

UO‘K: 004.932:621.383.51

<https://doi.org/10.70769/2181-4732.ITJ.2026-3.12>

FOTOVOLTAIK PANELLARDA IFLOSLANISH VA NUQSONLARNI CHUQUR O‘RGANISH ASOSIDA AVTOMATIK ANIQLASH

Bobokulov Shaxzod Ruziboy o‘g‘li, assistant, E-mail: bobokulovshakhzod200@gmail.com,
ORCID: 0000-0002-2180-2278, Science ID: MQD-0225-0018

Qarshi davlat universiteti, Qarshi, O‘zbekiston

Annotatsiya. Fotovoltaik (FV) panellar yuzasining changlanishi hamda mexanik va elektr nuqsonlari quyosh elektr stansiyalarida energiya ishlab chiqarishning pasayishiga olib keluvchi asosiy omillardan biri hisoblanadi. Bu muammo Markaziy Osiyo kabi qurg‘oqchil va changli mintaqalarda ayniqsa keskin namoyon bo‘ladi. Yirik FV massivlarni qo‘lda ko‘zdan kechirish ko‘p mehnat talab qiladi, uzoq vaqt oladi va xatoliklarga moyil. Ushbu maqolada FV panellar yuzasi holatini oltita sinfga — toza, changli, qush axlati bilan ifloslangan, elektr shikastlangan, mexanik shikastlangan va qor bilan qoplangan — avtomatik tasniflovchi yengil hamda natijalarini izohlash mumkin bo‘lgan chuqur o‘rganish tizimi taklif etilgan. ImageNet ma‘lumotlar to‘plamida oldindan o‘qitilgan MobileNetV3-Large neyron tarmog‘i asosida transfer o‘qitish usuli qo‘llanilgan. Sinflar o‘rtasidagi nomutanosiblikni kamaytirish uchun sinflar bo‘yicha vaznlangan yo‘qotish funksiyasi va ma‘lumotlarni kengaytirish usullaridan foydalanilgan. Shikastlangan sohalarni aniqlash uchun Grad-CAM usuli integratsiya qilingan. Tizim aniqlangan har bir holat uchun ilmiy adabiyotlarga asoslangan taxminiy quvvat yo‘qotish oralig‘ini belgilaydi va natijalarni o‘zbek tilidagi FastAPI veb-ilovasi orqali foydalanuvchiga taqdim etadi. 3065 ta tasvirdan iborat ochiq ma‘lumotlar to‘plamida model 78,81% validatsiya aniqligiga, 0,86 makro-o‘rtacha to‘liqlikka (recall) va 0,79 vaznlangan F1 ko‘rsatkichiga erishgan. Barcha hisoblashlar grafik tezlatkichsiz, faqat markaziy protsessorida bajarilgan. Kam uchraydigan, biroq ekspluatatsiya nuqtayi nazaridan muhim sinflar — qor bilan qoplangan (0,98), qush axlati bilan ifloslangan (0,90) va elektr shikastlangan (0,90) panellar — yuqori to‘liqlik bilan aniqlangan. Taklif etilgan yondashuv quyosh energetikasi jadal rivojlanayotgan mintaqalarda texnik xizmat ko‘rsatish bo‘yicha qarorlar qabul qilishni qo‘llab-quvvatlovchi arzon va amaliy vosita hisoblanadi.

Kalit so‘zlar: fotovoltaik panel, changlanish, nuqsonlarni tasniflash, chuqur o‘rganish, transfer o‘qitish, MobileNetV3, Grad-CAM, izohlanuvchan sun‘iy intellekt, quvvat yo‘qotilishi.

