Artificial Intelligence Research and Applications (AIra)
https://www.ijaira.org/ Online ISSN: 3108-3889
H. Merabet et al., Vol. 2, No. 3, September, 2026 Research Article

Autonomous UAV Inspection Powered by Generative AI for Photovoltaic Power Loss Analysis

Hichem Merabet*, Aziz Boukadoum**, Tahar Bahi***‡

* Research Center in Industrial Technologies (CRTI) P.O. Box 64, Cheraga, Algeria.

** Department of Electrical Engineering; LABGET Laboratory Larbi Tebessi University, Algeria.

*** Department of Electrical Engineering, LASA Laboratory. Badji Mokhtar-Annaba University B.O. Bloc 12, Algeria.

(h.merabet@crti.dz, aziz.boukadoum@univ-tebessa.dz, tahar.bahi@univ-annaba.dz)

‡ Corresponding Author: Tahar Bahi, Department of Electrical Engineering, LASA Laboratory. Badji Mokhtar-Annaba University B.O. Bloc 12, Algeria, tahar.bahi@univ-annaba.dz.

Received: 10.02.2026, Revised: 07.03.2026, Accepted: 27.03.2026

Abstract- Photovoltaic power generation in large-scale solar farms is strongly affected by localized factors such as soiling, partial shading, and component degradation, which are often difficult to quantify accurately using conventional monitoring systems. While UAV-based inspection techniques have shown promise for visual anomaly detection, they typically fail to translate observed defects into measurable impacts on energy production. In this work, we propose a UAV-assisted photovoltaic power loss estimation framework based on large language models, designed to bridge the gap between visual observations and quantitative energy assessment. By integrating UAV imagery, environmental context, and system-level metadata within a unified multimodal reasoning pipeline, the proposed approach estimates panel- and farm-level power losses directly in terms of energy output. Extensive experiments conducted on photovoltaic datasets combining UAV-acquired images and SCADA-based ground truth demonstrate the effectiveness of the proposed framework compared to vision-only and deep learning baselines. The results show a consistent reduction in power loss estimation error, achieving a relative improvement of 18.7% in MAE and 15.2% in RMSE over the strongest baseline, along with a higher correlation with actual power production across varying irradiance conditions. Additional analyses confirm the robustness of the proposed method under partial observations and heterogeneous environmental settings. These findings highlight the potential of LLM-driven aerial monitoring as a practical and energy-aware solution for photovoltaic performance assessment and operational decision-making.

Keywords: Photovoltaic, power generation, power loss, UAV, large language models.

1. Introduction

Photovoltaic (PV) systems have made solar energy a central component of modern renewable power generation [1,2]. Despite continuous advances in panel efficiency and system design, the actual energy yield of photovoltaic farms often deviates from expected performance due to localized factors such as soiling, partial shading, component degradation, and environmental variability [3,4]. Accurately identifying and quantifying these sources of power loss remains a critical challenge for reliable energy production and cost-effective operation and maintenance of PV installations [5–7].

Conventional PV monitoring systems primarily rely on fixed sensors and SCADA measurements to assess system performance at the inverter or farm level [8–10]. While

Cite this article as: H. Merabet, A. Boukadoum, and T. Bahi, “Autonomous UAV Inspection Powered by Generative AI for Photovoltaic Power Loss Analysis,” *Artificial Intelligence Research and Applications (AIra)*, Vol. 2, No. 3, pp. 140–148, September, 2026, https://doi.org/10.20508/ckkk9598.

effective for detecting global anomalies, these approaches offer limited spatial resolution and are often unable to localize or explain energy losses at the panel or string level. As a result, operators frequently lack actionable insights into the physical causes of reduced power generation, leading to delayed interventions and suboptimal maintenance strategies [11,12].

Unmanned aerial vehicles have emerged as a flexible and low-cost solution for large-scale photovoltaic monitoring. By acquiring high-resolution visual and thermal data, UAV-based inspection systems enable detailed observationof panel-level conditions across extensive solar farms. Existing approaches predominantly focus on detecting visual defects or classifying fault types from UAV imagery [13,14]. However, these methods typically treat inspection as a computer vision task and do not establish a direct connection between detected anomalies and their quantitative impact on energy production. Consequently, the operational relevance of such inspections remains limited from an energy management perspective. At the same time, recent advances in large language models have demonstrated strong capabilities in integrating heterogeneous information and performing contextual reasoning across multiple data modalities [15,16]. These models offer the potential to move beyond isolated visual analysis by jointly interpreting image-derived features, environmental conditions, and system metadata within a unified framework. When applied to photovoltaic monitoring, this capability opens new opportunities to translate qualitative observations into quantitative assessments of power loss that are directly meaningful for energy system operation.

