future internet

MDPI

Article

Deep Learning-Based Approach to Automated Monitoring of Defects and Soiling on Solar Panels

Ahmed Hamdi †, Hassan N. Noura *,† and Joseph Azar

FEMTO-ST Institute, University Marie et Louis Pasteur, F-90000 Belfort, France; ahmed.hamdi@femto-st.fr (A.H.); joseph.azar@univ-fcomte.fr (J.A.)

* Correspondence: hassan.noura@univ-fcomte.fr

† These authors contributed equally to this work.

Abstract

The reliable operation of photovoltaic (PV) systems is often compromised by surface soiling and structural damage, which reduce energy efficiency and complicate large-scale monitoring. To address this challenge, we propose a two-tiered image-classification framework that combines Vision Transformer (ViT) models, lightweight convolutional neural networks (CNNs), and knowledge distillation (KD). In Tier 1, a DINOv2 ViT-Base model is fine-tuned to provide robust high-level categorization of solar-panel images into three classes: Normal,Soiled, and Damaged. In Tier 2, two enhanced EfficientNetB0 models are introduced: (i) a KD-based student model distilled from a DINOv2 ViT-S/14 teacher, which improves accuracy from 96.7% to 98.67% for damage classification and from 90.7% to 92.38% for soiling classification, and (ii) an EfficientNetB0 augmented with Multi-Head Self-Attention (MHSA), which achieves 98.73% accuracy for damage and 93.33% accuracy for soiling. These results demonstrate that integrating transformer-based representations with compact CNN architectures yields a scalable and efficient solution for automated monitoring of the condition of PV systems, offering high accuracy and real-time applicability in inspections on solar farms.

Keywords: solar panels; damage detection; soiling detection; deep learning; vision transformers; Dinov2

check for updates

Academic Editor: Manuel José Cabral dos Santos Reis

Received: 27 July 2025
Revised: 17 September 2025
Accepted: 17 September 2025
Published: 23 September 2025

Citation: Hamdi, A.; Noura, H.N.; Azar, J. Deep Learning-Based Approach to Automated Monitoring of Defects and Soiling on Solar Panels. *Future Internet* **2025**, *17*, 433. https://doi.org/10.3390/fi17100433

Copyright: © 2025 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/).

1. Introduction

The integration of solar energy systems has emerged as a critical strategy in the global transition toward sustainable, low-carbon infrastructures [1]. Central to this shift are photovoltaic (PV) panels, whose efficiency can be significantly reduced from their design levels by environmental and physical factors.

To illustrate the operational context of this work, Figure 1 presents the main components and power flow of a typical on-grid solar PV system. Sunlight is captured by the solar array and converted into DC power. This energy is routed through the inverter, which supplies AC power to household and industrial loads. Surplus energy is stored in the battery bank for later use, while additional electricity can be drawn from the distribution grid (DISCOM) when solar generation is insufficient. Conversely, when local production exceeds demand, excess power can be exported to the grid. The inverter plays a central role in coordinating energy flows among the solar array, the battery bank, the grid (as both backup source and output destination), and the electrical loads, ensuring reliable and continuous power availability.

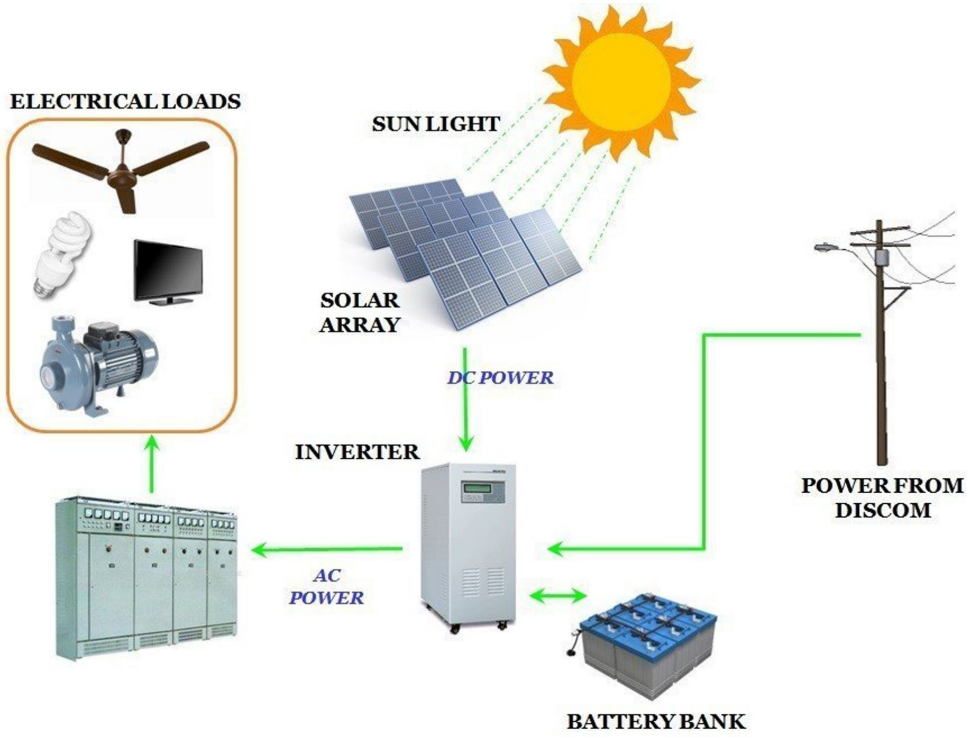

Figure 1. Schematic of an on-grid solar PV system, showing energy flows between the solar array, inverter, battery bank, distribution grid, and electrical loads.

While PV systems offer a renewable alternative to fossil fuels, their efficiency can be significantly reduced by soiling [2]. Dust [3], snow [4], and bird droppings [5] are frequent contaminants that block sunlight and reduce energy conversion, sometimes causing losses ranging from 5% to more than 60%, depending on severity and local conditions [6]. In addition, physical [7] and electrical [8] damage, such as cracks, delamination, hot spots, or diode breakdown, can drastically reduce panel output. Figure 2 shows representative examples of clean panels, soiling types, and damage modalities.

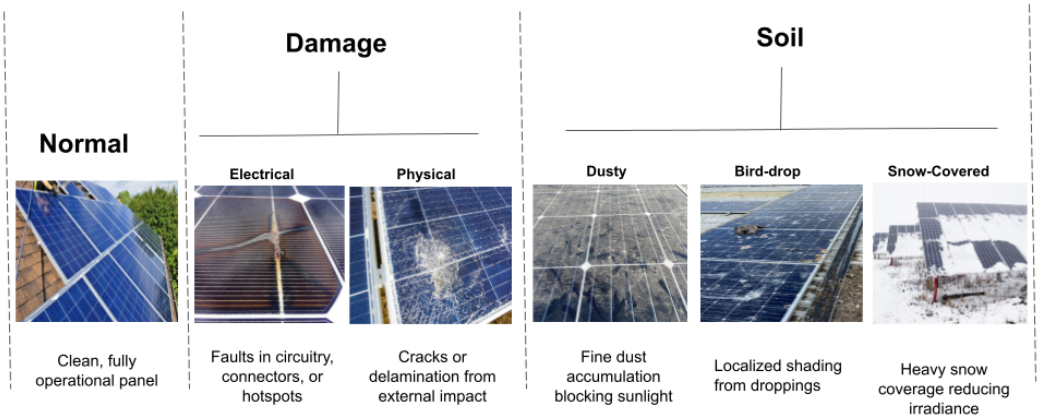

Figure 2. Examples of solar panel surfaces in various conditions, including normal, soiled (dust, bird droppings, snow), and damaged (electrical, physical).

Manual inspection of PV systems across large farms is labor-intensive and inefficient. Drone-based inspection systems have therefore gained popularity, as they can cover vast areas quickly while capturing high-resolution imagery. Figure 3 illustrates a drone operating over a solar array with AI-assisted anomaly detection.



Figure 3. Autonomous drone conducting visual inspection of solar panels using advanced sensing technology. (Source: [9]).

