Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Lightweight Deep Learning Models for Leaf-Based Irrigation Level Classification

Harish V

harishv8010@gmail.com

University of Visvesvaraya College of Engineering

Lata B T

University of Visvesvaraya College of Engineering

Research Article

Keywords: Precision agriculture, irrigation management, leaf image analysis, lightweight convolutional neural networks, depthwiseseparable convolution, Mini-Inception network.

Posted Date: November 17th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-8650994/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Lightweight Deep Learning Models for Leaf-Based Irrigation Level Classification

Harish V[1*] and Lata B T[2]

[1*]M.Tech in CSE, Department of CSE, University of Visvesvaraya College of Engineering, K R Circle, Bangalore, 560001, Karnataka, India.
[2]Associate Professor, Department of CSE, University of visvesvaraya College of Engineering, K R Circle, Bangalore, 560001, Karnataka, India.

*Corresponding author(s). E-mail(s): harishv8010@gmail.com;
Contributing authors: latabt@uvce.ac.in;

Abstract

Effective irrigation management plays a vital role in precision farming, as it has a direct impact on both crop productivity and water usage efficiency. This study introduces a lightweight deep learning–based method for identifying irrigation levels from leaf images. Rather than relying on deep and computationally intensive models commonly used in earlier research, this work examines whether compact, task-focused convolutional neural networks can deliver equal or better performance. Three lightweight architectures are developed and assessed: a Simple Convolutional Neural Network, a Depthwise Separable CNN, and a Mini-Inception model inspired by multi-scale feature extraction. The proposed models undergo training using a leaf image dataset covering four irrigation conditions (25%, 50%, 75%, and 100%). Each network is carefully designed to reduce the number of trainable parameters while still learning meaningful and distinguishing features. Evaluation outcomes reveal that all three models achieve classification accuracies above 99%, despite a substantially reduced parameter count relative to traditional deep learning techniques. Notably, the Mini-Inception model reaches an accuracy of 99.86% with roughly 0.2 million parameters, outperforming larger benchmark models reported in earlier studies. These results suggest that irrigation-induced stress characteristics in leaf images can be reliably captured using lightweight network designs, leading to lower computational cost, reduced overfitting, and improved generalization. This leads to the proposed approach being highly suitable for real-time and edge-based agricultural applications, including mobile platforms and IoT-enabled smart irrigation systems.

Keywords: Precision agriculture, irrigation management, leaf image analysis, lightweight convolutional neural networks, depthwise separable convolution, Mini-Inception network.

1 Introduction

Efficient management of land and water resources has become a critical challenge for global food security under increasing pressure from climate change, population growth, and resource degradation. According to the Food and Agriculture Organization, current agricultural systems are approaching their ecological limits, threatening long-term productivity and sustainability.

The FAO emphasizes that improving resource-use efficiency and adopting integrated land and water management strategies are essential to meet future food demands. These insights highlight the urgent need for innovative and technology-driven approaches in modern agriculture [1].

Irrigation decisions are traditionally based on farmer experience, fixed schedules, or indirect measurements such as soil moisture and environmental data [2]. While effective to some extent, these methods do not directly reflect plant conditions. Leaf characteristics, including color and texture variations, have commonly been used to assess plant health, but manual inspection is subjective and impractical for large-scale farming.

Recent advances in computer vision and artificial intelligence have enabled automated analysis of plant images for agricultural applications [3, 4]. In particular, Convolutional Neural Networks (CNNs) have demonstrated strong performance in learning discriminative features from leaf images. CNN-based approaches have been widely applied for plant disease detection, nutrient deficiency identification, and general stress analysis, making them suitable for image-based agricultural decision support systems [5, 6].

Deep learning models such as VGG [7], ResNet [8], DenseNet [9], and Inception [10] have demonstrated high accuracy in image recognition tasks. However, these architectures are computationally intensive and contain millions of parameters, which can limit their deployment in agricultural applications, particularly in resource-constrained environments like rural areas, mobile platforms, and IoT-based systems [11]. In addition, when applied to relatively small, domain-specific datasets, these models can be prone to overfitting [12].