УДК: 004.932:621.383.51

АВТОМАТИЧЕСКОЕ ОБНАРУЖЕНИЕ ЗАГРЯЗНЕНИЙ И ДЕФЕКТОВ ФОТОЭЛЕКТРИЧЕСКИХ ПАНЕЛЕЙ НА ОСНОВЕ ГЛУБОКОГО ОБУЧЕНИЯ

Бобокулов Шахзод Рузибой угли, ассистент

Каршинский государственный университет, г. Карши, Узбекистан

Аннотация. Загрязнение поверхности фотоэлектрических (ФЭ) панелей пылью, а также механические и электрические дефекты являются одними из основных факторов снижения выработки электроэнергии на солнечных электростанциях. Особенно остро эта проблема проявляется в засушливых и запылённых регионах, таких как Центральная Азия. Ручной осмотр крупных массивов ФЭ панелей требует значительных трудозатрат, занимает много времени и сопряжён с риском ошибок. В статье предложена легковесная и интерпретируемая система глубокого обучения для автоматической классификации состояния поверхности ФЭ панелей по шести классам: чистые, запылённые, загрязнённые птичьим помётом, с электрическими повреждениями, с механическими повреждениями и покрытые снегом. Применён метод трансферного обучения на основе нейронной сети MobileNetV3-Large, предварительно обученной на наборе данных ImageNet. Для уменьшения дисбаланса классов использованы взвешенная по классам функция потерь и методы аугментации данных. Для локализации повреждённых областей интегрирован метод Grad-

САМ. Система определяет для каждого распознанного состояния ориентировочный диапазон потерь мощности на основе научной литературы и предоставляет результаты пользователю через веб-приложение FastAPI с узбекязычным интерфейсом. На открытом наборе данных из 3065 изображений модель достигла точности на валидационной выборке 78,81%, макроусреднённой полноты (recall) 0,86 и взвешенного показателя F1, равного 0,79. Все вычисления выполнялись исключительно на центральном процессоре без графического ускорителя. Редкие, но важные с эксплуатационной точки зрения классы — панели, покрытые снегом (0,98), загрязнённые птичьим помётом (0,90) и имеющие электрические повреждения (0,90), — распознавались с высокой полнотой. Предложенный подход представляет собой недорогой и практичный инструмент поддержки принятия решений по техническому обслуживанию в регионах с интенсивным развитием солнечной энергетики.

Ключевые слова: фотоэлектрическая панель, загрязнение пылью, классификация дефектов, глубокое обучение, трансферное обучение, MobileNetV3, Grad-CAM, интерпретируемый искусственный интеллект, потери мощности.

UDC: 004.932:621.383.51

DEEP LEARNING-BASED AUTOMATIC DETECTION OF SOILING AND FAULTS ON PHOTOVOLTAIC PANELS

Shakhzod Bobokulov – Assistant

Karshi State University, Karshi, Uzbekistan

Abstract. Dust accumulation on the surfaces of photovoltaic (PV) panels, as well as mechanical and electrical defects, are among the main factors contributing to reduced energy generation at solar power plants. This problem is particularly pronounced in arid and dusty regions such as Central Asia. Manual inspection of large PV arrays is labour-intensive, time-consuming and prone to errors. This paper proposes a lightweight and explainable deep learning system for automatically classifying PV panel surface conditions into six classes: clean, dusty, contaminated with bird droppings, electrically damaged, mechanically damaged and snow-covered. Transfer learning is applied using the MobileNetV3-Large neural network pretrained on the ImageNet dataset. To mitigate class imbalance, a class-weighted loss function and data augmentation techniques are employed. The Grad-CAM method is integrated to localise damaged regions. For each detected condition, the system estimates a potential power loss range based on scientific literature and presents the results to users through a FastAPI web application with an Uzbek-language interface. On a publicly available dataset containing 3,065 images, the model achieved a validation accuracy of 78.81%, a macro-averaged recall of 0.86 and a weighted F1-score of 0.79. All computations were performed exclusively on a central processing unit without a graphics accelerator. Rare but operationally important classes — snow-covered panels (0.98), panels contaminated with bird droppings (0.90) and electrically damaged panels (0.90) — were identified with high recall. The proposed approach provides a low-cost and practical decision-support tool for maintenance operations in regions experiencing rapid development of solar energy.

Keywords: photovoltaic panel, dust accumulation, defect classification, deep learning, transfer learning, MobileNetV3, Grad-CAM, explainable artificial intelligence, power loss.