The integration of artificial intelligence (AI) into drone-based inspection workflows improves automation and accuracy. Deep learning models can recognize complex soiling patterns and defect types directly from images, enabling consistent large-scale monitoring. A major challenge entails balancing accuracy with constraints on model size and inference speed: transformer-based models such as Vision Transformers (ViTs) [10] achieve strong performance but are computationally heavy, while lightweight CNNs [11] such as EfficientNet are faster but may miss subtle signs of defects.

- Problem Statement and Motivation.
  Despite significant progress has been made in automated inspection, but existing methods still face important limitations. Many approaches are either computationally heavy transformer-based models that are not well-suited to real-time deployment on drones or lightweight CNNs that fail to capture subtle soiling and signs of damage. Moreover, a unified framework that effectively balances accuracy, efficiency, and scalability for large-scale solar panel inspection is still lacking. These challenges indicate the need for an approach that is both accurate and lightweight, making it suitable for real-world monitoring on solar farms.
- Aim of the Paper. This paper aims to design and evaluate a two-tiered classification framework that combines the strengths of DINOv2-based vision transformers and lightweight EfficientNet extensions to achieve both high accuracy and real-time suitability for the detection of defects and soiling on solar panels.
- Contributions. The main contributions of this work are summarized as follows:
  - We introduce a novel two-tiered classification framework for monitoring the condition of solar panels, where Tier 1 provides high-level categorization (Normal, Soiled Damaged) using DINOv2 ViT-Base, and Tier 2 performs fine-grained classification with lightweight EfficientNet-based models.
  - We enhance EfficientNetB0 with two complementary strategies: (i) Multi-Head Self-Attention (MHSA) for capturing global contextual information and (ii) Knowledge Distillation (KD) from a strong DINOv2 ViT-S teacher to improve feature learning in compact models.
  - We conduct extensive experiments on solar panel datasets under diverse conditions, demonstrating that the proposed models outperform state-of-the-art baselines in terms of accuracy, inference time, and parameter efficiency.

Organization

The remainder of this paper is structured as follows. Section 2 reviews recent developments in monitoring the condition of solar panels using deep learning, focusing on CNN-based methods, Vision Transformers, and hybrid techniques such as MHSA. Section 3 presents the background knowledge essential for understanding our approach, including common solar panel defects, drone-based inspection systems, CNN fundamentals, EfficientNet architecture, the DINOv2 model, and knowledge distillation. Section 4 describes the proposed two-tiered classification framework in detail, outlining the dataset, architectural design, and enhancement strategies that use KD and MHSA. Section 5 provides comprehensive experimental results that were used to evaluate classification performance, confusion matrices, and inference efficiency across various model configurations. Section 6 offers a discussion of key findings, error patterns, and the trade-offs between accuracy and model complexity. Finally, Section 7 concludes the paper and outlines potential future directions, such as integrating object detection, self-supervised learning, and hardware-aware model optimization. The abbreviations used in this paper are summarized in Table 1.

Table 1. Table of abbreviations.

Abbreviation	Full Term
AI	Artificial Intelligence
AC	Alternating Current
B0, B4, B7	EfficientNet Variants (B0 = baseline, B4/B7 = scaled versions)
CNN	Convolutional Neural Network
DC	Direct Current
DINO	Distillation with No Labels
DINOv2	Distillation with No Labels Version 2
EL	Electroluminescence
FC	Fully Connected (Layer)
FLOPs	Floating Point Operations
KD	Knowledge Distillation
KL Divergence	Kullback–Leibler Divergence
MBConv	Mobile Inverted Bottleneck Convolution
MHSA	Multi-Head Self-Attention
NAS	Neural Architecture Search
NLP	Natural Language Processing
PV	Photovoltaic
RGB	Red–Green–Blue (color image format)
SE	Squeeze-and-Excitation
SSL	Self-Supervised Learning
TPR	True Positive Rate
UAV	Unmanned Aerial Vehicle (Drone)
ViT	Vision Transformer
YOLO	You Only Look Once (Object Detection Model)

2. Related Work

This section reviews prior research on solar panel monitoring. It begins with the general problem (Section 2.1), then follows with a discussion of CNN-based methods (Section 2.2), Transformer-based approaches (Section 2.3), and recent work combining MHSA with lightweight CNNs (Section 2.4). Finally, it summarizes comparative studies and identifies remaining gaps (Section 2.5).

2.1. Introduction to the General Problem

Solar PV systems are central to sustainable energy infrastructure, making reliable monitoring essential [12,13]. Traditional inspection approaches [14,15], which primarily

rely on manual assessments, face major scalability and efficiency challenges on large-scale solar farms. These methods are often labor-intensive, time-consuming, prone to subjective inconsistencies, and logistically difficult in hazardous or hard-to-reach areas. The inherent limitations of manual inspection create a bottleneck, delaying the identification and rectification of performance-degrading issues.

As a result, AI-based methods have been proposed to automate monitoring. Still, environmental variability (illumination, weather, thin soiling layers, microcracks) makes robust generalization difficult [16]. The demand for real-time analysis for proactive maintenance [17] further imposes strict efficiency constraints, especially on edge devices [18]. These challenges motivate the development of advanced deep learning architectures combined with optimization techniques.

2.2. CNN-Based Methods

CNNs [19] learn spatial feature hierarchies directly from pixels, which has made them widely used in solar panel monitoring. Architectures such as AlexNet [20], VGG [21], ResNet [22], and EfficientNet [23] have been applied to classify panels as normal, soiled, or damaged [24,25]. CNNs are also used for defect detection with EL imagery [26] and object-detection tasks with YOLO.

Their main limitation is their reliance on local receptive fields, which restricts modeling of long-range dependencies. Even with deeper networks or larger kernels, global context remains difficult to capture. For analysis of solar panels, this can be problematic when large-scale or distributed anomalies must be recognized.

2.3. Transformer-Based Models

Vision Transformers (ViTs) [27], inspired by NLP [28], divide images into patches and use self-attention to capture relationships among all patches. This gives them an inherent global receptive field, often leading to better performance than that obtained with CNNs for tasks requiring a holistic understanding of images.

Self-supervised learning frameworks such as DINO [29] and DINOv2 [30] further strengthen ViTs, producing general-purpose features transferable to many tasks. However, their quadratic complexity in relation to patch number leads to high computational and memory costs, limiting deployment on resource-constrained edge devices used for inspection of solar panels. This motivates research into hybrid and optimized models.

2.4. MHSA in Lightweight CNNs

MHSA [31], originally from Transformers, has been integrated into CNNs [32] to combine local efficiency with global context modeling. When paired with knowledge distillation (KD), MHSA-augmented CNNs are better able to absorb globally-aware knowledge from ViT teachers. These hybrid CNNs achieve a balance between accuracy and efficiency, making them promising for real-time solar inspection on edge devices.

2.5. Comparative Works

Several works have applied deep learning for solar panel monitoring. CNN-based methods detect defects using IR imagery [33], classify soiling from optical images [34], and identify cracks or cell failures [35,36]. Object-detection models like YOLO have also been adapted [37].

However, most approaches focus narrowly on specific defect types or single imaging modalities, limiting generalizability. Few propose integrated frameworks for both broad condition categorization and fine-grained classification. In addition, deployment challenges remain, as lightweight CNNs sacrifice accuracy, while ViTs are too computationally heavy for use on edge devices.

Our work addresses these gaps by proposing a two-tiered classification framework: Tier 1 uses a pre-trained Vision Transformer (DINOv2 ViT-Base) for coarse categorization, and Tier 2 employs an EfficientNetB0 student enhanced with KD and MHSA for fine-grained classification. This design combines strong representation learning with computational efficiency, targeting real-world edge deployment.

3. Background and Preliminaries

This section first identifies common issues like soiling and damage (Section 3.1) that hinder solar panel performance. It then introduces drone-based inspection and the crucial role of AI (Section 3.2) for efficient analysis. The section explains the fundamentals of CNNs for visual tasks (Section 3.3) and presents the efficient EfficientNet architecture (Section 3.4). Furthermore, it explores DINOv2 for robust self-supervised learning (Section 3.5) and knowledge distillation for model enhancement (Section 3.6), laying the groundwork for the proposed methodology for solar panel inspection.

3.1. Common Issues in Solar Panel Performance: Soiling and Damage

While solar PV systems provide a clean and sustainable source of energy, their efficiency can be reduced by two major factors: surface soiling and physical/electrical damage.