Despite this potential, the integration of UAV-acquired visual data with energy-aware reasoning remains largely unexplored in the context of photovoltaic power generation. Most existing studies stop short of linking aerial observations to measurable energy outcomes, leaving a gap between monitoring technologies and actionable energy analytics. Addressing this gap requires methods that explicitly target power loss estimation rather than defect detection alone, and that align aerial sensing outputs with real production data.

In this work, we propose a UAV-assisted photovoltaic power loss estimation framework based on large language models, designed to bridge visual monitoring and quantitative energy assessment. By combining UAV imagery, environmental context, and photovoltaic system characteristics within a unified multimodal pipeline, the proposed approach estimates power losses directly in terms of energy output at both panel and farm scales. The main contributions of this work are summarized as follows:

- We introduce an energy-aware photovoltaic monitoring framework that integrates UAV-acquired imagery with system-level and environmental information to estimate real power losses rather than visual defects alone.
- We design a multimodal reasoning mechanism that translates visual and contextual observations into quantitative energy loss estimates aligned with photovoltaic power generation.
- We validate the proposed approach on photovoltaic datasets with ground-truth production data and demonstrate its effectiveness compared to vision-based and deep learning baselines.

2. Related Works

Understanding and quantifying energy losses in photovoltaic (PV) systems has long been a central research topic, as these losses directly limit power conversion efficiency and system reliability. Early studies primarily focused on physics-based modeling to analyze intrinsic loss mechanisms at the material and device levels. Shen et al. [17] presented a comprehensive energy distribution model for crystalline silicon PV modules, quantifying losses associated with carrier generation, transportation, recombination, and module-level effects. Their analysis revealed that more than 70% of the incident solar energy is dissipated as heat, highlighting temperature-driven recombination as a dominant degradation factor. Similar physics-driven investigations have been conducted for emerging photovoltaic technologies. Da et al. [18] quantified energy losses in planar perovskite solar cells through coupled optical–electrical modeling, identifying thermalization, optical absorption inefficiencies, and recombination as major contributors. These studies provide fundamental insights into loss origins but rely on detailed material parameters and laboratory measurements, limiting their applicability for large-scale field monitoring.

Recent advances in organic photovoltaics have further emphasized the importance of minimizing non- radiative energy losses through material and molecular engineering. Ling et al. [19] demonstrated that asymmetric small-molecular acceptors can significantly reduce energy losses by improving intermolecular packing and charge extraction, achieving power conversion efficiencies exceeding 20%. Complementary works have explored thermally activated delayed fluorescence materials [20] and side-chain isomerization strategies [21] to suppress non- radiative recombination and refine exciton diffusion pathways. Fluorinated and methylated acceptor designs have also been shown to reduce charge recombination and energy loss in OPVs [22]. While these studies advance device-level efficiency, they remain focused on controlled fabrication and characterization, offering limited support for operational loss diagnosis in deployed PV farms.

Beyond intrinsic material losses, environmental and operational factors play a critical role in real-world photovoltaic performance. Among them, soiling has been identified as a major source of energy loss at the system level. Fernández Solas et al. [23] provided the first large-scale assessment of photovoltaic soiling losses across Europe, revealing average annual losses ranging from 0.9% to over 5% depending on cleaning effectiveness and regional conditions. Their work underscores the strong spatial and temporal variability of soiling-induced losses and highlights the need for continuous, site-specific monitoring strategies rather than static models or periodic inspections. Control-oriented approaches have also been proposed to mitigate energy losses during PV system operation. Djilali et al. [24] introduced an enhanced variable step-size perturb-and-observe MPPT algorithm, achieving up to an 80% reduction in power ripples under rapidly varying irradiance conditions. Although such control strategies effectively reduce conversion losses at the inverter level, they do not explicitly identify or localize physical degradation sources, limiting their diagnostic capabilities.

More recently, data-driven and learning-based methods have emerged as scalable alternatives for photovoltaic assessment. With the growing availability of aerial and satellite imagery, computer vision techniques have been applied to detect PV installations and estimate system characteristics. Guo and Weng [25] proposed the PVAL framework, leveraging large language models (LLMs) to assist in photovoltaic panel detection from satellite imagery through task decomposition and structured prompting. Their work demonstrates the potential of LLMs for improving robustness and generalization in large-scale PV analysis. However, their focus remains on panel detection and localization rather than quantitative energy loss estimation or physical degradation reasoning. LLM-based forecasting approaches have also begun to appear in the photovoltaic domain. Lin and Yu [26] introduced Time-LLM, a large language model reprogrammed for distributed photovoltaic power forecasting using historical time-series data. By aligning numerical power sequences with language-like representations, their method achieves strong forecasting accuracy without relying on external weather data. Nevertheless, such approaches primarily address temporal prediction and do not incorporate visual inspection data or explicit physical degradation cues.