Introduction

The globally installed capacity of solar photovoltaic (PV) systems has grown at an unprecedented rate over the past decade [1]. Uzbekistan is becoming one of the most active solar energy markets in Central Asia, and national targets envisage raising utility-scale PV capacity to several gigawatts by 2030. However, the very climatic features that make the region attractive for solar energy — high irradiance, low precipitation and a long dry season — also make PV installations extremely prone to soiling. The literature reports that dust deposition alone reduces PV energy yield by 15–25% in arid climates, and that losses may exceed 50% when panels are left uncleaned for long

periods [2, 3, 4]. At plant scale, soiling translates into the loss of a substantial share of annual revenue, and scheduling cleaning becomes a first-order economic problem rather than a purely technical one [5].

Besides soiling, PV modules are affected by bird droppings that create localised hot spots, physical damage such as glass cracking, electrical faults in the form of burn marks and corrosion, and seasonal snow cover. Conventional monitoring of such conditions relies on periodic manual inspection or string-level electrical measurements; both scale poorly at large plants comprising hundreds of thousands of modules [6, 7]. Computer vision offers an alternative: inexpensive RGB cameras — handheld, fixed, or mounted on an unmanned aerial vehicle (UAV) — can capture panel images and allow them to be analysed automatically [8].

Recent studies have applied convolutional neural networks (CNNs) to soiling localisation [9], [10], the detection of uneven dust layers [11] and object-detection-based soiling recognition [12]. Approaches based on infrared thermography [13, 14] and electroluminescence imaging [15, 16] achieve high accuracy but require specialised and expensive equipment. Nevertheless, three gaps persist in practice. First, most of the proposed models are computationally heavy and unsuitable for edge deployment at remote plants with limited computing infrastructure. Second, most classifiers operate as black boxes and do not indicate where the detected problem is located, which undermines operator trust [17]. Third, classification outputs are rarely converted into operationally meaningful quantities such as expected power loss. In addition, in real image collections clean and dusty panels outnumber the rare physical and electrical damage classes by a wide margin, and this class imbalance biases the classifier towards the majority classes at the expense of precisely the most important ones [18].

The aim of this study is to develop a lightweight, explainable and fully deployed PV condition monitoring system that addresses the gaps outlined above. The scientific novelty of the work consists in: (1) a six-class classifier based on MobileNetV3-Large transfer learning with a class-weighted loss function; (2) an empirical analysis of the recall–precision trade-off induced by class weighting; (3) Grad-CAM integration that highlights the affected region directly on the image; (4) a rule-based module that converts the classification result into a power loss range; (5) a FastAPI web application with an Uzbek-language interface. To the best of our knowledge, this is one of the first PV inspection systems localised into Uzbek.

Materials and Methods

System architecture. The overall structure of the proposed system is shown in Fig. 1. An input RGB image of a PV panel is pre-processed and passed to a transfer-learned CNN backbone; the classifier head returns probabilities over the six condition classes. The Grad-CAM module produces a heat map highlighting the image region most responsible for the decision, while a rule-based post-processing stage maps the predicted class to a power loss range and an operational conclusion (clean, or soiled/faulty with a cleaning or repair recommendation). All components are served through a FastAPI server and presented in a single-page responsive web interface.




Fig. 1. Overall architecture of the proposed PV soiling and fault detection system.

Source: developed by the author.

Dataset. The experiments were carried out on the publicly available Kaggle dataset «Solar Panel Images: Clean and Faulty» [19]. The dataset consists of RGB photographs of PV panels grouped into six condition classes. After preparation it contained 3065 images in total; the class distribution is shown in Fig. 2 and summarised in Table 1 together with the power loss range assigned to each class on the basis of the soiling and fault literature [3, 4, 6]. The dataset is strongly imbalanced: the two

majority classes (clean — 48.8%; dusty — 34.9%) account for more than 83% of all images, whereas the operationally most critical fault classes are rare — electrical damage accounts for only 3.3% and physical damage for 2.3%. The images were split into training (80%) and validation (20%) sets using class-stratified sampling; the validation set preserved the class proportions and comprised 613 images.




Fig. 2. Class distribution in the PV panel image dataset.

Source: author's calculation based on the dataset [19].