Soiling refers to the accumulation of foreign material on PV panels and is a substantial cause of reduced energy yield. A wide array of contaminants, including dust, bird droppings, snow, and industrial pollutants, can deposit on the panel surface and obstruct incident sunlight. The extent of energy loss varies by contaminant type, accumulation duration, and local conditions, ranging from a few percentage points in clean environments to more than 30–50% in polluted or arid regions [6]. Different contaminants have distinct impacts: dust forms light-scattering layers; bird droppings cause localized hotspots; and snow can completely block sunlight, requiring removal strategies.

In addition to soiling, physical and electrical damage pose serious threats to panel integrity and safety. Physical damage includes micro-cracks, delamination, corrosion, or glass breakage from hail, debris, or thermal stress. Such defects reduce energy generation and may expose internal components to moisture, accelerating degradation and creating safety risks.

Electrical damage, though less visible, is equally critical. It includes hotspots, bypass diode failures, interconnect corrosion, and faulty wiring, which can reduce output and even pose fire hazards. Figure 2 illustrates examples of the types of soiling and damage that commonly affect solar panels. Addressing these issues through regular inspection and maintenance is crucial to ensure long-term reliability and safety.

3.2. Drone-Based Inspection and the Role of AI

Traditional inspection of large-scale solar PV installations is labor-intensive, slow, and sometimes hazardous, requiring technicians to walk across arrays. Unmanned aerial vehicles (drones) now provide a faster, safer, and more cost-effective alternative. Equipped with RGB and thermal cameras, they can efficiently survey large solar farms. RGB imagery reveals visible surface defects such as cracks, delamination, and soiling, while thermal imaging highlights hotspots from faulty cells, shading, or electrical issues.

While drones streamline data collection, the large volume of imagery requires automated analysis. AI, particularly computer vision, enables automatic detection and classification of soiling and damage, reducing inspection time and enabling proactive maintenance. The integration of drones with AI-powered image analysis thus provides a powerful solution for scalable, efficient solar panel monitoring.

3.3. CNNs in Visual Inspection

CNNs are a cornerstone of modern computer vision that are widely used for classification, detection, and segmentation tasks. Their ability to automatically learn hierarchical spatial features from raw pixel data makes them well-suited for use in visual inspection, including for anomaly and defect detection.

A typical CNN includes convolutional, pooling, and fully connected layers. Convolutional layers learn local features such as edges and textures; pooling layers reduce spatial dimensions and add translational invariance; fully connected layers combine learned features for classification.

CNNs have been successfully applied to diverse domains such as detection of manufacturing defects, medical imaging, and infrastructure monitoring. Their ability to learn discriminative features directly from data removes the need for manual feature engineering. This makes them effective for capturing subtle anomalies that traditional handcrafted methods often miss.

Building on this foundation, CNNs have been widely adopted for inspection of solar PV installations, where they provide the basis for efficient automated monitoring. In this work, we focus on the EfficientNetB0 architecture, which balances accuracy and efficiency for large-scale deployment.

3.4. EfficientNet: Compound Scaling for Efficient and Accurate CNNs

The EfficientNet family, introduced by Tan and Le [38], is designed to achieve a balance between accuracy and efficiency in vision models. Its key innovation is the compound scaling method, which jointly scales depth, width, and resolution in a balanced way.

The base architecture, EfficientNet-B0, was obtained through neural architecture search with an emphasis on accuracy–efficiency trade-offs. It relies on three main building blocks:

- Mobile Inverted Bottleneck Convolutions (MBConv): inverted residuals and linear bottlenecks enable efficient feature extraction [39].
- Depthwise Separable Convolutions: reduce parameters and FLOPs by factorizing convolution into depthwise and pointwise operations.
- Squeeze-and-Excitation (SE) units: adaptively re-weight feature channels to emphasize informative ones with minimal extra cost.

Larger EfficientNet variants (B1–B7) are derived from B0 via compound scaling, offering different balances of accuracy and efficiency.

For solar panel inspection, EfficientNet provides flexible options depending on deployment. Smaller models (e.g., B0, B1) are suited for resource-constrained drones or edge devices, while larger ones (e.g., B4–B7) can be used in server settings where accuracy is prioritized. The architectural design of MBConv and SE units supports robust feature learning, making EfficientNet well-suited to capturing subtle soiling patterns and signs of damage while remaining computationally efficient.

3.5. DINOv2: Self-Supervised Learning for Robust Visual Features

DINOv2 [30] is a state-of-the-art self-supervised framework that learns general-purpose visual features from massive unlabeled image datasets, making them transferable across diverse domains.

The method follows a student–teacher paradigm wherein two identical Vision Transformers process different augmented views of the same image. Unlike in traditional supervised distillation, no human labels are used; instead, the student is trained to match the teacher's predictions.

Trained on large-scale data with carefully designed strategies, DINOv2 produces semantically rich and robust features that transfer effectively to tasks such as classification, detection, and segmentation.

In the context of solar panel monitoring, DINOv2 provides features that capture both structural patterns and subtle variations linked to soiling or damage, even under changing viewpoints and illumination. This makes it well-suited to serving as the backbone for our Tier-1 classifier. By fine-tuning on a smaller labeled solar dataset, we adapt DINOv2's generic features to the domain, enabling accurate and efficient analysis.

3.6. Knowledge Distillation for Model Enhancement

Knowledge Distillation (KD) [40,41] transfers the knowledge of a larger teacher model to a smaller student, enabling the latter to achieve competitive accuracy with lower complexity. The student is trained using a weighted combination of the task loss and a distillation loss that aligns its predictions with the teacher's. Soft probability distributions from the teacher, obtained via temperature-scaled softmax, provide richer information than hard labels would and guide the student toward smoother and more generalizable decision boundaries. This process is illustrated in Figure 4.

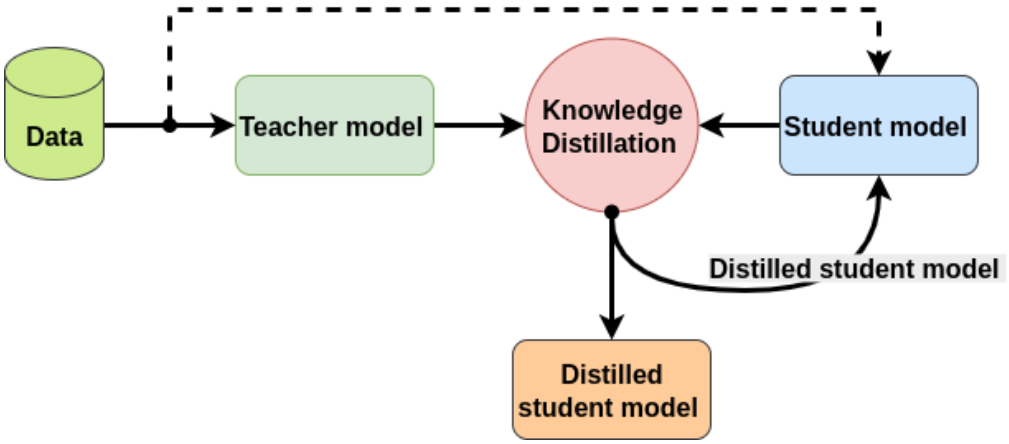

Figure 4. Simplified diagram of the Knowledge Distillation process.

The total loss is calculated as follows:

$$\mathcal{L}_{\text{total}} = \alpha \cdot \mathcal{L}_{\text{CE}}(y, \hat{y}_{\text{student}}) + (1-\alpha) \cdot T^2 \cdot \mathcal{L}_{\text{KL}}(\text{soft}_T(y_{\text{teacher}}), \text{soft}_T(\hat{y}_{\text{student}})) \quad (1)$$

with $\mathcal{L}_{\text{CE}}$ as cross-entropy, $\mathcal{L}_{\text{KL}}$ the KL divergence, and $\text{soft}_T(\cdot)$ the temperature-scaled softmax.

KD is especially beneficial when the teacher (e.g., DINOv2 ViT) encodes rich, generalizable features that a compact student (e.g., EfficientNetB0) can inherit. This boosts the accuracy of lightweight models while retaining efficiency, making them suitable for edge deployment.

