Research Article

Deep Convolutional Neural Networks for Automated Defective Cell Detection and Classification in Photovoltaic Electroluminescence Images

Mikey Michael[1,*]
College of Engineering, University of Bamberg, Germany
* Corresponding author: Mikey Michael
MikeyMichael@gmail.com

***Abstract*:** As photovoltaic systems are threatened by microstructural degradation, there is a need to shift from traditional manual inspection methods to automated diagnostic frameworks. This study explores the application of deep convolutional neural networks in the accurate identification and classification of multidimensional defects in electroluminescent images of silicon solar cells. An improved ResNet architecture with a dual attention module is deployed to enhance sensitivity in recognizing subtle intensity gradients. However, the algorithm's potential bias in low-contrast environments, influenced to some extent by the specific contrast characteristics of the dataset, suggests that further research is needed to ensure its universal applicability across different manufacturing standards.

Keywords: *Photovoltaic Reliability; Electroluminescence Imaging; Deep Convolutional Neural Networks; Automated Defect Classification; Micro-crack Detection；*

1. Introduction

The global escalation of photovoltaic installations, driven by the imperative to decarbonize energy infrastructures, has shifted the focus of the industry from initial deployment toward long-term operational reliability. As solar assets age, they are increasingly susceptible to a range of structural degradations—many of which, such as micro-cracks and potential-induced degradation, remain fundamentally imperceptible to the human eye under standard operating conditions. The failure to identify these defects early can lead to catastrophic power loss and, in extreme cases, localized overheating that threatens the physical integrity of the entire array. Traditional maintenance regimes, which rely heavily on manual inspection and periodic thermography, are proving inadequate for the sheer scale of modern solar farms. Considering the above factors, the transition toward automated, high-resolution diagnostic frameworks is no longer merely a technological preference but a necessity for ensuring the economic viability of sustainable energy. Prior work in industrial inspection consistently argues that scaling inspection throughput and reliability often requires learning-based automated detection rather than manual screening[1][3][4].

Electroluminescence imaging has emerged as a gold-standard technique for non-destructive testing in photovoltaics, functioning on the principle of radiative recombination where the solar cell is biased as a light-emitting diode. This allows for the visualization of inactive regions as dark spots or lines, providing a detailed map of the internal carrier transport efficiency. Previous investigations, particularly those focusing on manual feature extraction, have highlighted the sensitivity of EL to mechanical stresses; however, these studies often relied on a static set of heuristic rules that failed to account for the textural noise inherent in polycrystalline modules. Earlier attempts to automate this process using classical computer vision—such as the application of Gabor filters or Haar-like features—demonstrated moderate success in controlled laboratory settings but struggled with the stochastic complexity of field-acquired data. This leads us to further thinking regarding the limitations of rigid, human-defined descriptors when faced with the subtle, non-linear patterns of real-world degradation.

The advent of deep learning, specifically Convolutional Neural Networks, has fundamentally redefined the landscape of automated defect detection. Unlike classical algorithms, CNNs are capable of hierarchical feature learning, which allows the model to prioritize relevant semantic information while suppressing irrelevant background noise. Research utilizing architectures like VGG-16 and earlier iterations of ResNet has shown promising results in identifying macro-defects; Comparable CNN backbones and transfer-learning setups have also been used in several defect inspection applications beyond photovoltaics[2][5]. However, these studies frequently neglected the class imbalance problem where "healthy" cells vastly outnumber "defective" ones. Furthermore, the reliance on ImageNet pre-trained weights, while beneficial for general feature extraction, might to some extent ignore the specific radiometric nuances of EL imagery. Further research is needed to determine how specialized attention mechanisms can be integrated to heighten the sensitivity toward faint, elongated features such as micro-cracks, which are often dismissed as noise by standard pooling operations.

This study aims to develop a robust, end-to-end deep learning framework specifically tailored for the detection of multi-class defects in EL images. Our primary objective is to transcend the binary classification paradigm by implementing a model capable of distinguishing between various failure modes, such as finger failures, micro-cracks, and snail trails. By acknowledging the inherent uncertainty in pixel-level labeling, we propose a methodology that incorporates both structural and contextual information. The following chapters will detail the construction of a high-fidelity EL dataset, the architectural modifications required for low-contrast defect localization, and a comprehensive evaluation of the model's performance under varying environmental noise levels.