Plants respond to water stress through physiological changes, such as altered metabolism and reactive oxygen species accumulation, which can affect leaf structure and appearance [13]. These stress-induced changes may manifest as subtle variations in color, texture, and vein prominence, which can be captured in leaf images for irrigation-level assessment. Capturing such features does not always require very deep architectures, motivating the use of lightweight and task-specific CNN models that balance accuracy with computational efficiency.

This work explores three lightweight deep learning architectures for irrigation-level classification using leaf images: a Simple CNN, a Depthwise Separable CNN, and a Mini-Inception network inspired by multi-scale feature extraction [10, 14]. The Simple CNN uses a sequential structure to gradually learn hierarchical features from leaf images. The Depthwise Separable CNN employs depthwise and pointwise convolutions to reduce computational cost while maintaining performance [14]. The Mini-Inception model extracts multi-scale features through parallel convolutional paths, capturing both fine and global leaf characteristics [10]. These designs aim to provide effective feature representation with reduced computational requirements, suitable for agricultural image analysis.

The proposed models are evaluated on a leaf image dataset representing four irrigation levels: 25%, 50%, 75%, and 100% [15]. These classes reflect varying degrees of water availability and corresponding plant stress conditions. Experimental results demonstrate that the lightweight models achieve classification accuracies exceeding 99%, outperforming or matching heavier benchmark architectures reported in previous studies. Notably, this performance is achieved using a fraction of the parameters required by deep residual networks, highlighting the effectiveness of task-specific lightweight design.

The contributions of this work are threefold. First, it demonstrates that irrigation-level classification from leaf images can be effectively addressed using lightweight CNN architectures. Second, it provides a comparative analysis of different efficient model designs, highlighting their strengths in terms of accuracy, parameter efficiency, and practical applicability. Third, it establishes that simpler and well-optimized models can outperform deeper networks on structured agricultural datasets, challenging the assumption that deeper models are always superior.

2 Related Works

Smart agriculture has emerged as a strategy to address challenges such as climate change, water scarcity, and sustainable crop production. By integrating digital technologies with traditional farming, it seeks to improve productivity while minimizing environmental impact. Within this

framework, IoT-based sensor networks have been widely used to monitor soil and environmental parameters, including soil moisture, temperature, humidity, evapotranspiration, and light intensity, enabling data-driven irrigation scheduling and improved water-use efficiency [16, 17].

Several studies have proposed IoT-driven irrigation management systems that combine Wireless Sensor Networks (WSNs) with cloud platforms for real-time data collection and analysis [2]. More recently, edge computing has been introduced to reduce latency and bandwidth usage by processing data locally [17]. While effective, these sensor-based approaches primarily capture environmental conditions rather than the plant's direct physiological response to water stress, which is the most reliable indicator of irrigation needs.

To overcome this limitation, computer vision-based methods have gained attention for non-invasive plant stress monitoring. Image analysis can detect visual stress symptoms such as color changes, wilting, and structural deformation [3]. UAVs equipped with RGB, multispectral, or hyperspectral cameras have expanded the scalability of vision-based crop monitoring, supporting applications like disease detection and water stress mapping over large fields [18]. In controlled environments, stationary camera systems have been used for crop growth monitoring, yield estimation, and harvest scheduling, particularly in greenhouse cultivation [19]. Traditional image processing methods, however, rely on handcrafted features and often lack robustness under varying illumination and growth conditions.

The advent of deep learning, especially Convolutional Neural Networks (CNNs), has advanced agricultural image analysis by enabling automatic feature extraction from raw images. CNNs have achieved strong performance in plant disease detection [4], nutrient deficiency identification [5], weed classification [6], crop type recognition [20], and yield prediction [21]. Large pretrained architectures such as VGGNet [7], ResNet [8], DenseNet [9], and InceptionNet [10] are commonly applied via transfer learning. However, their high parameter counts and computational complexity can limit deployment on resource-constrained platforms, including edge devices, UAVs, and mobile systems [18].