In our framework, Tier-2 employs KD to enhance EfficientNetB0 by transferring knowledge from the DINOv2 teacher, which has been fine-tuned on Tier-1 outputs. This enables the student to perform fine-grained classification of soiling and damage with higher accuracy while remaining computationally efficient.

3.7. Performance Metrics Employed

To comprehensively evaluate the classification performance of the proposed models across both tiers, we employ a set of standard and widely accepted performance metrics. These metrics enable both global and class-wise analysis of the models' predictive capabilities, as described in Table 2.

These metrics jointly provide a holistic view of model performance, supporting fair comparisons across models with varying complexity and application constraints.

Table 2. Performance metrics with descriptions and formulas. TP = True Positive, TN = True Negative, FP = False Positive, FN = False Negative.

Metric	Description	Equation
Accuracy	Represents the overall correctness of the model by measuring the ratio of correctly predicted instances to the total number of instances.	$\frac{TP+TN}{TP+TN+FP+FN}$
Precision	Indicates the proportion of correctly predicted positive samples among all samples predicted as positive.	$\frac{TP}{TP+FP}$
Recall (Sensitivity)	Measures the model's ability to correctly identify actual positive samples.	$\frac{TP}{TP+FN}$
F1-Score	The harmonic mean of precision and recall, useful in class-imbalanced scenarios.	$2 \cdot \frac{\text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}}$
Confusion Matrix	A layout showing actual vs. predicted class distributions. Useful for error analysis.	$\begin{bmatrix} TP & FN \\ FP & TN \end{bmatrix}$
Inference Time	Average time (in milliseconds) to classify a single image; critical for edge deployment.	–

4. Proposed Solution

This section details the proposed two-tiered framework for classification of solar panel defects and soiling, including the dataset, overall design, and specific models used.

4.1. Dataset Composition

The dataset consists of 1131 annotated RGB images from a public Kaggle repository [42], organized into three conditions with six subclasses:

- Clean Panels (244 images)
- Surface Contaminants (504 images): Dust (191), Bird Droppings (192), Snow (124)
- Panel Degradation (378 images):Physical Damage (144), Electrical Faults (232, e.g., discoloration, burn marks, micro-cracks)

Representative samples are shown in Figure 5, with distributions across classes and subclasses.

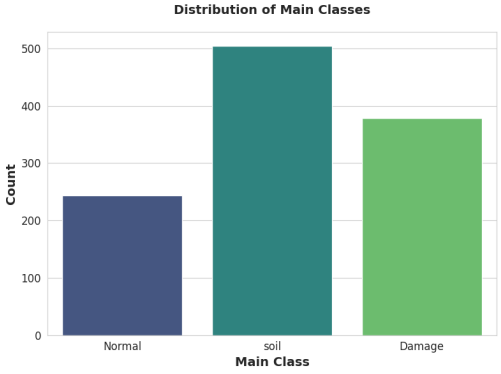

(**a**) Main class distribution

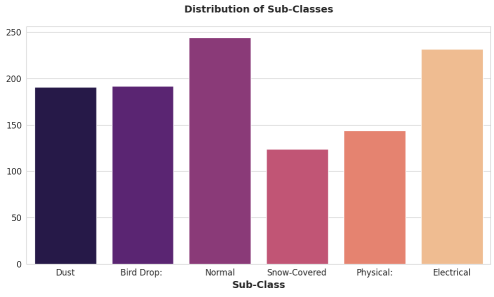

(**b**) Subclass distribution

Figure 5. Dataset distributions.

4.2. Overall Architecture

The framework is structured in two tiers:

- Tier 1: Transformer-based model classifies images as Normal, Damaged, or Soiled.
- Tier 2: Lightweight CNNs further classify subtypes as Damaged (Electrical vs. Physical) or Soiled (Dust, Bird Droppings, Snow).

If Tier-1 predicts Normal, classification ends. Otherwise, images are forwarded to Tier-2 for subtype prediction. Figure 6 illustrates the pipeline.

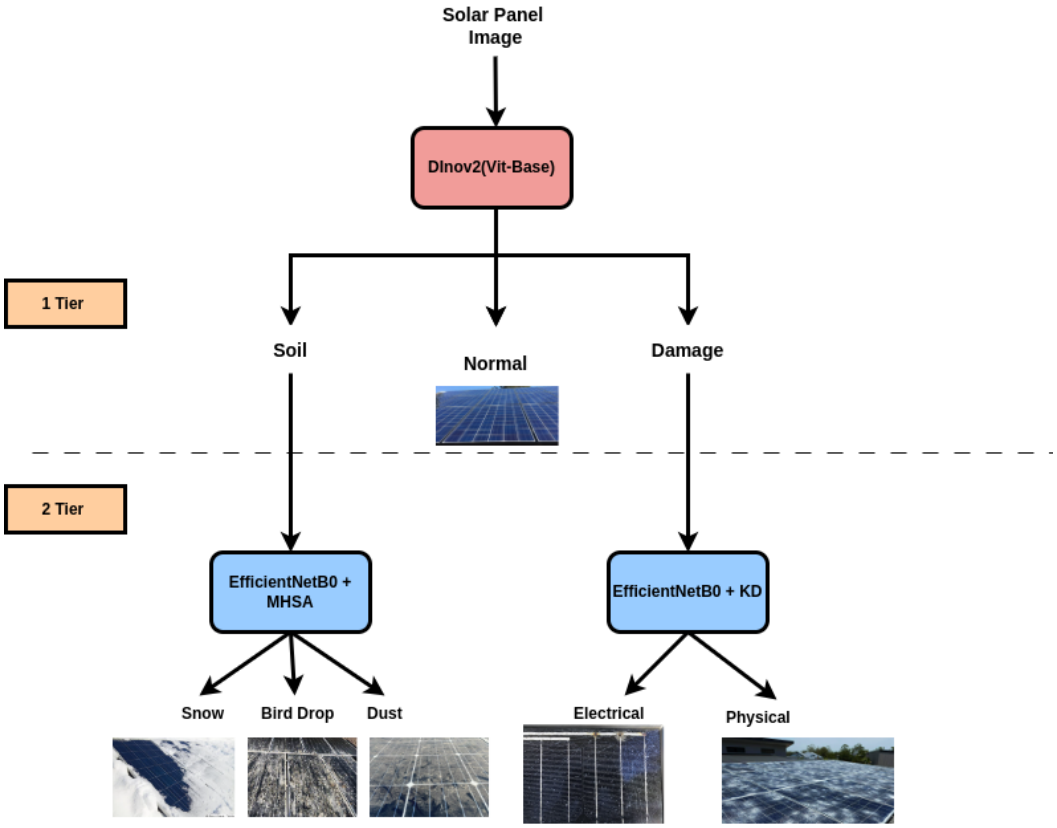

Figure 6. Proposed two-tiered framework.

4.3. Tier 1: DINOv2 ViT-Base

Tier-1 uses the DINOv2 ViT-Base/14 model. We fine-tune the classification head and `fc1` layers of MLP blocks while freezing the backbone. This selective update preserves pretrained representations while reducing training cost.

The model outputs one of the three high-level classes. The use of DINOv2 helps to manage variations in lighting, perspective, and environmental conditions in drone imagery.

4.4. Tier 2: EfficientNetB0 Enhancements

Tier-2 refines classification using EfficientNetB0, with two alternative enhancements, as summarized in Table 3:

Table 3. Tier-2 enhancement methods.

Method	Enhancement	Target Classes
EfficientNetB0 + MHSA	Global context via attention	Electrical/Physical; Dust/Bird/Snow
EfficientNetB0 + KD	Transfer of ViT features	Electrical/Physical; Dust/Bird/Snow

4.4.1. EfficientNetB0 with MHSA

We integrate a Multi-Head Self-Attention (MHSA) module after the final convolutional block. Each head computes attention using Q, K, and V projections, and outputs are concatenated and passed to a projection + classifier. This enables global context modeling, complementing convolution's local focus.