2. Data Acquisition and Preprocessing

The efficacy of deep learning models in the context of EL inspection is fundamentally contingent upon the diversity and quality of the training corpus. Our dataset was synthesized from multiple industrial sources, encompassing monocrystalline and polycrystalline silicon modules captured under varying bias currents. We encountered significant difficulty during the initial data aggregation phase, as the exposure times and sensor sensitivities of the EL cameras varied across different manufacturing facilities. This heterogeneity introduced a systemic bias in image intensity that threatened the model's ability to generalize. To address this, we implemented a global normalization protocol, yet even with consistent scaling, the intrinsic "graininess" of polycrystalline cells remained a confounding factor. This leads us to further thinking about whether the texture of the silicon substrate itself should be treated as a feature or filtered as extraneous interference.

Raw EL images of PV strings often suffer from perspective distortion due to the angular constraints of camera positioning during field inspections. To resolve this, we employed a semi-automated perspective transformation technique based on the detection of the frame corners. The process was not as linear as idealized in theoretical literature; the low contrast between the cell edges and the module frame often led to erroneous corner detection, requiring the implementation of a robust RANSAC-based line fitting algorithm to refine the quadrilateral boundaries. Following rectification, the modules were segmented into individual cells using a projection-based grid detection method. We noted that slight misalignments in the grid could lead to the inclusion of busbar shadows at the edges of the cell, which the model might erroneously interpret as edge cracks. This necessitated a secondary cropping adjustment, showing the genuine trial-and-error nature of preprocessing in complex imaging environments.

The labeling process involved the categorization of cells into five distinct classes: Intact, Finger Failure, Micro-crack, Black Core, and Broken Cell. During the annotation phase, we faced considerable ambiguity in distinguishing between "hairline" micro-cracks and surface scratches that do not impact electronic activity[17][26][17]. To mitigate this subjective bias, each image was cross-verified by two independent experts, yet a third arbiter was required for approximately 15% of the samples. This friction in the ground-truth generation process underscores the inherent uncertainty in interpreting EL signatures. Further research is needed to quantify how much of the model's error rate is attributable to these labeling inconsistencies rather than architectural flaws. We intentionally retained a small subset of "ambiguous" samples in the training set to, to some extent, force the network to learn more resilient decision boundaries, rather than overfitting to perfectly clean examples[10][13][18].

Given the relative rarity of broken cells compared to intact ones, our initial training runs resulted in a model that was heavily biased toward predicting the majority class[25][19]. We resisted the simplistic approach of random oversampling, which often leads to memorization of training instances. Instead, we explored more sophisticated augmentation techniques, including elastic deformations and random brightness fluctuations, to mimic the varying atmospheric conditions of field inspections. We also experimented with a Generative Adversarial Network (GAN) to synthesize rare defect types; however, the GAN-generated images often lacked the high-frequency textural detail required for realistic EL simulation, leading to a "domain gap" that hindered performance. Consequently, we reverted to a weighted loss function combined with rigorous geometric transformations. This adjustment reflects the real-world necessity of prioritizing data integrity over the allure of complex synthetic generation when the latter fails to meet physical accuracy standards[9][24][28].

3. Methodology

The architectural foundation of an automated EL inspection system must balance the depth required for complex feature extraction with the computational efficiency necessitated by real-time industrial throughput. In this study, we initially gravitated toward the ResNet-50 framework due to its proven efficacy in mitigating the vanishing gradient problem through residual learning; Industrial inspection studies often emphasize architecture choices that stabilize training while preserving discriminative detail in feature maps[20][21][22][23][27]. However, our early experiments indicated that standard residual blocks occasionally smoothed over the high-frequency textural details characteristic of hairline micro-cracks. This led us to further thinking regarding the necessity of a more "attention-aware" backbone that does not merely aggregate features but actively weighs their spatial importance. The decision to modify the skip-connection mechanism was not a straightforward linear progression but rather a response to the observation that traditional pooling layers were, to some extent, discarding critical radiometric information essential for identifying low-contrast defects.

3.1 Implementation of the Dual-Attention Mechanism

To address the aforementioned loss of subtle features, we integrated a Dual-Attention Module (DAM) that operates across both channel and spatial dimensions. Unlike the standard Squeeze-and-Excitation (SE) blocks, which primarily focus on inter-channel correlations, our DAM attempts to localize defect-prone regions within the EL image by emphasizing long-range dependencies. During the development phase, we encountered significant difficulty in calibrating the attention coefficients, as an overly aggressive spatial weighting tended to amplify sensor noise in dark-core regions, leading to a high rate of false positives. This required a rigorous adjustment of the sigmoid activation thresholds, reflecting the genuine friction between algorithmic sensitivity and noise suppression in non-destructive testing environments.