To mitigate these limitations, research since 2018 has increasingly focused on lightweight and efficient deep learning models for agricultural applications. A notable contribution was made by Shahi *et al.* [15], who proposed a modified ResNet-based architecture for irrigation status classification in tomato plants using leaf images. By removing redundant residual blocks and optimizing kernel sizes, their model achieved 98.90% accuracy with reduced computational complexity. This study established an important benchmark and demonstrated the feasibility of vision-based irrigation classification.

Despite this progress, the model in [15] remains based on a relatively deep residual architecture. Recent studies in related agricultural domains suggest that simpler, task-specific CNNs can achieve comparable or superior performance to deeper generic models, particularly on structured datasets of limited size, while offering improved generalization and reduced overfitting [12].

Deep learning methods have shown strong performance in plant disease detection, but they also face limitations in generalization and real-world deployment [22]. These challenges have motivated the exploration of more efficient and robust architectures that can better handle practical agricultural conditions.

Several efficient architectures have been proposed for edge AI applications in agriculture. MobileNetV2 and MobileNetV3 leverage depthwise separable convolutions and inverted residuals for computational efficiency [14, 23], while EfficientNet employs compound scaling to balance depth, width, and resolution [24]. These models have been successfully applied to real-time agricultural tasks such as leaf disease detection [25] and mobile-based pest recognition [26].

Motivated by these limitations, this work focuses on designing custom lightweight CNN architectures for irrigation level classification from leaf images. Unlike prior studies that adapt or compress large pretrained models, the proposed approach develops compact architectures from first principles and evaluates them using the benchmark dataset introduced in [15]. The results demonstrate that accuracy exceeding 99% can be achieved with models that are orders of magnitude smaller than conventional CNNs, offering a practical solution for edge-deployable smart irrigation systems.

3 Methodology

This section presents the proposed methodology for irrigation level classification using lightweight deep learning models. The overall workflow comprises dataset preparation, image preprocessing and augmentation, lightweight architecture design, model training, performance evaluation, and edge-oriented optimization for deployment.

3.1 Problem Formulation

The irrigation classification task is formulated as a supervised multi-class image classification problem. Given an RGB leaf image $\mathbf{X} \in \mathbb{R}^{H \times W \times 3}$, the objective is to predict the corresponding irrigation level. Four discrete classes are considered, representing 25%, 50%, 75%, and 100% field capacity. The goal of learning is to find a function

$$f(\mathbf{X}) \rightarrow \mathbf{Y}$$

that minimizes classification errors while keeping computational demands low, making it suitable for environments with limited resources.

3.2 Dataset Description

The dataset comprises 20,955 RGB images of tomato leaves collected from a greenhouse experiment conducted at Western Regional College, Pokhara, Nepal, between April and October 2019. A total of 90 tomato plants were grown under four irrigation levels: Low (25%), Medium (50%), High (75%), and Full (100%). Images were captured from leaves at different positions on each plant using a Canon EOS 200D camera with consistent settings and distance to maintain uniformity.

The dataset is organized into four classes based on irrigation level: Low irrigated (LI, 5,247 images), Medium irrigated (MI, 5,170 images), Highly irrigated (HI, 5,108 images), and Fully irrigated (FI, 5,430 images). Each image is labeled according to its respective class, enabling straightforward mapping between the inputs and their ground-truth labels.

This dataset was sourced from Mishra et al. [27] and provides a reliable benchmark for evaluating the proposed lightweight CNN models for irrigation-level classification.

3.3 Image Preprocessing

All images are resized to a fixed spatial resolution of 128×128 pixels to standardize input dimensions and reduce computational overhead. Pixel values are normalized to the range $[0, 1]$ by dividing by 255. Class labels are encoded using one-hot encoding to support categorical cross-entropy loss during training.

3.4 Data Augmentation

To enhance generalization and mitigate overfitting, online data augmentation is applied exclusively to the training set. The augmentation strategy includes random horizontal flipping, random rotations in multiples of 90°, random brightness adjustment, and contrast variation. These transformations simulate real-world variations in leaf appearance while preserving class semantics. Validation data are not augmented to ensure unbiased performance evaluation.