Table 1. Dataset composition and the power loss ranges assigned to the classes

Class	Number of images	Share, %	Estimated power loss, %	Typical cause
Clean	1495	48.8	0	Normal operating condition
Dusty	1070	34.9	15–25	Airborne dust, arid climate
Bird droppings	205	6.7	10–30	Bird droppings, local shading
Snow-covered	125	4.1	40–90	Seasonal snow cover
Electrical damage	100	3.3	30–60	Hot spots, burn marks, corrosion
Physical damage	70	2.3	20–50	Cracked glass, broken cells
Total	3065	100.0	–	–

Source: compiled by the author based on [3], [4], [6], [19]

Pre-processing and data augmentation. All images were resized to 224×224 pixels and normalised using the ImageNet [20] channel statistics. To improve generalisation, random cropping (scale 0.7–1.0), horizontal flipping, rotation of up to $\pm 15^\circ$ and random variation of brightness, contrast and saturation were applied during training. These transformations imitate the variations in camera viewpoint, panel orientation and illumination encountered in the field. Only deterministic resizing and centre cropping were applied to the validation images.

Network architecture and transfer learning. MobileNetV3-Large [21], pre-trained on ImageNet [20], was selected as the backbone. The original classification head was replaced by a dropout layer (rate 0.2) and a fully connected layer with six outputs, and the whole network was retrained end-to-end, since the texture statistics of PV panel images differ considerably from the natural ImageNet classes. The architecture contains approximately 5.4 million parameters and is therefore light enough for both training and inference on an ordinary CPU without a graphics accelerator; this requirement is one of the key constraints of the project, because such infrastructure is often unavailable at remote PV plants in Central Asia. The implementation additionally supports EfficientNet-B0 [22] as an interchangeable backbone.

Training procedure. The model was trained entirely on the CPU with the Adam optimiser, an initial learning rate of 10^{-3} and a batch size of 32, for at most 25 epochs. To counter the severe imbalance described above, the cross-entropy loss was weighted inversely proportionally to the class

frequencies [18, 23], which deliberately increases the penalty for misclassifying the rare fault classes. The learning rate was reduced by a factor of 0.1 whenever the validation loss did not improve for two consecutive epochs, and early stopping with a patience of five epochs prevented overfitting; the checkpoint with the lowest validation loss was retained. All experiments were implemented in PyTorch.

Interpretation via Grad-CAM. For visual explanation, Grad-CAM [24] is computed over the last convolutional block of the backbone. The gradients of the predicted class score with respect to the feature maps are globally averaged to yield importance weights; the resulting coarse localisation map is upsampled to the input image size and overlaid as a heat map. This allows the operator to verify at a glance that, for example, a «dusty» decision was driven by the actually soiled region rather than by background artefacts, which increases trust in the decisions of the system [17].

Power loss estimation and deployment. Each predicted class is mapped to the power loss range given in Table 1: 15–25% for dust [3, 4], 10–30% for bird droppings depending on the degree of coverage, 30–60% for electrical damage, 20–50% for physical damage and 40–90% for snow cover. The severity level (low, medium or high) is derived from the classifier confidence. The complete pipeline is deployed as a FastAPI web service with endpoints for batch prediction and Grad-CAM generation; the frontend supports the simultaneous upload of several images, per-class confidence charts, session history and both dark and light themes. The system follows the principle of graceful degradation: if the model file is missing, the interface displays training instructions instead of raising an error, while files larger than 15 MB or corrupted files are rejected individually without aborting the whole batch.

Results and Discussion

Training dynamics. Fig. 3 presents the loss and accuracy curves. The best validation accuracy reached 78.81% before the early stopping criterion was triggered. The persistent gap between the training and validation curves reflects the difficulty of the task under severe class imbalance and non-uniform image quality: the dataset mixes close-up photographs of individual modules with wide array views taken under highly variable illumination. Importantly, all training was performed on the CPU alone, which confirms the practical applicability of the lightweight backbone in infrastructure-constrained settings.




Fig. 3. Training and validation (a) loss and (b) accuracy of the MobileNetV3-Large model.
Source: author's calculations.