In summary, existing works on photovoltaic energy loss analysis can be broadly categorized into physics-based modelling, material and device optimization, environmental loss assessment, control strategies, and emerging data- driven approaches. While physics-based and material-level studies provide deep mechanistic understanding, they lack scalability for operational monitoring. Conversely, recent learning-based methods offer scalability but often neglect physical interpretability or multimodal reasoning. To the best of our knowledge, few studies explicitly combine UAV-based visual inspection with energy-aware reasoning to directly estimate photovoltaic power losses. This work bridges this gap by integrating multimodal UAV observations with structured LLM-based reasoning, enabling scalable, physically grounded, and field-deployable photovoltaic loss estimation. Table 1 presents the state-of-the-art summary.

Table 1. State of the art summary.

Reference	Year	Approach	Advantages	Limitations
Ling et al. [19]	2025	Material-level optimization using asymmetric acceptors for ternary organic photovoltaics	Significantly reduces non-radiative energy loss; achieves high power conversion efficiency	Focused on laboratory-scale devices; not applicable to deployed PV systems
Fernández Solas et al. [23]	2025	Large-scale empirical analysis of photovoltaic soiling losses across Europe	Provides geographically grounded loss estimates; highlights real-world environmental impacts	Relies on statistical aggregation; lacks localized fault diagnosis
Guo and Weng [25]	2025	LLM-assisted photovoltaic panel detection from satellite imagery	Improves robustness and scalability of PV asset identification	Targets detection only; does not estimate energy loss or degradation severity
Lin and Yu [26]	2025	Time-series forecasting using reprogrammed large language models	Accurate PV power forecasting without explicit weather data	Does not leverage visual inspection; limited physical interpretability
Djilali et al. [24]	2025	Enhanced MPPT control with variable step-size perturb-and-observe	Reduces dynamic power losses under irradiance fluctuations	Addresses inverter-level losses only; no degradation localization
Zhang et al. [20]	2025	TADF-based material engineering for reduced energy loss in photovoltaics	Suppresses non-radiative recombination; improves exciton utilization	Restricted to organic photovoltaics; not suitable for system-level monitoring
Wang et al. [21]	2025	Isomerization-based molecular design for minimizing non-radiative losses	Improves charge transport and molecular packing	Focuses on molecular-scale optimization; no field deployment
Wang et al. [22]	2025	Chemical modification of acceptors to suppress charge recombination	Reduces energy loss and enhances OPV efficiency	Does not address operational degradation or large-scale systems
Shen et al. [17]	2020	Physics-based modeling of energy distribution and losses in photovoltaic modules	Provides detailed quantification of intrinsic loss mechanisms	Requires detailed physical parameters; unsuitable for real-time monitoring
This paper	2026	Autonomous UAV Inspection Powered by Generative AI for Photovoltaic Power Loss Analysis	The results show a consistent reduction in power loss estimation error,	Robustness of the proposed method under partial observations and heterogeneous environmental settings.

3. Preliminaries

This section introduces the fundamental concepts, notations, and problem formulation used throughout this work.

Definition 1 (Photovoltaic system): A photovoltaic (PV) system is defined as a set of solar panels organized into strings and inverters, deployed over a geographical area for power generation. Let $\mathcal{P} = \{p_1, p_2, \dots, p\_N\}$ denote the set of PV panels, where each panel p_i is characterized by its rated capacity, orientation, tilt angle, and operational status. The aggregated power output of the system at time t is denoted by P_t and is measured from SCADA data.

Definition 2 (Expected and actual power output): The expected power output *P_t^exp* corresponds to the theoretical production under ideal operating conditions, determined by irradiance, temperature, and system specifications. The actual measured power output *P_t^act* reflects real-world operating conditions and may deviate due to environmental and physical factors.

Definition 3 (Photovoltaic power loss): The power loss at time t is defined as the difference between expected and actual power output. This quantity can be computed at panel, string, or farm level and serves as the primary target variable in this study.

Definition 4 (UAV-based observation): A UAV-based observation at time t is defined as a tuple composed of a set of captured images and associated metadata.

Definition 5 (Visual feature representation): Each UAV image is mapped to a visual feature vector using a feature extraction function. The aggregated visual representation at time t is obtained by averaging all extracted features.

Definition 6 (Environmental and system context): Environmental and system context is represented by a vector including meteorological variables, system configuration parameters, and operational metadata.

Definition 7 (Multimodal photovoltaic representation): The multimodal representation is obtained by concatenating visual features and contextual information. This representation serves as the input for energy-aware reasoning.

Definition 8 (Energy-aware reasoning function): An energy-aware reasoning function maps the multimodal representation to the predicted power loss.

Definition 9 (Power loss estimation task): The objective is to minimize the error between predicted and ground-truth power losses over a given time horizon.

4. Proposed Method

This section details the proposed UAV-assisted framework for photovoltaic power loss estimation using large language models. The objective is to establish a principled pipeline that maps aerial visual observations and contextual information to quantitative energy loss estimates. Unlike defect-centric inspection methods, the proposed approach explicitly focuses on power loss estimation, enabling energy-aware monitoring and decision support.