4.4.2. EfficientNetB0 with Knowledge Distillation

Alternatively, we apply KD to train EfficientNetB0 as a student of the Tier-1 DINOv2 teacher. The loss combines cross-entropy with KL divergence between teacher and student outputs. This allows the student to inherit richer decision boundaries from the teacher, improving generalization while remaining lightweight.

Outcome

Tier 1 (DINOv2) provides robust coarse classification, while Tier 2 (EfficientNetB0 + MHSA or +KD) refines the results for deployment on constrained devices.

5. Experimental Results

This section presents a comprehensive evaluation of our two-tiered classification framework designed for solar panel monitoring. The evaluation focuses on the effectiveness of the DINOv2 ViT-Base model in Tier-1 classification and the performance of two enhancement strategies applied to EfficientNetB0 in Tier 2: Knowledge Distillation and MHSA.

5.1. Tier-1 Classification

In the first stage of the proposed framework, the DINOv2 ViT-Base/14 model was fine-tuned to classify solar panel images into three high-level categories: Normal, Damaged, and Soiled. This model, selected for its balance between performance and computational efficiency, outperformed larger transformer variants explored in prior studies. By freezing the backbone and fine-tuning only the classification head, we leveraged robust, self-supervised features while minimizing computational overhead.

A quantitative comparison with six state-of-the-art models is presented in Table 4 and employs four standard metrics: Accuracy, Precision, Recall, and F1-Score. The DINOv2 ViT-Base achieved an F1-Score of 95.2% and an accuracy of 95.5%, outperforming all other contenders, including ViT-Large and EfficientNetB4.

Table 4. Performance Comparison of DINOv2 ViT-Base with Top Tier-1 Models.

Model	Params	Accuracy	Precision	Recall	F1-Score
EfficientNetB4 [43]	19M	0.933	0.934	0.927	0.929
NASNetMobile [43]	5.3M	0.928	0.933	0.919	0.923
EfficientNetB7 [43]	66M	0.923	0.930	0.911	0.917
ViT-Base [43]	86M	0.940	0.938	0.937	0.937
ViT-Large [43]	307M	0.941	0.944	0.933	0.937
DeiT-Base-Distilled [43]	86M	0.927	0.931	0.921	0.924
DINOv2 ViT-Base	86M	**0.955**	**0.943**	**0.962**	**0.952**

The confusion matrix in Figure 7 offers further insight into the model's per-class prediction performance. The model demonstrates a high true-positive rate across all three classes, with minimal misclassification. In particular, the high recall (96.2%) reflects the model's effectiveness in correctly identifying defective and soiled panels, a critical requirement for minimizing false negatives in maintenance systems for solar panels.

Qualitative results are presented in Figure 8, showcasing example predictions across diverse panel conditions. These examples highlight the model's robustness under varied lighting, soiling patterns, and damage manifestations.

Overall, the DINOv2 ViT-Base demonstrates excellent generalization and reliability, establishing a strong foundation for fine-grained classification in Tier 2. Its performance justifies its use over larger models and aligns with the goal of maintaining computational efficiency without compromising accuracy.

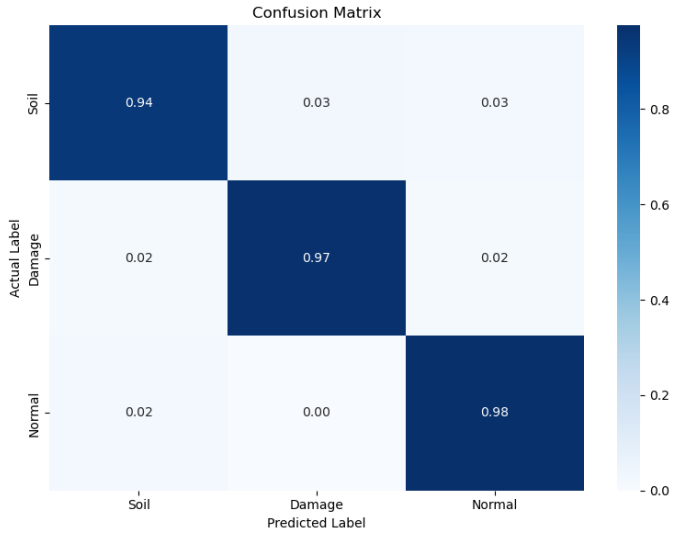

Figure 7. Normalized confusion matrix for DINOv2 ViT-Base on Tier-1 classification (Normal, Damaged, Soiled).



Figure 8. Tier-1 prediction examples from the DINOv2 ViT-Base model across Normal, Damaged, and Soiled categories.

5.2. Tier-2 Classification

In this section, we present a fine-grained analysis of classification of solar panel images, focusing on soiling and damage types. Two enhancement strategies, KD and MHSA, are explored in conjunction with lightweight CNN architectures to improve performance while maintaining efficiency.

5.2.1. Performance in Soiling Classification

To assess the effectiveness of our fine-grained soiling-classification module, we evaluated the impact of two enhancement strategies applied to the EfficientNetB0 architecture: KD and MHSA. The goal was to improve recognition of three common soiling types on PV panels: Dust, Bird Droppings, and Snow.

In the KD setup, a DINOv2 ViT-S/14 model served as the teacher, transferring its rich feature representations to the student EfficientNet-based models. In parallel, the MHSA mechanism was introduced to capture global spatial dependencies. The standalone performance of the DINOv2 ViT-S/14 teacher on this task is reported in Table 5, while Table 6 presents the comparative evaluation of the enhanced models against several lightweight baselines.

Table 5. Performance of DINOv2 ViT-S/14 (teacher) in classification of soiling types (Tier 2).

Model	Task	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)
DINOv2 ViT-S/14	Soiling Classification	94.29	95.15	95.15	95.00

Table 6. Performance Comparison of Soiling Classification Models.

Model	Params	Accuracy	Precision	Recall	F1-Score	Inference Time (ms)
EfficientNetB0 [43]	5.3M	0.907	0.922	0.911	0.913	0.142
MobileNetV2 [43]	3.4M	0.813	0.843	0.833	0.820	0.095
EfficientNetB4 [43]	19.0M	0.893	0.917	0.900	0.903	0.283
EfficientNetB7 [43]	66.0M	0.907	0.921	0.913	0.914	0.735
B0 KD Model	5.3M	0.924	0.928	0.934	0.931	0.142
B4 KD Model	19.0M	0.840	0.863	0.851	0.849	0.283
B7 KD Model	66.0M	0.914	0.921	0.911	0.915	0.735
MobileNetV2 KD Model	3.4M	0.895	0.912	0.894	0.901	0.095
B0 + MHSA	6.0M	0.933	0.941	0.942	0.942	0.152
B4 + MHSA	19.7M	0.867	0.869	0.874	0.871	0.278
B7 + MHSA	66.7M	0.924	0.933	0.934	0.933	0.745
MobileNetV2 + MHSA	4.1M	0.918	0.923	0.923	0.921	0.108

The confusion matrices in Figure 9 provide a class-wise analysis of the best-performing KD and MHSA models based on EfficientNetB0. Notable findings include:

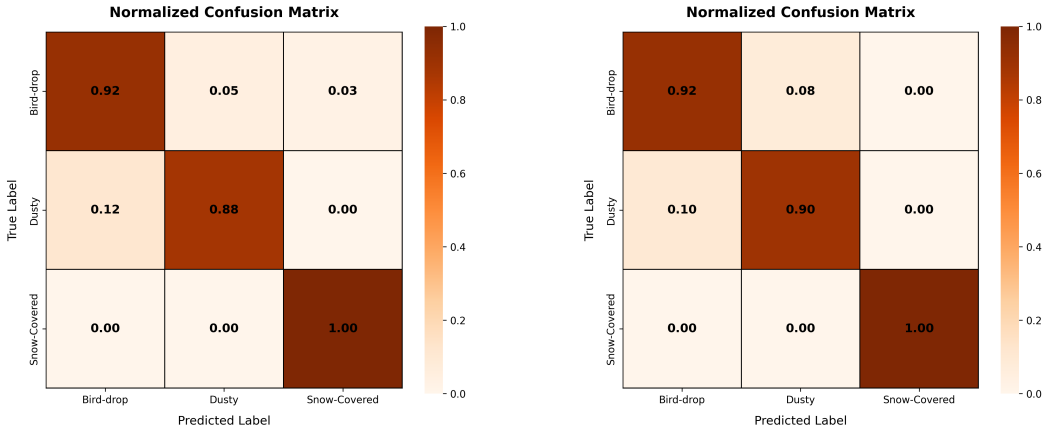

(**a**) (**b**)

Figure 9. Normalized confusion matrices for soiling classification. Columns: predicted labels; rows: true labels. (**a**) KD Model (EfficientNetB0). Accuracy: 92.38%, F1-Score: 93.05%. (**b**) MHSA Model (EfficientNetB0). Accuracy: 93.33%, F1-Score: 94.17%.