Classification performance. Table 2 reports the per-class precision, recall and F1-score on the validation set of 613 images; Fig. 4 shows the corresponding confusion matrix and Fig. 5 the per-class metrics. The overall validation accuracy is 78.81% and the weighted F1-score is 0.79; the macro-averaged recall of 0.86 is considerably higher than the macro-averaged precision of 0.74, which is the expected outcome of class-weighted training on an imbalanced dataset [18, 23].

Table 2. Per-class validation metrics of the MobileNetV3-Large model

Class	Precision	Recall	F1-score	Support
Bird droppings	0.43	0.90	0.58	41
Clean	0.90	0.81	0.85	299
Dusty	0.77	0.67	0.71	214
Electrical damage	0.72	0.90	0.80	20
Physical damage	0.80	0.86	0.83	14
Snow-covered	0.81	0.98	0.89	25
Macro average	0.74	0.86	0.78	613
Weighted average	0.81	0.78	0.79	613

Source: author's calculations.




Fig. 4. Confusion matrix on the validation set (n = 613).

Source: author's calculations.




Fig. 5. Per-class precision, recall and F1-score on the validation set.

Source: author's calculations.

Three important conclusions follow from these results. First, the rare but operationally decisive classes are detected with very high recall: not a single snow-covered panel was missed (recall 0.98), while the bird droppings and electrical damage classes reached a recall of 0.90 despite accounting for only 6.7% and 3.3% of the data respectively. In a maintenance context this behaviour is exactly the desirable one, since a missed fault leaves a degraded panel in operation, whereas a false alarm merely calls for an additional inspection.

Second, the cost of this sensitivity is concentrated in the precision of the bird droppings class (0.43): the confusion matrix shows that a noticeable share of clean and dusty panels is labelled as bird droppings. At a resolution of 224×224 pixels, light reflections on clean glass or accumulated dust patches visually resemble droppings, and the class-weighted loss further reinforces the tendency of the classifier to choose this small class. Third, the dominant error mode among the majority classes is the mutual confusion of clean and dusty panels: 38 clean panels were predicted as dusty and 32 dusty panels as clean. This boundary is objectively ambiguous, because light and uniform soiling produces only subtle changes in surface appearance and these changes depend strongly on illumination [10]; consequently the recall of the dusty class (0.67) is the lowest among all classes.

Error analysis and practical significance. From an operational point of view the structure of the errors is favourable. When the six classes are reduced to a binary decision — clean or any anomaly — most of the remaining errors either stay within the anomaly group (for example, bird droppings instead of dust, or electrical instead of physical damage) or belong to the clean/dusty boundary, where the difference in power loss is smallest. Severe fault categories are practically never missed: in the validation set not a single electrically or physically damaged panel was classified as clean. A qualitative review of the Grad-CAM heat maps showed that the network attends to the soiled or damaged region itself rather than to the panel frame or the background; this supports the reliability of the classification and enables the operator to dismiss obvious false alarms, such as a heat map placed over a light reflection. The CPU inference time of the deployed model is a few tens of milliseconds per image, so the web application comfortably supports interactive batch analysis and the real-time processing of UAV image streams even without graphics hardware.

Interpretation of power loss. Fig. 6 shows the power loss ranges corresponding to each condition class. Although these estimates are derived from the literature [3, 4, 5] rather than from electrical measurements on the photographed panels, they turn an abstract class label into an actionable maintenance signal: a plant operator can prioritise work by placing panels labelled as electrically damaged or heavily snow-covered, whose potential loss exceeds 40%, ahead of lightly soiled ones. Combined with the high recall achieved precisely on the high-loss classes, the system acts as an effective triage tool.




Fig. 6. Estimated power loss ranges by predicted condition class.

Source: compiled by the author based on [3], [4], [5].

Limitations. Five limitations of the study should be noted. First, the low precision of the bird droppings class (0.43), although operationally tolerable, generates false alarms that increase the inspection workload; a higher input resolution, focal loss [25], class-balanced loss [23] or a two-stage verification model are promising remedies. Second, the clean/dusty boundary remains ambiguous at the image level; combining images with irradiance-normalised power measurements would help to resolve light soiling. Third, the power loss estimates are class-level literature values rather than per-panel measurements. Fourth, all images contain a single dominant panel, whereas dense aerial photographs of a full array would require a detection stage prior to classification. Fifth, the results are reported for a single stratified split of one public dataset; repeated experiments with different random

seeds, k-fold cross-validation and a systematic ablation against unweighted training and alternative backbones would strengthen the statistical conclusions.