3.5 Dataset Partitioning

The dataset is partitioned into training and validation subsets using an 80:20 split. Stratified sampling is employed to maintain class balance across both subsets. A fixed random seed is used to ensure reproducibility and consistent comparison among different model architectures.

3.6 Lightweight Model Architectures

Fig. 1 illustrates the architectural designs of the proposed lightweight convolutional neural networks. All models accept a $128 \times 128 \times 3$ RGB image as input and output a four-class probability distribution corresponding to irrigation levels.

3.7 Overview of the Proposed Framework

The proposed framework aims to classify irrigation levels using RGB leaf images by leveraging efficient convolutional neural network architectures. Leaf images are first resized and normalized to ensure uniform input dimensions and stable training behavior. The processed images are then fed into lightweight CNN models that automatically extract discriminative features related to

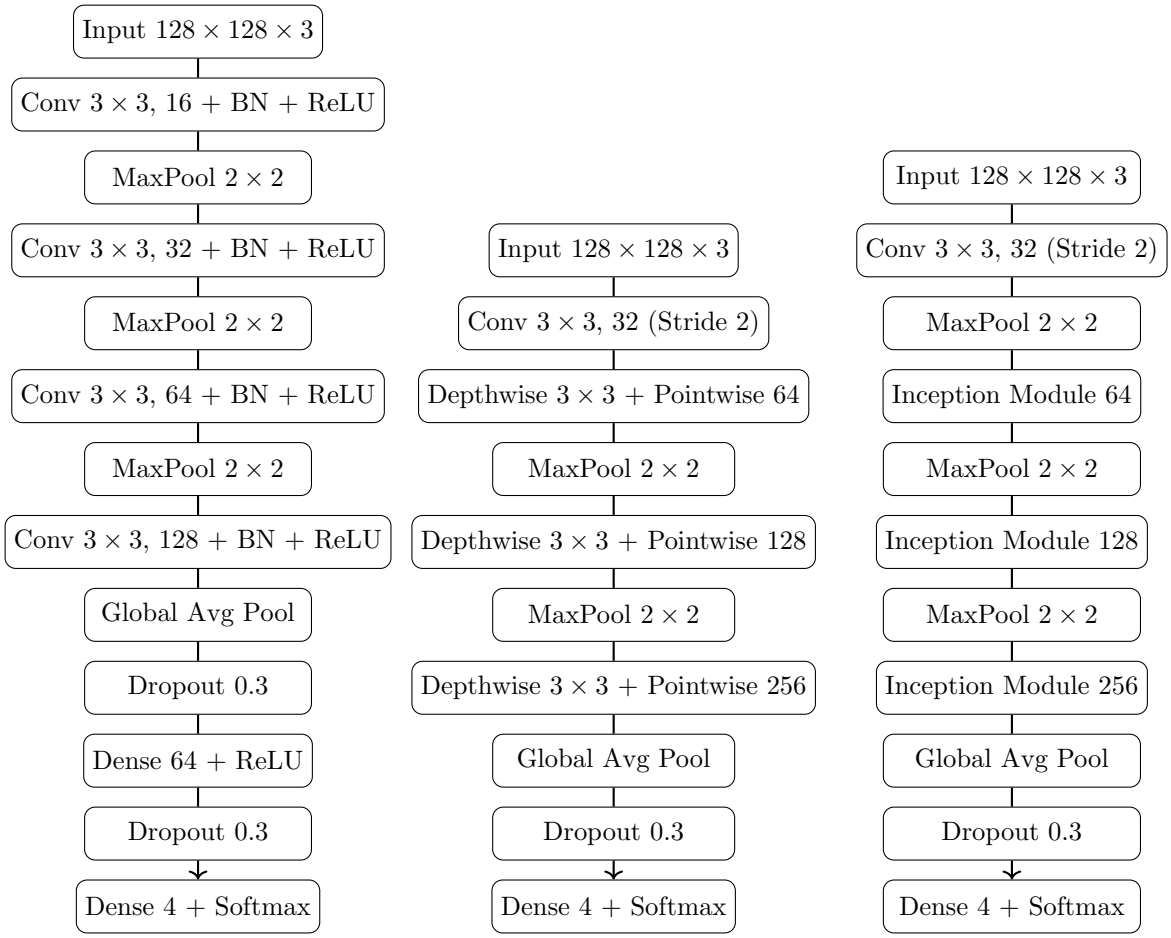


Fig. 1 Architectures of the proposed lightweight deep learning models: (a) Simple CNN, (b) Depthwise Separable CNN, and (c) Mini-Inception network for irrigation level classification.