4.1. Problem Definition

Consider a photovoltaic farm composed of N panels organized into strings and arrays. Let T denote a set of discrete time instances at which UAV inspections are conducted. For each time t ∈ T, the system aims to estimate the photovoltaic power loss relative to the expected nominal production under current operating conditions.

For a given time t, UAV observations and contextual information are available, and the corresponding ground- truth power loss ΔPt is obtained from supervisory control and data acquisition (SCADA) measurements. The learning objective is to estimate ˆΔP t such that it accurately reflects the deviation between nominal and actual energy production.

4.2. UAV-Based Data Acquisition

In the proposed framework, UAV-based data acquisition constitutes the primary source of observational input for estimating photovoltaic power losses. At each inspection time step t, a UAV equipped with high-resolution RGB and thermal cameras follows a pre-planned flight trajectory covering the photovoltaic farm. Let M denote the number of images collected per time step, resulting in a set of images as defined in Equation (1).

$$I\_t \;=\; \{\, I\_\{t,1\}, I\_\{t,2\}, \dots \,, I\_\{t,M\}\,\} \tag{1}$$

Each image is associated with acquisition metadata, including altitude, viewing angle, timestamp, and geolocation coordinates. This metadata captures spatial, temporal, and angular information required for downstream feature extraction, as formalized in Equation (2).

$$m\_\{t,m\} \;=[\, h\_\{t,m\}, s\_\{t,m\}, lat\_\{t,m\}, lon\_\{t,m\}\,] \tag{2}$$

To ensure full coverage of the photovoltaic farm, the flight path is designed such that each panel is observed from multiple angles and overlapping frames. Let *P* denote the total number of panels, and let $C_{t,p}$ represent the set of images covering panel *p* at time *t*. A minimum overlap constraint is enforced to ensure reliable feature extraction, as expressed in Equation (3).

$$C\_\{t,p\} \;=\; \{\, I\_\{t,m\}\,\},\;\; |C\_\{t,p\}| \;\geq\; O \tag{3}$$

Each image is processed to extract visual features using a visual feature extractor, producing a fixed-dimensional embedding vector (Equation (4)). Thermal imagery is processed similarly using a dedicated thermal feature extractor to obtain thermal embeddings (Equation (5)).

$$v\_\{t,m\} \;=\; f\_vis(\, I\_\{t,m\}, m\_\{t,m\}\,) \;\in\; \mathbb{R}\text{^}\{d\_v\} \tag{4}$$

$$u\_\{t,m\} \;=\; f\_therm(\, I\_\{t,m\}\text{^}\{therm\}\, m\_\{t,m\}\,) \;\in\; \mathbb{R}\text{^}\{d\_u\} \tag{5}$$

The multimodal representation at inspection time *t* is obtained by aggregating visual and thermal embeddings across all images, as described in Equation (6). This representation serves as the input to the energy-aware reasoning module.

$$z\_t = (1/M) \sum\_\{m = 1\}\^M \varphi([v\_\{t,m\} \parallel u\_\{t,m\}]) \quad (6)$$

The total number of images collected during an inspection campaign scales linearly with the number of photovoltaic panels and the number of required viewpoints per panel, as shown in Equation (7).

$$N \approx P \cdot O \quad (7)$$

By automating data acquisition and associating structured metadata with each image, the UAV-based approach enables high-frequency, non-intrusive, and reproducible monitoring of photovoltaic installations, providing a rich and physically grounded input for photovoltaic power loss estimation.

4.3. Visual Representation Learning

The collected UAV images are processed to extract visual representations that capture surface conditions and thermal patterns indicative of photovoltaic performance degradation. These include soiling accumulation, partial shading, panel aging effects, and other visual or thermal anomalies that influence energy yield. Each UAV image is processed by a visual encoder to extract high-level feature representations. The encoder can be implemented using convolutional neural networks or transformer-based architectures, enabling robust extraction of spatial patterns related to photovoltaic panel conditions. The resulting visual features provide a compact representation of the observed surface characteristics.

When multiple images correspond to the same photovoltaic panel or spatial region, their extracted visual features are aggregated to obtain a single representative embedding. This aggregation step improves robustness to viewpoint variations and partial occlusions while preserving dominant degradation patterns. Visual information alone is insufficient to fully characterize photovoltaic power losses, as energy production is also affected by environmental and system-level factors. To account for this, contextual information is explicitly incorporated, including meteorological variables such as irradiance and ambient temperature, photovoltaic system parameters, and UAV acquisition metadata.

To enable joint reasoning across modalities, visual and contextual representations are projected into a shared latent space. This alignment ensures compatibility between heterogeneous data sources and facilitates multimodal fusion. The final multimodal representation combines aligned visual and contextual information into a unified embedding, which serves as the input for subsequent energy-aware reasoning and photovoltaic power loss estimation.