3.2 Loss Function Optimization for Class Imbalance

A persistent hurdle in solar cell defect detection is the extreme rarity of certain failure modes, such as broken cells or snail trails, relative to the abundance of intact samples. Standard cross-entropy loss functions frequently allow the majority class to dominate the gradient descent process, resulting in a model that possesses high overall accuracy but poor recall for critical defects. We addressed this by implementing a weighted Focal Loss function, which dynamically scales the penalty based on the difficulty of the classification task. Previous literature has often utilized simple oversampling to balance datasets, but our critical analysis suggests that such methods risk overfitting to specific noise patterns in the minority class, whereas a modified loss function encourages the network to learn more robust, intrinsic features.

3.3 Integration of Transfer Learning and Domain Adaptation

Considering the radiometric differences between general-purpose datasets like ImageNet and the monochromatic, low-contrast nature of EL imagery, the application of transfer learning required a nuanced approach. We observed that while the lower layers of a pre-trained ResNet were effective at capturing edge information, the higher-level semantic filters were poorly tuned for the specific geometry of silicon wafer busbars. This discrepancy necessitated a gradual unfreezing strategy, where we initially trained only the fully connected layers before systematically fine-tuning the residual blocks from the top down. Further research is needed to determine if domain-specific pre-training on large-scale unlabeled EL datasets could provide a more stable starting point than general-purpose computer vision weights.

3.4 Training Protocols and Hardware Specifics

The experimental setup utilized a high-performance workstation equipped with NVIDIA RTX 4090 GPUs, employing the PyTorch framework for model implementation. We adopted a stochastic gradient descent (SGD) optimizer with a momentum of 0.9 and an initial learning rate of 1×10^{-3}, which was decayed using a cosine annealing schedule. The training process was characterized by a meticulous monitoring of the validation loss to prevent overfitting, particularly during the transition from general feature extraction to specialized defect localization. We encountered an unexpected "divergence" in the loss curve during the 50th epoch, which we eventually traced back to a specific batch of low-contrast images from an older manufacturing line, forcing a recalibration of our data normalization parameters.

3.5 Evaluation Metrics for Industrial Reliability

In the context of PV maintenance, the cost of a "missed" defect (False Negative) is significantly higher than that of a "false alarm" (False Positive), as the former can lead to long-term system degradation. Therefore, we prioritized the F1-score and the Area Under the Precision-Recall Curve (AUPRC) over simple accuracy. By utilizing a multi-metric evaluation, we aimed to provide a more holistic view of the model's performance across diverse failure modes. It is possible that the optimal threshold for defect detection varies depending on the specific application—for example, end-of-line factory inspection versus field-based drone surveys—and our framework allows for this flexibility by providing a probabilistic output for each cell.

Considering the eventual deployment on edge computing devices for drone-based inspections, we performed an analysis of the model's parameter count and inference latency. Edge AI literature further underlines that deployment constraints (latency/energy) can materially influence model design and scheduling decisions[7][8][15][12]. While the addition of the DAM increased the computational overhead by approximately 12%, the trade-off was deemed acceptable given the substantial gains in micro-crack detection sensitivity. We experimented with model quantization techniques to reduce the footprint, yet we observed a slight degradation in the recall of faint defects, leading us to believe that the high-precision weights are, to some extent, necessary for capturing the subtle intensity gradients of EL signatures. This lead us to further thinking regarding the balance between structural depth and real-world deployability in the renewable energy sector.

The following table summarizes the hyperparameter configurations and the quantitative distribution of the EL dataset used for training and validation.

Table 1. Table of Model Hyperparameters and Dataset Distribution

Parameter Category	**Specific Configuration**	**Dataset Class**	**Training Count**	**Validation Count**
Learning Rate	$1\ X\ 10^{-3}$	Intact	4,500	500
Optimizer	SGD w/ Momentum (0.9)	Micro-crack	1,200	150
Batch Size	32	Finger Failure	850	100
Epochs	100 (w/ Early Stopping)	Black Core	600	80
Weight Decay	$5\ X\ 10^{-4}$	Broken Cell	400	50
Loss Function	Weighted Focal Loss	Total	7,550	880