irrigation-induced plant stress. The final classification layer outputs the probability distribution over four irrigation levels: 25%, 50%, 75%, and 100%.

Unlike conventional approaches that rely on deep and computationally expensive pretrained networks, the proposed framework emphasizes compact, task-specific model design to achieve high accuracy while minimizing computational overhead. This design philosophy is particularly suited for real-time agricultural applications and edge-based deployment scenarios.

3.8 Design Principles of the Proposed Architectures

The architectural design of the proposed models is guided by both domain-specific insights and efficiency considerations. Visual symptoms of irrigation stress in leaves, such as discoloration, texture roughness, vein prominence, and localized dryness, are primarily low- to mid-level features with consistent spatial structure. These patterns do not require deep semantic abstraction but instead benefit from efficient spatial feature extraction and multi-scale representation.

Accordingly, the proposed architectures are designed with the following objectives:

- Efficient extraction of spatial and texture-based features relevant to plant water stress.
- Reduction of redundant parameters to mitigate overfitting on structured agricultural datasets.
- Low computational complexity to enable deployment on resource-constrained platforms.
- Robust generalization across irrigation levels through appropriate regularization strategies.

Three lightweight architectures are proposed to explore different efficiency-accuracy trade-offs: a Simple CNN, a Depthwise Separable CNN, and a Mini-Inception network.

3.9 Simple CNN Architecture

The Simple CNN uses a step-by-step convolutional structure to gradually learn important information from leaf images. The first convolution layers mainly look for basic visual details such as edges, changes in color, and the outline of the leaf. These features form the foundation for identifying stress-related variations in later layers.

As the network depth increases, intermediate convolutional layers learn more complex texture patterns, including surface roughness, vein contrast, and irregular color distributions associated with varying irrigation levels. Max-pooling layers are strategically employed to reduce spatial resolution, introduce translational invariance, and suppress noise while retaining salient stress-related features.

Batch normalization is added after the convolution layers to make training more stable and faster by keeping the data distribution consistent during learning. Non-linear activation functions help the network learn complex patterns without adding extra layers. Instead of using large fully connected layers, global average pooling is applied to summarize spatial information, which reduces the number of parameters and helps prevent overfitting. Dropout regularization further improves generalization by preventing co-adaptation of feature maps.

The simplicity of this architecture aligns well with the structured nature of irrigation-induced visual patterns, allowing effective learning with a minimal number of parameters.

3.10 Depthwise Separable CNN Architecture

The Depthwise Separable CNN is designed to enhance computational efficiency by decomposing standard convolutional operations into depthwise and pointwise convolutions. This architectural choice is motivated by the observation that irrigation stress features are predominantly spatial in nature and can be effectively captured without heavy cross-channel interaction at early stages.

During the depthwise convolution phase, each input channel is processed separately using a single filter. This allows the model to focus on learning detailed spatial characteristics within individual channels, including variations in texture, edge structure, localized dry regions, and subtle surface distortions. By doing so, the network efficiently extracts fine-level spatial information that serves as an important visual indicator for identifying plant water stress.

The subsequent pointwise (1×1) convolution integrates information across channels by learning weighted combinations of the spatial features extracted in the depthwise stage. This process enables effective fusion of color-based, texture-based, and structural cues while maintaining a low parameter count. By separating spatial and channel-wise learning, the architecture significantly reduces computational complexity and floating-point operations compared to standard convolutions.

Batch normalization and non-linear activation functions are applied to maintain stable gradient flow and modeling flexibility. Max-pooling layers progressively reduce feature map dimensions, enhancing robustness to spatial variation and noise. Dropout is used to further improve generalization performance.