- The KD model shows strong performance for the Bird Droppings class, with a 95.2% true-positive rate and low false-positive rate (4.2%).
- The most common error for the KD model involves confusion between Snow and Dust (8.4%).
- The MHSA-enhanced model significantly improves Snow detection (TPR: 94.7%) and reduces confusion between Dusty and Snow to 5.1%.
- Overall, the MHSA model achieves consistently high true-positive rates across all classes, indicating balanced and robust classification.

Qualitative examples of soiling predictions are provided in Figure 10, further illustrating the model's effectiveness in distinguishing visually similar soiling conditions under varying lighting and occlusion patterns.



Figure 10. Sample predictions for soiling classification across the Dust, Bird Droppings, and Snow categories.

5.2.2. Performance in Damage Classification

To evaluate the robustness and generalization of the proposed models on the binary classification task of identifying Electrical versus Physical defects in solar panels, we applied both KD and MHSA techniques. The KD framework leveraged a DINOv2 ViT-S/14 teacher model with 16-head Multi-Head Attention, transferring its strong representational capacity to lightweight student architectures such as EfficientNetB0, B4, and B7. In parallel, we evaluated the performance gains achieved by incorporating MHSA alone into these student models.

The standalone performance of the teacher model is summarized in Table 7, which shows perfect classification metrics across all evaluated indicators.

Table 7. Performance of DINOv2 ViT-S/14 (teacher) in damage classification (Tier 2).

Model	Task	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)
DINOv2 ViT-S/14	Damage Classification	99.98 ± 0.1	99.98 ± 0.1	99.98 ± 0.1	99.98 ± 0.1

The results comparing KD-enhanced and MHSA-based models against baseline architectures are presented in Table 8, highlighting improvements in both accuracy and inference efficiency.

Table 8. Comparison of performance of damage-classification models.

Model	Params	Accuracy	Precision	Recall	F1-Score	Inference Time (ms)
EfficientNetB0 [43]	5.3M	0.967	0.976	0.961	0.968	0.142
MobileNetV2 [43]	3.4M	0.901	0.898	0.907	0.898	0.095
EfficientNetB4 [43]	19.0M	0.952	0.933	0.966	0.949	0.283
EfficientNetB7 [43]	66.0M	0.943	0.956	0.936	0.945	0.735
B0 KD Model	5.3M	0.986	0.982	0.989	0.985	0.142
B4 KD Model	19.0M	0.961	0.954	0.961	0.957	0.283
B7 KD Model	66.0M	0.987	0.982	0.990	0.986	0.735
MobileNetV2 KD Model	3.4M	0.947	0.937	0.950	0.943	0.095
B0 + MHSA	6.0M	0.987	0.982	0.990	0.986	0.152
B4 + MHSA	19.7M	0.992	0.990	0.991	0.990	0.278
B7 + MHSA	66.7M	0.981	0.975	0.982	0.978	0.745
MobileNetV2 + MHSA	4.1M	0.943	0.927	0.952	0.936	0.108

To support a better understanding of class-wise performance, normalized confusion matrices for the top KD and MHSA models are presented in Figure 11. These visualizations highlight key improvements in both recall and precision for the Electrical and Physical damage classes.

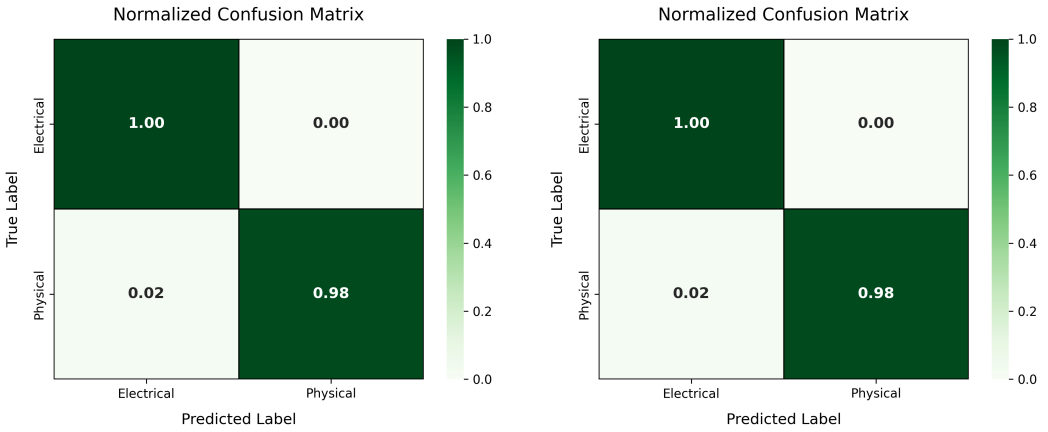

(**a**) (**b**)

Figure 11. Normalized confusion matrices for damage-classification models. Columns: predicted labels; rows: true labels. (**a**) KD Model (EfficientNetB7). Accuracy: 98.67%, F1-Score: 98.56%. (**b**) MHSA Model (EfficientNetB4). Accuracy: 99.20%, F1-Score: 99.00%.

Key observations include the following:

- The KD model based on EfficientNetB7 achieves a recall of 99.0% for Electrical damage and a low false-positive rate (1.8%) for Physical damage, improving over the EfficientNetB0 baseline.
- The MHSA model based on EfficientNetB4 achieves a near-perfect recall of 99.1% for both classes and reduces the confusion between classes to below 1%.
- Both models significantly outperform the baseline EfficientNetB0 in terms of both precision and recall while maintaining efficient inference times.

Qualitative examples of predictions across both damage categories are shown in Figure 12, illustrating the models' ability to generalize to unseen samples.

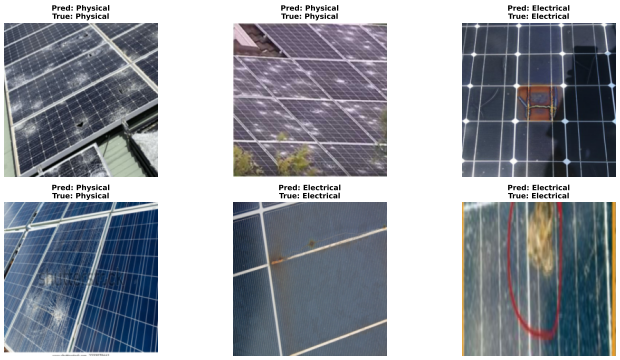

Figure 12. Sample predictions for binary damage classification across Electrical and Physical categories.

5.3. Model Efficiency Analysis

In this section, we analyze the efficiency of the proposed KD and MHSA models in terms of parameter count, classification error patterns, and performance improvements. These results are compared with the EfficientNetB0 baseline to examine differences in model complexity and predictive accuracy.

- Parameter Efficiency:
  - The KD model exhibits a streamlined architecture with 4.66 million parameters, which is notably close to the EfficientNetB0 baseline (5.3M, Table 6).
  - MHSA-enhanced models introduce a moderate increase in model complexity, ranging from 6.01M to 7.26M parameters. This increase is justified by consistent performance gains (as evidenced in Tables 6 and 8).
  - Specifically for soiling classification, the KD model achieves superior accuracy and a higher F1-Score with parameter counts comparable to the baseline, highlighting efficient representational learning. MHSA models, though slightly larger, deliver further accuracy improvements at modest computational cost.
- Error Patterns:
  - For soiling classification, both the KD and MHSA models exhibit residual confusion between the Snow and Dust categories, which appears to be a persistent challenge due to their visual similarity. While baseline errors are not visualized, the KD and MHSA models substantially mitigate misclassification rates.
  - In damage classification, the confusion matrices in Figure 11 demonstrate low false-positive and false-negative rates, where panel (a) corresponds to KD and panel (b) to MHSA. This represents a notable improvement over the baseline model, whose error rates are reported in Table 8.
- Performance Gains:
  - As detailed in Table 6, the KD model yields an accuracy gain of +1.68% and an F1-score improvement of +1.75% for soiling classification relative to the EfficientNetB0 baseline. MHSA integration further increases accuracy by +2.63% and the F1-score by +2.87%.
  - In damage classification (Table 8), the KD model outperforms the baseline by +1.97% in accuracy and +1.76% in F1-score. Similarly, the MHSA-enhanced variant shows an accuracy gain of +2.03% and an F1-score improvement of +2.21%.