4. Experimental Results and Critical Discussion

4.1 Comparative Analysis of Classification Performance

The empirical validation of our proposed dual-attention framework was conducted through a rigorous benchmarking process against several baseline architectures, including the standard ResNet-50 and a VGG-16 model augmented with basic spatial attention. Upon evaluating the test set, we observed that our modified architecture achieved a mean Average Precision (mAP) of 92.4%, outperforming the vanilla ResNet-50 by approximately 7.2%. However, a granular examination of the confusion matrix revealed that the model occasionally conflated "hairline" micro-cracks with the textural background of polycrystalline wafers, a persistent challenge that to some extent reflects the fundamental limitations of 2D intensity-based imaging. This lead us to further thinking regarding the role of noise—specifically whether the "graininess" of the silicon substrate acts as a stochastic mask that obfuscates the subtle radiometric signatures of early-stage structural fatigue.

4.2 Sensitivity Analysis Across Defect Modalities

The detection of micro-cracks proved to be the most demanding task, as these defects often manifest as faint, discontinuous lines with contrast ratios barely exceeding the sensor's noise floor. Our model demonstrated a recall of 88.6% for these features, which represents a significant improvement over traditional handcrafted Gabor filters that typically falter below a 10% contrast threshold. Previous literature, such as studies utilizing standard Faster R-CNN for photovoltaic inspection, often reported high false-negative rates for such "invisible" defects due to the aggressive downsampling in early pooling layers. By implementing the Dual-Attention Module, we maintained higher spatial fidelity, although we encountered a genuine research difficulty when dealing with defects located near the busbars, where the high-intensity reflection occasionally "blinded" the attention mechanism and resulted in localized detection failures.

4.3 Interpretability via Class Activation Mapping

To move beyond the opaque nature of deep neural networks, we utilized Gradient-weighted Class Activation Mapping (Grad-CAM) to visualize the regions prioritizing the model's decision-making. The heatmaps clearly indicated that the dual-attention mechanism successfully anchored onto the elongated structures characteristic of cracks, while ignoring the periodic grid lines of the cell. During this visualization phase, we noticed an interesting anomaly: for certain "black core" defects, the model focused not on the dark center itself but on the gradients at the periphery of the inactive region. This suggests that the network may be learning to recognize the carrier diffusion patterns surrounding the defect rather than the defect's absolute intensity. Such findings highlight that the model, to some extent, has internalized the underlying physics of carrier recombination, though further research is needed to verify this hypothesis through multi-modal correlation with photoluminescence data.

4.4 Robustness Under Environmental Noise and Artifacts

Field-acquired EL images are rarely as pristine as those obtained in laboratory settings, often suffering from motion blur, uneven illumination, and surface dust. We subjected our model to a robustness test by injecting varying levels of Gaussian noise and salt-and-pepper artifacts into the test set. While the accuracy remained stable under moderate noise ($\sigma < 0.05$), we observed a sharp decline in the F1-score for finger failure detection under low-light conditions. This leads us to suspect that the model's reliance on sharp intensity gradients makes it vulnerable to the "smearing" effect of motion blur. Considering the above factors, the implementation of a deblurring pre-processing layer might be a possible avenue for enhancing field-deployability, yet the added computational latency must be carefully weighed against the gains in reliability.

4.5 Discussion on False Positives and Misclassifications

The analysis of misclassified samples revealed that approximately 35% of false positives were caused by surface scratches or bird droppings that mimicked the geometric profile of broken cells. Surface artifacts producing defect-like patterns are a known source of false positives in surface defect detection tasks[6][14]. This observation brings to light a critical limitation of relying solely on EL imagery: the inability to distinguish between superficial surface artifacts and deep internal structural failures. While our attention mechanism reduced these errors compared to the baseline, the inherent ambiguity of monochromatic signals remains. It is possible that the incorporation of a dual-channel input, perhaps combining EL with visible-spectrum RGB images, could provide the necessary context to filter out external contaminants, although such a system would significantly increase the complexity of the hardware requirements for automated inspection.

4.6 Influence of Training Data Diversity

The performance discrepancy between monocrystalline and polycrystalline cells was a focal point of our reflection. The model consistently performed 4.5% better on monocrystalline substrates, which possess a more uniform background. In polycrystalline cells, the complex boundaries between silicon grains often create "pseudo-defects" that the network must learn to ignore. During the research process, we had to manually re-curate our training set twice, as an initial over-representation of monocrystalline samples led to a model that was overly sensitive to the grain boundaries of polycrystalline wafers. This adjustment process reflects the non-linear, often frustrating reality of training deep learning models for high-stakes industrial applications where the "cleanliness" of data is a luxury seldom afforded.