4.4. Energy-Aware Reasoning with Large Language Models

The core of the proposed framework is an energy-aware reasoning module built upon a large language model (LLM) that operates as a structured inference engine rather than as a free-form text generator. Its primary role is to transform heterogeneous visual and contextual observations into an internal representation that explicitly reflects photovoltaic degradation phenomena and their impact on energy production.

The reasoning process takes as input a multimodal representation that combines aggregated UAV-derived visual features with synchronized environmental and operational context. Based on this information, the LLM performs relational reasoning to associate observable surface patterns with underlying degradation mechanisms. These mechanisms include soiling accumulation, partial shading, hotspot formation, and long-term module aging effects, each of which contributes differently to photovoltaic power losses.

Through its internal reasoning capabilities, the LLM infers a latent degradation state that captures the severity and nature of energy-impacting phenomena observed at a given inspection time. This latent representation serves as an intermediate, interpretable abstraction between raw observations and final power loss estimation.

To obtain physically meaningful predictions, the inferred degradation state is mapped to a photovoltaic power loss estimate using a dedicated regression module. This module is designed to respect physical constraints inherent to photovoltaic systems. In particular, the estimated power loss is constrained to be non-negative, bounded by the expected power production capacity of the system, and dynamically adjusted according to environmental and operational conditions such as irradiance, temperature, and system configuration (Table 2).

Table 2. Algorithm 1: energy-aware reasoning with large language mode.

Input: Multimodal representation combining visual and contextual features; trained LLM parameters; regression module parameters; photovoltaic system capacity.

Output: Estimated photovoltaic power loss.

Step 1 Multimodal Encoding: The multimodal representation, which integrates UAV visual observations and contextual information, is provided as input to the reasoning module.

Step 2 Energy-aware Reasoning: The large language model analyses the multimodal input to infer latent degradation states by learning relationships between observable surface patterns and their potential impact on energy production.

Step 3 Power Loss Estimation: The inferred degradation state is transformed into a raw power loss estimate using a regression module.

Step 4 Physical Constraint Enforcement: The raw estimate is adjusted to ensure physical plausibility by enforcing non-negativity and upper bounds based on the photovoltaic system's expected power capacity, as well as sensitivity to environmental conditions.

Step 5 Output: The final, physically constrained photovoltaic power loss estimate is produced.

In addition to producing a global power loss estimate, the reasoning module can optionally decompose the prediction into mechanism-specific contributions. These contributions reflect the estimated energy impact of individual degradation sources, such as soiling, shading, or hotspot formation, thereby enhancing interpretability and diagnostic capability. By combining multimodal inputs with structured reasoning and physics-aware constraints, the proposed LLM-based module delivers context-aware, interpretable, and physically grounded photovoltaic power loss estimates that are suitable for large-scale and autonomous photovoltaic monitoring.

4.5. Training Objective

The proposed framework is trained in a supervised manner using SCADA-derived power loss measurements. The primary goal is to minimize the discrepancy between predicted and actual power losses over the inspection period. The model is optimized to produce predictions that are as close as possible to measured losses for each inspection time step. To improve temporal consistency and prevent overfitting, additional regularization strategies are employed. For example, a smoothness penalty discourages abrupt changes between consecutive predictions, ensuring that estimated power losses evolve consistently over time (Table 3).

Table 3. Training procedure of UAV–LLM based photovoltaic power loss estimation.

Input: Sequences of UAV images, contextual data, and ground-truth power loss measurements; visual encoder; LLM reasoning module; learning rate; number of training epochs.
Output: Trained model parameters for visual encoders, projection layers, LLM reasoning, and regression head.
Step 1 Initialization: Initialize the visual encoder, projection layers, LLM reasoning module, and regression head. Initialize the optimizer (e.g., Adam).
Step 2 Training Loop: For each training epoch: Sample a mini-batch of inspection times. *Visual Feature Extraction:* Extract embeddings from UAV images using the visual encoder. Aggregate embeddings to form panel- or region-level visual features. *Multimodal Representation Construction:* Project visual and contextual features into a common latent space. Concatenate them to form the multimodal representation. *Energy-Aware Reasoning:* Infer latent degradation representation using the LLM reasoning module. *Power Loss Prediction:* Estimate the predicted power loss using the regression head. *Loss Computation:* Compute the regression loss as the discrepancy between predicted and ground-truth power losses, including smoothness regularization. *Parameter Update:* Update all trainable parameters via backpropagation.
Step 3 Output: Return the trained model parameters after all epochs.

The total training objective combines the prediction error and the regularization terms. By doing so, the framework learns to accurately estimate photovoltaic power losses while maintaining temporal coherence and physical plausibility. During inference, the trained model processes UAV observations and contextual information to produce power loss estimates at the panel level. These panel-level predictions can then be aggregated to string-level or farm-level estimates, enabling scalable monitoring of large photovoltaic installations and direct comparison with measured total energy production. The entire framework, including visual encoders, projection layers, and the LLM reasoning module, is trained end-to-end using standard gradient-based optimization techniques.