Conclusion

This paper has proposed a lightweight and explainable deep learning system for the automatic detection of soiling and faults on photovoltaic panels. Using MobileNetV3-Large transfer learning, class-weighted training and data augmentation, the system achieved a validation accuracy of 78.81% and a macro-averaged recall of 0.86 on a strongly imbalanced dataset of 3065 images; all computation and deployment were performed on the CPU alone. Most importantly for maintenance tasks, the rare high-severity classes are detected with a recall of 0.86–0.98 and no damaged panel is classified as clean. Grad-CAM heat maps localise the affected region, while the per-class power loss estimates translate the predictions into quantities that are meaningful for technical service. The complete system is deployed as an Uzbek-language web application intended for the operators of PV plants in Uzbekistan, where solar capacity is expanding rapidly and dust-related losses are a first-order operational problem.

Further research will continue in five directions: (1) improving the precision of the bird droppings class and the separation of the clean and dusty classes through higher-resolution input images, focal loss and hard example mining; (2) collecting and annotating a field dataset from operating PV plants in Uzbekistan and re-tuning the model to local dust composition and illumination conditions; (3) linking image-based predictions to inverter-level electrical measurements in order to calibrate a per-panel power loss regression; (4) extending the pipeline with a YOLO-type detector stage for panel localisation, enabling UAV-based inspection of a full array; (5) integrating the module into a unified intelligent monitoring platform for Central Asian renewable energy infrastructure.

References

- [1] IRENA, “Renewable Capacity Statistics 2025,” Abu Dhabi, 2025.
- [2] T. Sarver, A. Al-Qaraghuli, and L. L. Kazmerski, “A comprehensive review of the impact of dust on the use of solar energy: History, investigations, results, literature, and mitigation approaches,” *Renewable and Sustainable Energy Reviews*, vol. 22, pp. 698–733, 2013, doi: 10.1016/j.rser.2012.12.065.
- [3] A. Sayyah, M. N. Horenstein, and M. K. Mazumder, “Energy yield loss caused by dust deposition on photovoltaic panels,” *Solar Energy*, vol. 107, pp. 576–604, 2014, doi: 10.1016/j.solener.2014.05.030.
- [4] M. R. Maghami, H. Hizam, C. Gomes, M. A. Radzi, M. I. Rezadad, and S. Hajighorbani, “Power loss due to soiling on solar panel: A review,” *Renewable and Sustainable Energy Reviews*, vol. 59, pp. 1307–1316, 2016, doi: 10.1016/j.rser.2016.01.044.
- [5] K. Ilse *et al.*, “Techno-economic assessment of soiling losses and mitigation strategies for solar power generation,” *Joule*, vol. 3, no. 10, pp. 2303–2321, 2019, doi: 10.1016/j.joule.2019.08.019.
- [6] D. S. Pillai and N. Rajasekar, “A comprehensive review on protection challenges and fault diagnosis in PV systems,” *Renewable and Sustainable Energy Reviews*, vol. 91, pp. 18–40, 2018, doi: 10.1016/j.rser.2018.03.082.
- [7] A. Mellit and S. A. Kalogirou, “Artificial intelligence and internet of things to improve efficacy of diagnosis and remote sensing of solar photovoltaic systems,” *Renewable and Sustainable Energy Reviews*, vol. 143, p. 110889, 2021, doi: 10.1016/j.rser.2021.110889.
- [8] Y. Zefri, A. ElKettani, I. Sebari, and S. Ait Lamallam, “Thermal infrared and visual inspection of photovoltaic installations by UAV photogrammetry: Application case Morocco,” *Remote Sens. (Basel)*, vol. 10, no. 12, p. 1874, 2018, doi: 10.3390/rs10121874.
- [9] S. Mehta, A. P. Azad, S. A. Chemmengath, V. Raykar, and S. Kalyanaraman, “DeepSolarEye: Power loss prediction and weakly supervised soiling localization via fully convolutional networks for