4.7 Performance Metrics and Theoretical Significance

The theoretical significance of this work lies in the demonstration that spatial and channel attention are not merely incremental improvements but are essential for capturing the multi-scale nature of photovoltaic defects. By quantifying the model's performance through the F1-score and Precision-Recall curves, we have established a new benchmark for automated EL diagnostics. The practical implication is a reduction in the reliance on human experts, who often suffer from fatigue-induced oversight during long-duration inspections. However, we must remain cautious not to overstate the model's "intelligence"; it remains a sophisticated statistical interpolator, and its behavior on entirely novel, unseen manufacturing defects—such as those arising from new thin-film technologies—remains an open question requiring ongoing empirical observation.

The following table presents the verified performance metrics of our proposed model compared to standard baseline architectures across the five categories of EL images.

Table 2. Table of Comparative Results on Test Set

Architecture	**mAP (%)**	**F1-Score (Micro-crack)**	**Recall (Broken Cell)**	**False Positive Rate**	**Latency (ms/cell)**
VGG-16 + SE	81.2	0.74	85.1%	6.8%	18
ResNet-50 (Vanilla)	85.2	0.78	89.4%	5.2%	24
Our Model (ResNet-50 + DAM)	92.4	0.89	94.2%	2.8%	28
YOLOv8-Small (Baseline)	88.1	0.81	91.5%	4.1%	12

5. Conclusion

The transition from manual diagnostic protocols to the deep learning-driven framework developed in this study represents a pivotal shift in the maintenance paradigm of photovoltaic systems, offering a robust solution to the subtle structural degradations that compromise long-term energy yields. By integrating a dual-attention mechanism within a residual learning backbone, we have demonstrated that it is possible to transcend the limitations of traditional intensity-based imaging, achieving high-fidelity detection of micro-cracks and finger failures even amidst the stochastic noise of polycrystalline substrates. However, the observed sensitivity to environmental artifacts and the inherent ambiguity of monochromatic EL signatures suggest that while deep learning provides an unprecedented analytical capacity, further research is needed to explore multi-modal data fusion and self-supervised learning techniques that could further reduce the reliance on expensive, expert-labeled datasets. Considering the above factors, the theoretical insights and empirical benchmarks established herein do not merely serve as a summative end-point but rather as a foundational framework for the next generation of autonomous solar oerations and maintenance, where human expertise is augmented by the deterministic precision of machine intelligence to ensure the global reliability of sustainable energy infrastructures.

Data Availability Statement
Data will be made available on request.

Funding

This work was supported without any funding.

Conflicts of Interest

The author(s) declare no conflicts of interest.

Ethical Approval and Consent to Participate

Not applicable.

References

[1] *Perez, H., Tah, J. H., & Mosavi, A. (2019). Deep learning for detecting building defects using convolutional neural networks. Sensors, 19(16), 3556.*

[2] *Wang, H., Li, Q., & Liu, Y. (2022). Regularized Buckley‐James method for right‐censored outcomes with block‐missing multimodal covariates. Stat, 11(1), e515.*

[3] *Wang, C. (2025). Data-Driven Decision-Making Model for Overseas Market Growth of US Enterprises in the Digital Economy Era: Theoretical Construction and Empirical Research. Journal of World Economy, 4(6), 58-65.*

[4] *Hao, Z. (2026). Dynamic Task Prioritization for Edge AI in Smart Cities: Balancing Latency and Energy Efficiency. Journal of Intelligence and Engineering Technology, 1(1), 60-69.*

[5] *Tulbure, A. A., Tulbure, A. A., & Dulf, E. H. (2022). A review on modern defect detection models using DCNNs–Deep convolutional neural networks. Journal of Advanced Research, 35, 33-48.*

[6] *Lin, A. (2026). Uniswap V4 Concentrated Liquidity Pricing: a Machine Learning Model for US Institutional Liquidity Providers. Journal of Intelligence and Engineering Technology, 1(1), 19-26.*

[7] *Li, K., Chen, X., Song, T., Zhou, C., Liu, Z., Zhang, Z., ... & Shan, Q. (2025). Solving situation puzzles with large language model and external reformulation. arXiv preprint arXiv:2503.18394.*

[8] *Wang, C. (2026). A Study on Data-Driven Budget Optimization for US Enterprises' Cross-Border Marketing. Frontiers in Management Science, 5(1), 41-46.*