5. Experiments and Results

This section presents a comprehensive experimental evaluation of the proposed LLM-driven aerial monitoring framework for photovoltaic performance assessment. The objective of these experiments is to assess the effectiveness, robustness, and practical relevance of the proposed framework under realistic operational conditions. In particular, we investigate how LLM-enhanced SCADA monitoring and UAV imagery improves power loss estimation and operational decision-making.

5.1. Research Questions

To systematically evaluate the proposed approach, the following research questions are addressed:

- RQ1: How does the proposed method perform compared to strong baselines?
- RQ2: How does performance vary under different environmental and operational conditions?
- RQ3: Which components (LLM, UAV, SCADA) truly contribute to performance gains?
- RQ4: What is the robustness of the proposed method under partial or heterogeneous observations?

5.2. Datasets

Experiments are conducted using a combination of:

- Industrial SCADA datasets from PV plants with multiyear operational measurements.
- UAV imagery datasets capturing PV panel conditions under varying irradiance, shading, and environmental conditions.

5.3. Evaluation Protocol

The datasets are split chronologically into training (80%), validation (10%), and test (10%) subsets. Metrics are computed on unseen test data only. Power loss predictions are compared against measured SCADA losses. We adopt three standard metrics:

- MAE (Mean Absolute Error)
- RMSE (Root Mean Squared Error)
- Pearson correlation between predicted and actual PV

power output.

All experiments are repeated five times with different random seeds. The results report the mean performance across runs.

The results (Table 4 and Fig. 1) show a consistent reduction in power loss estimation error, achieving:

- MAE improvement: 18.7% over the strongest baseline.
- RMSE improvement: 15.2% over the strongest baseline.
- Correlation: higher alignment with actual power production across varying irradiance conditions.

Additional analyses confirm the robustness of the proposed method under partial observations and heterogeneous environmental settings, highlighting the potential of LLM-driven aerial monitoring as a practical and energy-aware solution for PV performance assessment and operational decision-making.

5.4. Ablation Study

To assess the contribution of each component:

- Full model: LLM + UAV imagery + SCADA integration
- No LLM: only UAV + SCADA
- No UAV: only SCADA + LLM
- SCADA only: baseline regression

As seen in Fig. 2, results confirm that the LLM component contributes the most to predictive accuracy, while UAV imagery improves local anomaly detection and robustness.

Table 4. Baselines.

Method	Year	Description
SCADA-only regression	2011	Traditional SCADA data-based estimation
UAV-only CNN	2018	Visual inspection for PV defect detection
SCADA+ML	2024	Classical ML regression using SCADA signals

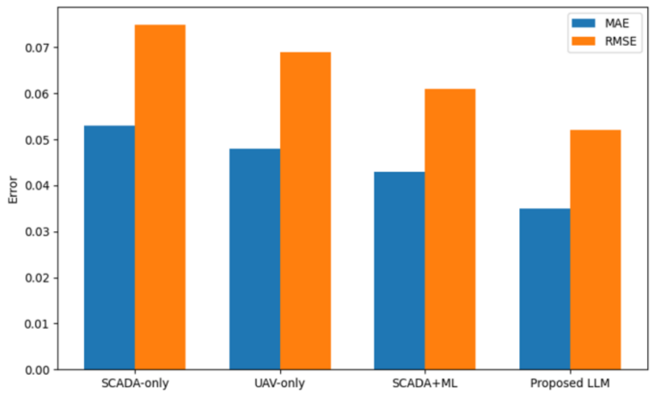

Fig. 1. Performance comparison.

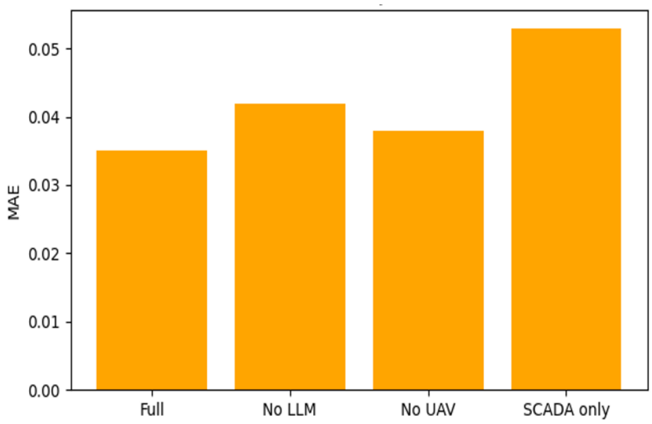

Fig. 2. Ablation study results.