- solar panels,” in *Proceedings of the IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2018, pp. 333–342. doi: 10.1109/WACV.2018.00043.
- [10] M. Yang, J. Ji, and B. Guo, “Soiling quantification using an image-based method: Effects of imaging conditions,” *IEEE J. Photovolt.*, vol. 10, no. 6, pp. 1780–1787, 2020, doi: 10.1109/JPHOTOV.2020.3018493.
- [11] S. Fan, Y. Wang, S. Cao, B. Zhao, T. Sun, and P. Liu, “A deep residual neural network identification method for uneven dust accumulation on photovoltaic (PV) panels,” *Energy*, vol. 239, p. 122302, 2022, doi: 10.1016/j.energy.2021.122302.
- [12] A. Et-Taleby, Y. Chaibi, M. Boussetta, and M. Benslimane, “Application of deep learning based detector YOLOv5 for soiling recognition in photovoltaic modules,” in *Proceedings of the 3rd International Conference on Innovative Research in Applied Science, Engineering and Technology (IRASET)*, Mohammedia, Morocco, 2023, pp. 1–6. doi: 10.1109/IRASET57153.2023.10152990.
- [13] M. W. Akram, G. Li, Y. Jin, X. Chen, C. Zhu, and A. Ahmad, “Automatic detection of photovoltaic module defects in infrared images with isolated and develop-model transfer deep learning,” *Solar Energy*, vol. 198, pp. 175–186, 2020, doi: 10.1016/j.solener.2020.01.055.
- [14] A. H. Herraiz, A. P. Marugán, and F. P. G. Márquez, “Photovoltaic plant condition monitoring using thermal images analysis by convolutional neural network-based structure,” *Renew. Energy*, vol. 153, pp. 334–348, 2020, doi: 10.1016/j.renene.2020.01.148.
- [15] S. Deitsch *et al.*, “Automatic classification of defective photovoltaic module cells in electroluminescence images,” *Solar Energy*, vol. 185, pp. 455–468, 2019, doi: 10.1016/j.solener.2019.02.067.
- [16] J. Wang *et al.*, “Deep-learning-based automatic detection of photovoltaic cell defects in electroluminescence images,” *Sensors*, vol. 23, no. 1, p. 297, 2023, doi: 10.3390/s23010297.
- [17] N. Ledmaoui, A. El Maghraoui, M. El Aroussi, and R. Saadane, “Enhanced fault detection in photovoltaic panels using CNN-based classification with PyQt5 implementation,” *Sensors*, vol. 24, no. 22, p. 7407, 2024, doi: 10.3390/s24227407.
- [18] M. Buda, A. Maki, and M. A. Mazurowski, “A systematic study of the class imbalance problem in convolutional neural networks,” *Neural Networks*, vol. 106, pp. 249–259, 2018, doi: 10.1016/j.neunet.2018.07.011.
- [19] Kaggle, “Solar Panel Images: Clean and Faulty.” [Online]. Available: <https://www.kaggle.com/datasets>
- [20] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, “ImageNet: A large-scale hierarchical image database,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009, pp. 248–255. doi: 10.1109/CVPR.2009.5206848.
- [21] J. Yosinski, J. Clune, Y. Bengio, and H. Lipson, “How transferable are features in deep neural networks?,” in *Advances in Neural Information Processing Systems (NeurIPS)*, 2014, pp. 3320–3328.
- [22] M. Tan and Q. V. Le, “EfficientNet: Rethinking model scaling for convolutional neural networks,” in *Proceedings of the 36th International Conference on Machine Learning (ICML)*, PMLR, 2019, pp. 6105–6114.
- [23] Y. Cui, M. Jia, T.-Y. Lin, Y. Song, and S. Belongie, “Class-balanced loss based on effective number of samples,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019, pp. 9268–9277. doi: 10.1109/CVPR.2019.00949.
- [24] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, “Grad-CAM: Visual explanations from deep networks via gradient-based localization,” in *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2017, pp. 618–626. doi: 10.1109/ICCV.2017.74.
- [25] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár, “Focal loss for dense object detection,” in *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2017, pp. 2980–2988. doi: 10.1109/ICCV.2017.324.