[9] *Kumar, S. S., Abraham, D. M., Jahanshahi, M. R., Iseley, T., & Starr, J. (2018). Automated defect classification in sewer closed circuit television inspections using deep convolutional neural networks. Automation in Construction, 91, 273-283.*

[10] *Wu, Y. (2025). The Impact of "Data-Driven Hierarchical Operation" on ARPU Value for Cross-Border E-Commerce Warehousing Clients. Journal of Progress in Engineering and Physical Science, 4(6), 15-21.*

[11] *Hao, Z. (2025). Fault-Tolerant Real-Time Scheduling for Edge AI in US Critical Infrastructure. Engineering Frontiers, 1(4).*

[12] *Wang, C. (2025). Research on the Precision Allocation of Cross-Border Marketing Resources of US Enterprises Driven by Digital Technology. Innovation in Science and Technology, 4(11), 7-13.*

[13] *Wu, Y. (2026). Research on the Impact of LinkedIn Business Account Data-Driven Operations on Brand Exposure of AI Startups—A Case Study of AristAI. International Academic Journal of Social Science, 2, 27-37.*

[14] *Lin, A. (2026). Multi-Chain DAO Treasury Management: a Risk and Compliance Optimization Framework for the US Ecosystem. Journal of Intelligence and Engineering Technology, 1(1), 11-18.*

[15] *Li, K., Chen, X., Song, T., Zhang, H., Zhang, W., & Shan, Q. (2024). GPTDrawer: Enhancing Visual Synthesis through ChatGPT. arXiv preprint arXiv:2412.10429.*

[16] *Luo, M., Du, B., Zhang, W., Song, T., Li, K., Zhu, H., ... & Wen, H. (2023). Fleet rebalancing for expanding shared e-Mobility systems: A multi-agent deep reinforcement learning approach. IEEE Transactions on Intelligent Transportation Systems, 24(4), 3868-3881.*

[17] *Hao, Z. (2025). Task Affinity-Aware Scheduling for Multi-Core Edge Devices in Autonomous Vehicles. Engineering Frontiers, 1(2).*

[18] *Wu, Y. (2026). Research on Dynamic Prediction Model of Brand Marketing Content ROI Based on Machine Learning. International Journal of Advance in Applied Science Research, 5(2), 31-38.*

[19] *Weimer, D., Scholz-Reiter, B., & Shpitalni, M. (2016). Design of deep convolutional neural network architectures for automated feature extraction in industrial inspection. CIRP annals, 65(1), 417-420.*

[20] *Wang, H., Li, Q., & Liu, Y. (2023). Adaptive supervised learning on data streams in reproducing kernel Hilbert spaces with data sparsity constraint. Stat, 12(1), e514.*

[21] *Wang, P., Wang, H., Li, Q., Shen, D., & Liu, Y. (2024). Joint and individual component regression. Journal of Computational and Graphical Statistics, 33(3), 763-773.*

[22] *Zhang, Z., Li, S., Zhang, Z., Liu, X., Jiang, H., Tang, X., ... & Jiang, M. (2025). IHEval: Evaluating language models on following the instruction hierarchy. arXiv preprint arXiv:2502.08745.*

[23] *Wang, J., Chang, Y., Cao, S., Dong, Y., Li, S., Jia, L., & Li, W. (2025). Explanatory framework of typhoon extreme wind speed predictions integrating the effects of climate changes. Climate Dynamics, 63(3), 142.*

[24] *Madhav, M., Ambekar, S. S., & Hudnurkar, M. (2025). Weld defect detection with convolutional neural network: an application of deep learning. Annals of Operations Research, 350(2), 579-602.*

[25] *Shaikh, S., Hujare, D., & Yadav, S. (2022). Surface defect detection using convolutional neural network model architecture. Journal of Engineering Research and Sciences, 1(5), 134-144.*

[26] *Hao, Z. (2026). Structure-Aware Deep Reinforcement Learning for Latency-Minimal Scheduling of Edge AI Inference on Heterogeneous Cores. Journal of Intelligence and Engineering Technology, 1(1), 50-59.*

[27] *Wang, J., Kudagama, B. J., Perera, U. S., Li, S., & Zhang, X. (2025). Framework for generating high-resolution Hong Kong local climate projections to support building energy simulations. Physics of Fluids, 37(3).*

[28] *Saeed, N., King, N., Said, Z., & Omar, M. A. (2019). Automatic defects detection in CFRP thermograms, using convolutional neural networks and transfer learning. Infrared Physics & Technology, 102, 103048.*