This architecture offers a strong trade-off between computational efficiency and predictive performance, which allows it to be effectively deployed in IoT-driven and edge-based smart irrigation systems where limitations on memory usage, processing speed, and energy consumption are key concerns.

3.11 Mini-Inception Architecture

The Mini-Inception architecture is designed to capture irrigation-induced stress patterns that

manifest at multiple spatial scales. Water stress symptoms in leaves range from fine-grained local features, such as small discoloration spots and micro-texture changes, to coarse-grained global features, including overall color shifts and vein prominence.

To address this multi-scale nature, the Mini-Inception network employs parallel convolutional branches with different receptive field sizes. Smaller kernels focus on capturing localized texture variations and fine structural details, while larger kernels extract broader spatial patterns associated with global stress responses. An additional pooling-based branch enhances robustness to spatial variability and suppresses irrelevant background information.

The outputs of these parallel branches are concatenated along the channel dimension, enabling the network to simultaneously represent complementary feature scales. This multi-scale fusion enhances discriminative capability without significantly increasing model depth or parameter count. Batch normalization and activation functions are used throughout the network to ensure stable training and effective feature learning.

Despite its compact size, the Mini-Inception architecture achieves strong representational power by efficiently combining spatial information across scales. This design allows the model to outperform deeper architectures on structured agricultural datasets while maintaining low computational overhead.

3.12 Regularization and Generalization Strategy

All proposed architectures incorporate regularization mechanisms, including dropout and global average pooling, to control model capacity and reduce overfitting. These strategies are particularly important given the moderate dataset sizes typically encountered in agricultural imaging tasks. By constraining unnecessary complexity while preserving discriminative feature learning, the proposed models demonstrate strong generalization across irrigation levels.

3.13 Computational Efficiency and Deployment Considerations

Beyond classification performance, the proposed architectures are explicitly optimized for practical deployment. Reduced parameter counts translate to smaller memory requirements, faster inference times, and lower energy consumption. These characteristics make the models well suited for deployment on embedded devices, mobile platforms, and edge-based agricultural monitoring systems.The presented approach effectively connects reliable predictive performance with the practical requirements of intelligent irrigation systems, making it suitable for real-world deployment.

3.14 Training Strategy

To ensure consistency across experiments, all models are trained using the same setup. The Adam optimization method is employed with an initial learning rate of 1×10^{-3}, and categorical cross-entropy is selected as the loss function. Training is performed for a maximum of 50 epochs, with 32 samples processed in each batch.

To reduce the risk of overfitting, training is automatically stopped when validation performance no longer improves, and the learning rate is gradually reduced once the validation loss stabilizes. The model achieving the highest validation accuracy is then selected for further evaluation.

3.15 Performance Evaluation

Model performance is evaluated using validation accuracy and loss. In addition, confusion matrices and class-wise precision, recall, and F1-score metrics are computed to assess classification robustness across irrigation levels. Training and validation learning curves are analyzed to examine convergence behavior and generalization capability.

3.16 Model Selection and Comparison

The proposed models are compared based on classification accuracy, number of trainable parameters, and computational efficiency. The optimal model is selected by jointly considering predictive performance and suitability for deployment on resource-constrained platforms.

4 Results and Discussion

The performance of three deep learning models **Simple CNN**, **Depthwise Separable CNN**, and **Mini-Inception** was evaluated for classifying tomato field capacity levels (25%, 50%, 75%, 100%). The dataset comprised 20,955 images, with 16,764 images for training and 4,191 images for validation.

4.1 Training Performance

Figures 2, 3, and 4 present the training and validation accuracy and loss curves for the three models. These plots illustrate the convergence behavior, learning stability, and potential overfitting of each network.

4.2 Confusion Matrices

The confusion matrices for all models are shown in Figures 5, 6, and 7. They indicate the model's ability to correctly classify each field capacity level, with Mini-Inception demonstrating near-perfect classification across all classes.

