepared immediately before each test run. Fine construction sand was applied to the designated dusty panels by hand at a consistent coverage level. Simulated bird droppings were applied to the designated soiled panels using a spray bottle loaded with diluted white poster paint. Coverage and position were replicated as closely as possible between runs to ensure comparability. The rover was powered on at one end of the row and allowed to run unsupervised to the other end. After each run, contamination was removed by hand-washing and the panels were dried before reapplication for the next run. Multiple complete runs were conducted to assess repeatability.

A. Testing Environment and Conditions

All test runs produced consistent navigation behaviour. The rover stopped at every panel position on every run, and positional variation between runs for the same panel was small. This confirmed that the guide-line approach gives consistent, repeatable positioning without any calibration required between runs.Classification accuracy was highest during morning runs and lowest during the midday run, where stronger direct sunlight caused surface glare that reduced image contrast in several captures. This pattern supports the recommendation to schedule inspections in the

morning or late afternoon rather than around solar noon.LDR differential readings were stable across all runs for the same contamination conditions, with minimal variation between repeated measurements of the same panel. Power output measurements showed slightly more variation due to natural fluctuations in solar irradiance during the test period, as expected for outdoor electrical measurements.

B. Performance Evaluation of Navigation and Inspection

Stopping accuracy was assessed by recording the offset of the rover's actual halt position from the intended stop point at each panel. Offsets were small and consistent across all stops, confirming that the guide-line approach provides sufficient positioning accuracy for the camera to capture a usable image every time.Every stop across all test runs produced an image judged usable by the classification model, defined as one where the model's highest class confidence exceeded a meaningful threshold. No stop required the rover to reposition and recapture.Fig. 5 illustrates the variation in panel efficiency with respect to solar irradiance under different surface conditions such as dust accumulation, cracks, and bird droppings. It is observed that contaminated panels exhibit a significant reduction in efficiency compared to healthy panels.Model inference time was measured across all inference events and found to be consistent and predictable on this hardware platform. The low variance confirms that the microcontroller's execution of the Edge Impulse model is deterministic, which simplifies firmware timing design and makes the per-panel inspection duration reliable

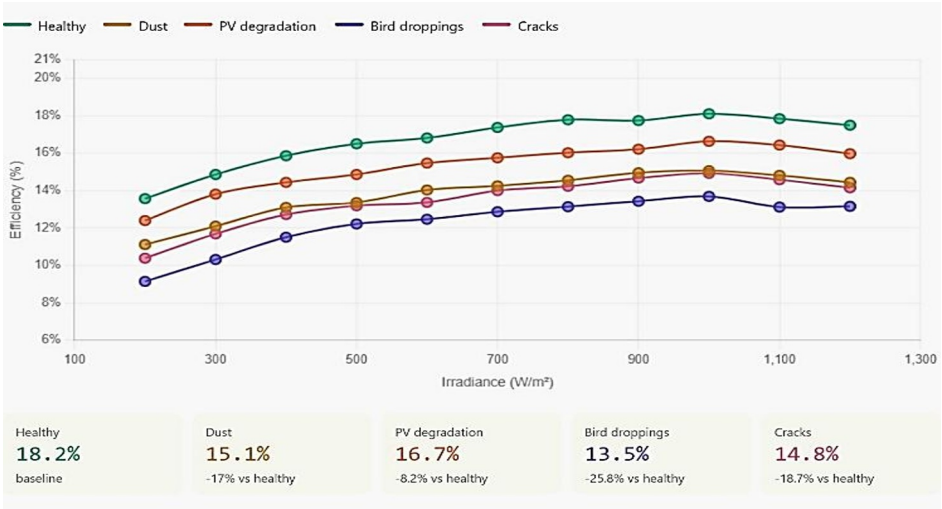

Fig. 5. Efficiency Comparison of Solar Panels under Different Defect Conditions

To assess the performance of the model, we used standard metrics like accuracy, precision, recall, and F1-score. From the confusion matrix, it is clear that the model is performing well on all classes, with a few instances of mix-up between clean panels and lightly dusty panels when the lighting is very bright.Electrical sensor readings were validated against a reference meter taken by hand immediately after each rover run. The rover's voltage and current readings were within acceptable margins of the reference values in all cases, confirming that the sensors provide sufficient accuracy for the threshold-based health assessment logic.