6. Discussion

The experimental results across both classification tiers demonstrate the effectiveness of combining advanced transformer architectures with model efficiency techniques for the detection of solar panel defects. In Tier 1, the fine-tuned DINOv2 ViT-Base model

substantially outperformed conventional architectures, achieving an F1-score of 95.2% with only 86M parameters; this result highlights its suitability for general triage tasks in field deployments. The confusion matrix revealed high sensitivity across all three classes (Normal, Damaged, and Soiled), supporting its robustness in real-world scenarios.

In Tier 2, the application of KD and MHSA demonstrated tangible benefits across both soiling and damage classification subtasks. The KD-based EfficientNetB0 model achieved a notable accuracy improvement of +1.68% and F1-score gain of +1.75% for soiling classification, despite maintaining a parameter count similar to the baseline. The MHSA-enhanced variant further improved performance with marginal increases in complexity.

For damage classification, MHSA yielded the highest accuracy (99.2%) and F1-score (99.0%) among all Tier 2 models, highlighting its strength in capturing complex inter-class relationships. Confusion matrix analysis confirmed reduced misclassification rates for both Electrical and Physical defect classes, indicating improved class separability.

Efficiency analysis revealed that performance gains were achieved without incurring excessive computational costs. All enhanced models remained within a practical range of parameter counts and inference latency, making them viable for deployment on edge devices in inspection systems for photovoltaic installations.

Overall, the results validate the proposed two-tiered framework as a balanced solution for accurate and efficient defect detection. While DINOv2 serves as a strong base classifier, its distilled knowledge and attention mechanisms can be effectively transferred to lightweight student models, enabling real-time, high-accuracy predictions in constrained environments.

7. Conclusions

This paper presented a two-tiered deep learning framework for efficient and accurate classification of solar panel defects, leveraging a combination of Vision Transformers, Knowledge Distillation, and Multi-Head Self-Attention. In the first tier, a fine-tuned DINOv2 ViT-Base model provided high-performance triage of solar panel images into general categories (Normal, Damaged, and Soiled). The second tier incorporated KD and MHSA techniques to refine classification across six specialized defect classes, achieving state-of-the-art performance while preserving computational efficiency.

Experimental evaluations across multiple CNN and transformer-based baselines demonstrated that the proposed models outperform existing architectures in terms of accuracy, precision, recall, and F1-score. Notably, KD-based EfficientNetB0 and MHSA-enhanced variants consistently surpassed the performance of their non-enhanced counterparts, validating the effectiveness of knowledge transfer and attention mechanisms in improving class discrimination.

Furthermore, the proposed models maintain low parameter counts and fast inference times, making them well-suited for deployment in edge-based photovoltaic monitoring systems. The confusion matrix analysis confirmed the models' ability to generalize well across visually similar defect types; they show a reduced tendency towards the common misclassification trends observed in baseline models.

In future work, this two-tiered framework can be extended with real-time object detection for localized defect identification and adapted to other renewable infrastructure domains such as wind turbines. Additionally, integrating self-supervised learning and hardware-aware optimization (e.g., quantization, pruning) will further enhance its deployment readiness for resource-constrained environments.

Author Contributions: Conceptualization, A.H. and H.N.N.; methodology, A.H. and H.N.N.; investigation, A.H. and H.N.N.; writing—original draft preparation, A.H. and H.N.N.; writing—review and editing, A.H. and H.N.N.; visualization, A.H. and H.N.N.; supervision, H.N.N. and J.A.; project administration, H.N.N. All authors have read and agreed to the published version of the manuscript.

Funding: This paper was supported by funds from the Bourgogne-Franche-Comté Region (AAP Région 2022 dispositif ANER-IA Limentaire project) and by the EIPHI Graduate School (contract "ANR-17-EURE-0002").

Data Availability Statement: No new data were created or analyzed in this study. Data sharing is not applicable to this article.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. U.S. Department of Energy. Solar Photovoltaic Cell Basics. 2025. Available online: https://www.energy.gov/eere/solar/solar-photovoltaic-cell-basics (accessed on 24 April 2025).
2. Adekanbi, M.L.; Alaba, E.S.; John, T.J.; Tundealao, T.D.; Banji, T.I. Soiling loss in solar systems: A review of its effect on solar energy efficiency and mitigation techniques. *Clean. Energy Syst.* **2024**, *7*, 100094. [CrossRef]
3. Bassil, J.; Noura, H.; Salman, O.; Chahine, K.; Guizani, M. Deep Learning Image Classification Models for Solar Panels Dust Detection. In Proceedings of the 2024 International Wireless Communications and Mobile Computing (IWCMC), Ayia Napa, Cyprus, 27–31 May 2024; pp. 1516–1521.
4. Ozturk, O.; Hangun, B.; Eyecioglu, O. Detecting snow layer on solar panels using deep learning. In Proceedings of the 2021 10th International Conference on Renewable Energy Research and Application (ICRERA), Istanbul, Turkey, 26–29 September 2021; pp. 434–438.
5. Sisodia, A.K.; Mathur, R.K. Impact of bird dropping deposition on solar photovoltaic module performance: A systematic study in Western Rajasthan. *Environ. Sci. Pollut. Res.* **2019**, *26*, 31119–31132. [CrossRef]
6. Hussain, A.; Batra, A.; Pachauri, R. An experimental study on effect of dust on power loss in solar photovoltaic module. *Renew. Wind Water Sol.* **2017**, *4*, 9. [CrossRef]
7. Akram, M.W.; Li, G.; Jin, Y.; Chen, X. Failures of Photovoltaic modules and their Detection: A Review. *Appl. Energy* **2022**, *313*, 118822. [CrossRef]
8. Hasan, A.A.; Ahmed Alkahtani, A.; Shahahmadi, S.A.; Nur E. Alam, M.; Islam, M.A.; Amin, N. Delamination-and electromigration-related failures in solar panels—A review. *Sustainability* **2021**, *13*, 6882. [CrossRef]
9. Renewable Energy Magazine. Above Partners with Leading Universities to Develop Next-Generation Drone Technology for Intelligent Solar Plant Inspections. 2025. Available online: https://www.renewableenergymagazine.com/pv_solar/above-partners-with-leading-universities-to-develop-20210611 (accessed on 22 April 2025).
10. Bounabi, M.; Azmi, R.; Chenal, J.; Diop, E.B.; Ebnou Abdem, S.A.; Adraoui, M.; Hlal, M.; Serbouti, I. Smart PV Monitoring and Maintenance: A Vision Transformer Approach within Urban 4.0. *Technologies* **2024**, *12*, 192. [CrossRef]
11. Amiri, A.F.; Kichou, S.; Oudira, H.; Chouder, A.; Silvestre, S. Fault detection and diagnosis of a photovoltaic system based on deep learning using the combination of a convolutional neural network (cnn) and bidirectional gated recurrent unit (Bi-GRU). *Sustainability* **2024**, *16*, 1012. [CrossRef]
12. Aboagye, B.; Gyamfi, S.; Ofosu, E.A.; Djordjevic, S. Investigation into the impacts of design, installation, operation and maintenance issues on performance and degradation of installed solar photovoltaic (PV) systems. *Energy Sustain. Dev.* **2022**, *66*, 165–176. [CrossRef]
13. Ansari, S.; Ayob, A.; Lipu, M.S.H.; Saad, M.H.M.; Hussain, A. A review of monitoring technologies for solar PV systems using data processing modules and transmission protocols: Progress, challenges and prospects. *Sustainability* **2021**, *13*, 8120. [CrossRef]
14. Shafiullah, M.; Ahmed, S.D.; Al-Sulaiman, F.A. Grid integration challenges and solution strategies for solar PV systems: A review. *IEEE Access* **2022**, *10*, 52233–52257. [CrossRef]
15. Hijjawi, U.; Lakshminarayana, S.; Xu, T.; Fierro, G.P.M.; Rahman, M. A review of automated solar photovoltaic defect detection systems: Approaches, challenges, and future orientations. *Sol. Energy* **2023**, *266*, 112186. [CrossRef]
16. Mohammad, A.; Mahjabeen, F. Revolutionizing solar energy with ai-driven enhancements in photovoltaic technology. *Bullet J. Multidisiplin Ilmu* **2023**, *2*, 1174–1187.
17. Cheragee, S.H.; Hassan, N.; Ahammed, S.; Islam, A. A study of IoT based real-time solar power remote monitoring system. *Int. J. Ambient Syst. Appl.* **2021**, *9*, 27–36.
18. Shuvo, M.M.H.; Islam, S.K.; Cheng, J.; Morshed, B.I. Efficient acceleration of deep learning inference on resource-constrained edge devices: A review. *Proc. IEEE* **2022**, *111*, 42–91. [CrossRef]