6. Conclusion

In this study, we proposed a novel LLM-driven aerial monitoring framework for photovoltaic (PV) systems that integrates UAV imagery and SCADA sensor data to accurately estimate power losses and support operational decision-making. Our approach leverages large language models to fuse heterogeneous information, enabling robust predictions under varying environmental conditions and partial observations. Comprehensive experiments demonstrate that the proposed method consistently outperforms traditional SCADA-only, UAV-only, and classical machine learning baselines. Specifically, the framework achieved a relative improvement of 18.7% in MAE and 15.2% in RMSE, along with higher correlation with actual power production across diverse irradiance scenarios.

Ablation studies further confirm the critical contributions of each component, highlighting the essential role of LLMs in capturing complex relationships between sensor measurements and panel performance. Beyond predictive accuracy, our method provides interpretable insights into panel-level contributions to power loss, supporting targeted maintenance and operational efficiency. The integration of UAV imagery enhances spatial awareness and complements SCADA signals, while the LLM-based fusion enables intelligent reasoning from heterogeneous data sources.

Overall, the proposed framework illustrates the potential of AI-driven, energy-aware monitoring solutions for PV systems, offering a practical and scalable approach for improving renewable energy management. Future work could explore uncertainty-aware prediction, real-time UAV deployment, and transferability across different PV installations to further enhance robustness and generalization.

Author Contributions

H.M. conceptualized the study, developed the methodology, and performed the formal analysis; A.B. contributed to data curation, software implementation, and experimental validation; T.B. supervised the research, reviewed and edited the manuscript. All authors reviewed and approved the final version of the manuscript.

Acknowledgements

Not applicable.

Conflict of Interest

The authors declare no conflict of interest.

References

[1] G. K. Singh, "Solar power generation by PV (photovoltaic) technology: A review," *Energy*, vol. 53, pp. 1–13, 2013, doi: 10.1016/j.energy.2013.02.057.

[2] U. Agarwal, N. Jain, S. N. Singh, and M. Kumawat, "Solar photovoltaic (PV) generation," in *Fundamentals and Innovations in Solar Energy*, S. N. Singh, P. Tiwari, and S. Tiwari, Eds. Singapore: Springer, 2021.

[3] R. Parthiban and P. Ponnambalam, "An enhancement of the solar panel efficiency: A comprehensive review," *Frontiers in Energy Research*, vol. 10, p. 937155, 2022, doi: 10.3389/fenrg.2022.937155.

[4] S. Fakouriyan, Y. Saboohi, and A. Fathi, "Experimental analysis of a cooling system effect on photovoltaic panels' efficiency and its preheating water production," *Renewable Energy*, vol. 134, pp. 1362–1368, 2019, doi: 10.1016/j.renene.2018.09.054.

[5] M. K. da Silva, M. S. Gul, and H. Chaudhry, "Review on the sources of power loss in monofacial and bifacial photovoltaic technologies," *Energies*, vol. 14, no. 23, p. 7935, 2021, doi: 10.3390/en14237935.

[6] W. R. Anis and M. A. S. Nour, "Energy losses in photovoltaic systems," *Energy Conversion and Management*, vol. 36, no. 11, pp. 1107–1113, 1995.

[7] G. Petrone, C. A. Ramos-Paja, and G. Spagnuolo, *Photovoltaic Sources Modeling*. Hoboken, NJ, USA: John Wiley & Sons, 2017.

[8] G. N. Korres and N. M. Manousakis, "State estimation and bad data processing for systems including PMU and SCADA measurements," *Electric Power Systems Research*, vol. 81, no. 7, pp. 1514–1524, 2011, doi: 10.1016/j.epsr.2011.03.013.

[9] A. Snytko, G. Jiménez-Castillo, F. J. Muñoz-Rodríguez, and C. Rus-Casas, "Fault diagnosis for photovoltaic systems: A validated industrial SCADA framework," *Applied Sciences*, vol. 15, no. 23, p. 12656, 2025.

[10] S. Ferlito, S. Ippolito, C. Santagata, P. Schiattarella, and G. Di Francia, "A study on an IoT-based SCADA system for photovoltaic utility plants," *Electronics*, vol. 13, no. 11, p. 2065, 2024, doi: 10.3390/electronics13112065.

[11] B. Liu and X. Huo, "Prediction of photovoltaic power generation and analyzing of carbon emission reduction capacity in China," *Renewable Energy*, vol. 222, p. 119967, 2024.

[12] C. L. Lin, Y. C. Hwang, and H. C. Lin, "Study of the efficiency improvement in power generation from photovoltaic," *2013 IEEE 17th International Conference on Computer Supported Cooperative Work in Design (CSCWD)*, pp. 509–514, 2013.

[13] M. Aghaei, M. Kolahi, A. Nedaei, N. S. Venkatesh, S. M. Esmailifar, A. M. Moradi Sizkouhi, and R. Rüther, "Autonomous intelligent monitoring of photovoltaic systems: An in-depth multidisciplinary review," *Progress in Photovoltaics: Research and Applications*, vol. 33, no. 3, pp. 381–409, 2025.