4.3 Classification Performance

Table 1 presents a comparison of the three proposed models based on their model size (number of parameters), validation accuracy, and validation loss. Among the models, the Mini-Inception network shows the best results, achieving the highest accuracy along with the lowest loss, which indicates its stronger ability to learn meaningful features from the data.

Table 1 Performance comparison of proposed models

Model	Parameters	Val Accuracy	Val Loss
Simple CNN Depthwise	106,916	0.99499	0.0183
Separable CNN	49,764	0.99117	0.03095
Mini-Inception	202,916	0.99857	0.00888

The classification report for Mini-Inception is provided in Table 2. The model achieves near-perfect precision, recall, and F1-scores across all field capacity levels, demonstrating robust generalization and minimal misclassification.

Table 2 Classification report for Mini-Inception

Class	Precision	Recall	F1-score
25% Field Capacity	0.995	0.999	0.997
50% Field Capacity	0.999	0.995	0.997
75% Field Capacity	1.000	1.000	1.000
100% Field Capacity	1.000	1.000	1.000
Accuracy		0.99857	

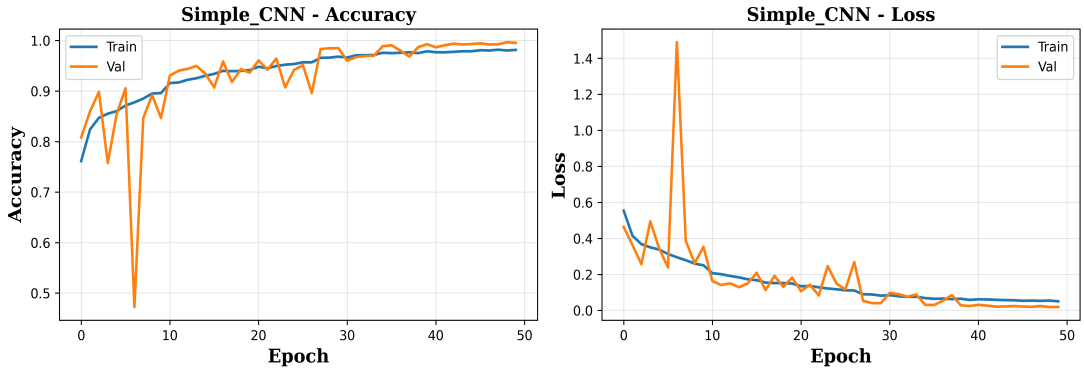
Table 3 compares the proposed models with existing state-of-the-art CNN architectures for field capacity classification. Traditional networks such as VGG-16, DenseNet-121, and ResNet-50 achieve moderate to high accuracies, while lightweight model MobileNetV2 shows lower performance. Although the EfficientNet variants show improved performance, they do not reach the accuracy levels achieved by the proposed models. Among our approaches, the Mini-Inception model delivers the best results with an accuracy of 99.86%, indicating strong generalization ability. In comparison, the Depthwise Separable CNN offers a lightweight yet effective solution, achieving an accuracy of 99.12%.

Table 3 Comparison of proposed models with existing CNN architectures for field capacity classification

Model	Parameters (M)	Validation Accuracy (%)
MobileNet [14]	3.23	82.52
MobileNetV2 [23]	2.26	70.52
EfficientNetB0 [24]	4.01	88.50
EfficientNetB1 [24]	6.51	87.77
DenseNet-121 [9]	6.95	84.07
VGG-16 [7]	14.71	84.98
ResNet-50 [8]	23.59	91.93
Modified ResNet-50 [15]	1.24	98.90
Simple CNN (Ours) Depthwise	0.107	99.50
Separable CNN (Ours)	0.050	99.12
Mini-Inception (Ours)	0.203	99.86