19. Alzubaidi, L.; Zhang, J.; Humaidi, A.J.; Al-Dujaili, A.; Duan, Y.; Al-Shamma, O.; Santamaría, J.; Fadhel, M.A.; Al-Amidie, M.; Farhan, L. Review of deep learning: Concepts, CNN architectures, challenges, applications, future directions. *J. Big Data* **2021**, *8*, 53. [CrossRef] [PubMed]
20. Zhang, X. The AlexNet, LeNet-5 and VGG NET applied to CIFAR-10. In Proceedings of the 2021 2nd International Conference on Big Data & Artificial Intelligence & Software Engineering (ICBASE), Zhuhai, China, 24–26 September 2021; pp. 414–419.
21. Muhammad, U.; Wang, W.; Chattha, S.P.; Ali, S. Pre-trained VGGNet architecture for remote-sensing image scene classification. In Proceedings of the 2018 24th International Conference on Pattern Recognition (ICPR), Beijing, China, 20–24 August 2018; pp. 1622–1627.
22. Xu, W.; Fu, Y.L.; Zhu, D. ResNet and its application to medical image processing: Research progress and challenges. *Comput. Methods Programs Biomed.* **2023**, *240*, 107660. [CrossRef] [PubMed]
23. Koonce, B. EfficientNet. In *Convolutional Neural Networks with Swift for Tensorflow: Image Recognition and Dataset Categorization*; Springer: Berlin/Heidelberg, Germany, 2021; pp. 109–123.
24. Maity, R.; Shamaun Alam, M.; Pati, A. An approach for detection of dust on solar panels using CNN from RGB dust image to predict power loss. In *Cognitive Computing in Human Cognition: Perspectives and Applications*; Springer: Berlin/Heidelberg, Germany, 2020; pp. 41–48.
25. Sharma, J.; Khattar, S.; Verma, T. Design and Development of Soil Detection Framework From Solar Panel Using Deep Learning. In Proceedings of the 2024 OPJU International Technology Conference (OTCON) on Smart Computing for Innovation and Advancement in Industry 4.0, Raigarh, India, 5–7 June 2024; pp. 1–7.
26. Hassan, S.; Dhimish, M. A survey of CNN-based approaches for crack detection in solar PV modules: Current trends and future directions. *Solar* **2023**, *3*, 663–683. [CrossRef]
27. Khan, S.; Naseer, M.; Hayat, M.; Zamir, S.W.; Khan, F.S.; Shah, M. Transformers in vision: A survey. *ACM Comput. Surv. (CSUR)* **2022**, *54*, 1–41. [CrossRef]
28. Wolf, T.; Debut, L.; Sanh, V.; Chaumond, J.; Delangue, C.; Moi, A.; Cistac, P.; Rault, T.; Louf, R.; Funtowicz, M.; et al. Transformers: State-of-the-art natural language processing. In Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations, Online, 16–20 November 2020; pp. 38–45.
29. Caron, M.; Touvron, H.; Misra, I.; Jégou, H.; Mairal, J.; Bojanowski, P.; Joulin, A. Emerging properties in self-supervised vision transformers. In Proceedings of the IEEE/CVF International Conference on Computer Vision, Montreal, BC, Canada, 11–17 October 2021; pp. 9650–9660.
30. Oquab, M.; Darcet, T.; Moutakanni, T.; Vo, H.; Szafraniec, M.; Khalidov, V.; Fernandez, P.; Haziza, D.; Massa, F.; El-Nouby, A.; et al. Dinov2: Learning robust visual features without supervision. *arXiv* **2023**, arXiv:2304.07193.
31. Lin, Y.; Wang, C.; Song, H.; Li, Y. Multi-head self-attention transformation networks for aspect-based sentiment analysis. *IEEE Access* **2021**, *9*, 8762–8770. [CrossRef]
32. Li, X.; Xu, M.; Liu, S.; Sheng, H.; Wan, J. Ultra-Lightweight Feature-Compressed Multi-Head Self-Attention Learning Networks for Hyperspectral Image Classification. *IEEE Trans. Geosci. Remote Sens.* **2024**, *62*, 1–14.
33. Pierdicca, R.; Paolanti, M.; Felicetti, A.; Piccinini, F.; Zingaretti, P. Automatic faults detection of photovoltaic farms: SolAIr, a deep learning-based system for thermal images. *Energies* **2020**, *13*, 6496. [CrossRef]
34. Mehta, S.; Azad, A.P.; Chemmengath, S.A.; Raykar, V.; Kalyanaraman, S. Deepsolareye: Power loss prediction and weakly supervised soiling localization via fully convolutional networks for solar panels. In Proceedings of the 2018 IEEE Winter Conference on Applications of Computer Vision (WACV), Lake Tahoe, NV, USA, 12–15 March 2018; pp. 333–342.
35. Umar, S.; Nawaz, M.U.; Qureshi, M.S. Deep learning approaches for crack detection in solar PV panels. *Int. J. Adv. Eng. Technol. Innov.* **2024**, *1*, 50–72.
36. Parikh, H.R.; Buratti, Y.; Spataru, S.; Villebro, F.; Reis Benatto, G.A.D.; Poulsen, P.B.; Wendlandt, S.; Kerekes, T.; Sera, D.; Hameiri, Z. Solar cell cracks and finger failure detection using statistical parameters of electroluminescence images and machine learning. *Appl. Sci.* **2020**, *10*, 8834. [CrossRef]
37. Huang, J.; Zeng, K.; Zhang, Z.; Zhong, W. Solar panel defect detection design based on YOLO v5 algorithm. *Heliyon* **2023**, *9*, e18826. [CrossRef]
38. Tan, M.; Le, Q. Efficientnet: Rethinking model scaling for convolutional neural networks. In Proceedings of the International Conference on Machine Learning, Long Beach, CA, USA, 9–15 June 2019; pp. 6105–6114.
39. Sandler, M.; Howard, A.; Zhu, M.; Zhmoginov, A.; Chen, L.C. Mobilenetv2: Inverted residuals and linear bottlenecks. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Salt Lake City, UT, USA, 18–22 June 2018; pp. 4510–4520.
40. Hinton, G.; Vinyals, O.; Dean, J. Distilling the knowledge in a neural network. *arXiv* **2015**, arXiv:1503.02531. [CrossRef]
41. Gou, J.; Yu, B.; Maybank, S.J.; Tao, D. Knowledge distillation: A survey. *Int. J. Comput. Vis.* **2021**, *129*, 1789–1819. [CrossRef]

42. hemanthsai7. Solar Panel Dust Detection. 2022. Available online: https://www.kaggle.com/datasets/pythonafroz/solar-panel-images (accessed on 26 February 2024).
43. Noura, H.N.; Chahine, K.; Bassil, J.; Abou Chaaya, J.; Salman, O. Efficient combination of deep learning models for solar panel damage and soiling detection. *Measurement* **2025**, *251*, 117185. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.