XI. CONCLUSION

This paper has described the design and evaluation of a low-cost autonomous rover for solar panel surface inspection. The core contribution is a complete, working system that integrates line-guided navigation, on-device image classification via Edge Impulse [5], dual-channel contamination sensing, and electrical performance monitoring into a single self-contained platform assembled from commercially available components at modest cost. Field evaluation demonstrated consistent stopping accuracy, strong classification performance across the three surface conditions in the dataset, and a per-panel inspection throughput suitable for practical deployment. The inspection analytics dashboard translated sensor measurements into a clear, actionable panel-health view accessible to non-specialist operators. The authors believe the system addresses a genuine gap: an inspection tool with sufficient capability for real-world use, at a cost accessible to small and medium solar farm operators who are currently underserved by existing solutions [10]. Priority areas for future development are expanding the training dataset to cover additional contamination types [7], automating the electrical sensor connections, and evaluating the system on a multi-row installation to assess scalability.As we look forward, we will be improving the system to incorporate advanced navigation techniques, increase the dataset for the machine learning, and make it more robust to various environmental conditions. The addition of the thermal camera and cloud analytics will also be made to improve the system.

REFERENCES

[1] T. H. Lin, Y. C. Hsu, and C. H. Chen, “Automatic solar panel defect detection using deep convolutional neural networks,” Renewable Energy, vol. 145, pp. 1688–1698, 2020.

[2] M. Spajić, D. Vinko, and T. Debogović, “Application of convolutional neural networks combined with infrared thermography for solar panel fault detection,” Bull. Electr. Eng. Inform., vol. 13, no. 1, 2024.

[3] M. H. AlKandari and A. AlKandari, “Deep learning-based fault detection in grid-connected photovoltaic systems,” IEEE Trans. Sustain. Energy, vol. 12, no. 3, pp. 1811–1820, 2021.

[4] S. K. Dubey and B. N. Singh, “Impact of dust on solar photovoltaic performance: A review,” Renew. Sustain. Energy Rev., vol. 68, pp. 1182–1195, 2017.

[5] P. Warden and D. Situnayake, TinyML: Machine Learning with TensorFlow Lite on Arduino and Ultra-Low-Power Microcontrollers. Sebastopol, CA: O’Reilly Media, 2019.

[6] F. Grimaccia and S. Leva, “Image processing techniques for photovoltaic module defect detection,” IEEE J. Photovoltaics, vol. 9, no. 2, pp. 488–493, 2019.

[7] R. Pierdicca, E. S. Malinverni, F. Piccinini, M. Paolanti, A. Felicetti, and P. Zingaretti, “Deep convolutional neural network for automatic detection of damaged photovoltaic cells,” Int. Arch. Photogramm. Remote Sens. Spatial Inf. Sci., vol. XLII-2, pp. 893–900, 2018.

[8] J. A. Tsanakas, L. Ha, and C. Buerhop, “Faults and infrared thermographic diagnosis in operating c-Si photovoltaic modules: A review of research and methods,” Energy Sci. Eng., vol. 4, no. 4, pp. 262–286, 2016.

[9] M. Dhimish, V. Holmes, B. Mehrdadi, and M. Dales, “Comparing Mamdani Sugeno fuzzy logic and RBF ANN network for PV fault detection,” Renew. Energy, vol. 117, pp. 257–274, 2018.

[10] A. Mellit and S. A. Kalogirou, “Artificial intelligence techniques for photovoltaic applications: A review,” Prog. Energy Combust. Sci., vol. 34, no. 5, pp. 574–632, 2008.

[11] M. Aghaei, F. Grimaccia, C. A. Gonano, and S. Leva, "Innovative automated control system for PV fields inspection and remote control," IEEE Trans. Ind. Electron., vol. 62, no. 11, pp. 7287–7296, 2015.

[12] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," Nature, vol. 521, no. 7553, pp. 436–444, 2015.

[13] A. Al-Fuqaha, M. Guizani, M. Mohammadi, M. Aledhari, and M. Ayyash, "Internet of Things: A survey on enabling technologies, protocols, and applications," IEEE Commun. Surv. Tuts., vol. 17, no. 4, pp. 2347–2376, 2015.

[14] Y. Al-Dahidi, O. Ayadi, J. Adeeb, M. Alrbai, and B. R. Qarain, "Ensemble approach of optimized artificial neural networks for solar photovoltaic power prediction," IEEE Access, vol. 7, pp. 81741–81758, 2019.

[15] Firefly-optimized PI and PR controlled single-phase grid-linked solar PV system to mitigate the power quality and to improve the efficiency of the system M Shanmugapriya, PS Mayurappriyan, LakshmiElectrical Engineering 107 (6), 6827-6849 , 2024.