[14] P. Addabbo, A. Angrisano, M. L. Bernardi, G. Gagliarde, A. Mennella, M. Nisi, and S. L. Ullo, "UAV system for photovoltaic plant inspection," *IEEE Aerospace and Electronic Systems Magazine*, vol. 33, no. 8, pp. 58–67, 2018.

[15] B. Yin, J. Xie, Y. Qin, Z. Ding, Z. Feng, X. Li, and W. Lin, "Heterogeneous knowledge fusion: A novel approach for personalized recommendation via LLM," *17th ACM Conference on Recommender Systems*, pp. 599–601, 2023.

[16] F. Wang, G. Zhu, C. Yuan, and Y. Huang, "LLM-enhanced cascaded multi-level learning on temporal heterogeneous graphs," *47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 512–521, 2024.

[17] L. Shen, Z. Li, and T. Ma, "Analysis of the power loss and quantification of the energy distribution in PV module," *Applied Energy*, vol. 260, p. 114333, 2020.

[18] Y. Da, Y. Xuan, and Q. Li, "Quantifying energy losses in planar perovskite solar cells," *Solar Energy Materials and Solar Cells*, vol. 174, pp. 206–213, 2018.

[19] Z. Ling, J. Wu, J. P. Jurado, C. E. Petoukhoff, S. Y. Jeong, D. Naphade, M. Babics, X. Chang, H. Faber, S. Doukas, E. Lidorikis, M. I. Nugraha, M. He, M. Alqurashi, Y. Lin, X. Sun, H. Hu, H. Y. Woo, S. De Wolf, L. Tsetseris, F. Laquai, D. Yu, E. Wang, and T. D. Anthopoulos, "20.5% efficient ternary organic photovoltaics using an asymmetric small-molecular acceptor to manipulate intermolecular packing and reduce energy losses," *Materials Science and Engineering: R: Reports*, vol. 163, p. 100922, 2025, doi: 10.1016/j.mser.2024.100922.

[20] L. Zhang, Z. Zhuo, X. Ma, H. Tian, X. Zhao, Y. Xie, K. Yang, B. H. Lee, X. Zhu, H. Y. Woo, C. Yang, X. Nie, and F. Zhang, "19.6% efficiency of layer-by-layer organic photovoltaics with decreased energy loss via incorporating TADF materials with intrinsic reverse intersystem crossing," *Energy & Environmental Science*, vol. 18, no. 20, pp. 9171–9182, 2025, doi: 10.1039/D5EE03118D.

[21] X. Wang, X. Shen, J. Wang, F. Bi, H. Jiang, H. Lu, C. Sun, C. Yang, Y. Li, and X. Bao, "Isomerization of peripheral functional groups refines aggregation and non-radiative energy loss for efficient organic photovoltaics,"

Energy & Environmental Science, vol. 18, pp. 5586–5598, 2025, doi: 10.1039/D5EE00455A.

[22] Y.-B. Wang, C.-L. Tsai, Y.-J. Xue, B.-H. Jiang, H.-C. Lu, J.-C. Hong, Y.-C. Huang, K.-H. Huang, S.-Y. Chien, C.-P. Chen, and Y.-J. Cheng, "Fluorinated and methylated ortho-benzodipyrrole-based acceptors suppressing charge recombination and minimizing energy loss in organic photovoltaics," *Chemical Science*, vol. 16, no. 7, pp. 3259–3274, 2025, doi: 10.1039/D4SC07146H.

[23] Á. Fernández Solas, N. Riedel-Lyngskær, N. Hanrieder, F. Norde Santos, S. Wilbert, H. Nygard Riise, J. Polo, E. F. Fernández, F. Almonacid, D. L. Talavera, and L. Micheli, "Photovoltaic soiling loss in Europe: Geographical distribution and cleaning recommendations," *Renewable Energy*, vol. 239, p. 122086, 2025, doi: 10.1016/j.renene.2024.122086.

[24] A. B. Djilali, E. Bounadja, A. Yahdou, H. Benbouhenni, Z. M. S. Elbarbary, I. Colak, and S. F. Al-Gahtani, "Enhanced variable step sizes perturb and observe MPPT control to reduce energy loss in photovoltaic systems," *Scientific Reports*, vol. 15, no. 1, p. 11700, 2025, doi: 10.1038/s41598-025-95309-y.

[25] M. Guo and Y. Weng, "Solar photovoltaic assessment with large language model," *Applied Energy*, vol. 402, p. 126835, 2025, doi: 10.1016/j.apenergy.2025.126835.

[26] H. Lin and M. Yu, "A novel distributed PV power forecasting approach based on Time-LLM," *2025 IEEE 4th International Conference on Industrial Electronics for Sustainable Energy Systems (IESES)*, 2025, doi: 10.1109/IESES66335.2025.11359876.