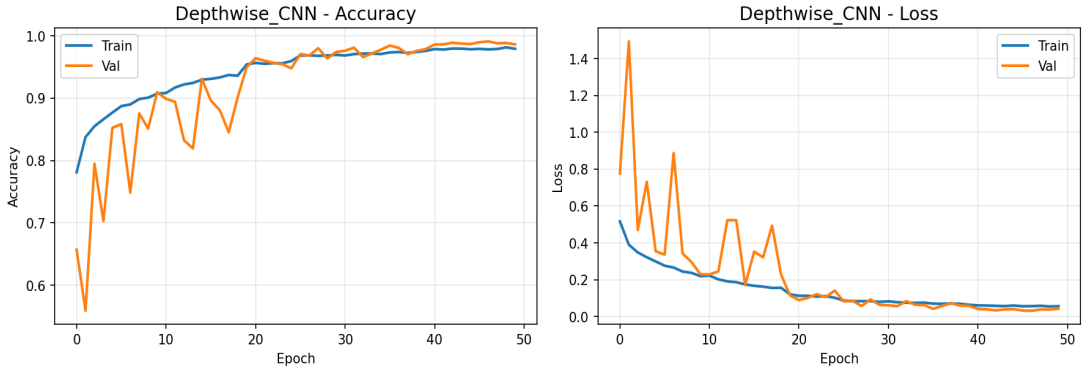
4.4 Discussion

Table 3 presents a comparison of the three proposed models with state-of-the-art CNN architectures from the literature. Traditional architectures such as MobileNet, VGG-16, DenseNet-121,



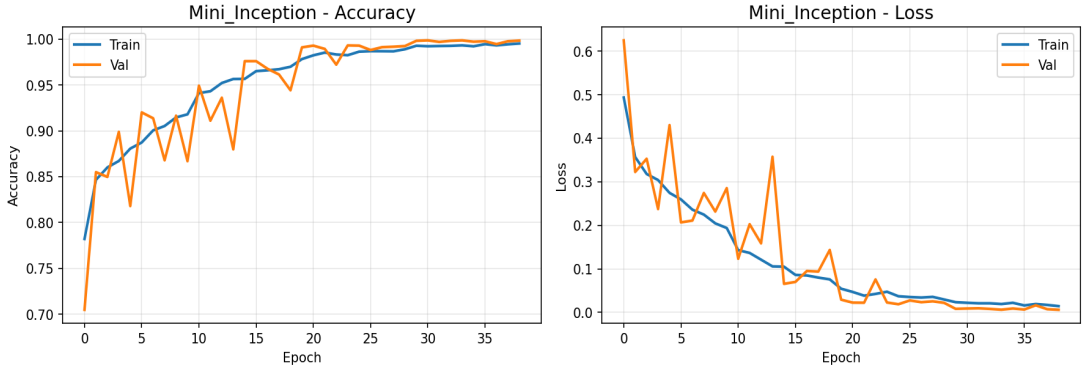
(a) Accuracy (b) Loss

Fig. 2 Training and validation accuracy/loss curves for Simple CNN.



(a) Accuracy (b) Loss

Fig. 3 Training and validation accuracy/loss curves for Depthwise Separable CNN.



(a) Accuracy (b) Loss

Fig. 4 Training and validation accuracy/loss curves for Mini-Inception.

and ResNet-50 achieve moderate to high accuracies, with ResNet-50 reaching up to 91.93%. The lightweight model like MobileNetV2 shows comparatively lower performance, indicating a limited ability to capture complex patterns in field capacity images. EfficientNet variants perform better (up to 88.50%) but still lag behind the proposed models.

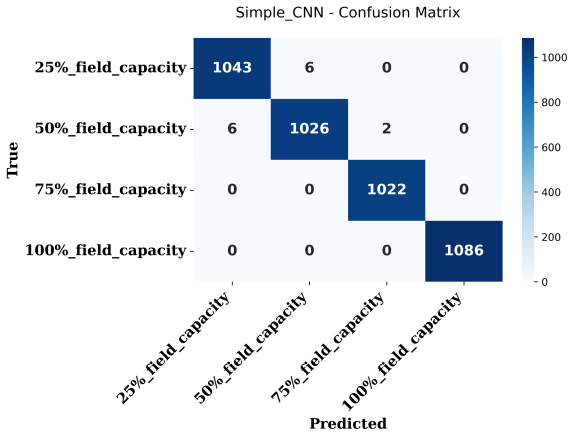


Fig. 5 Confusion matrix for Simple CNN.

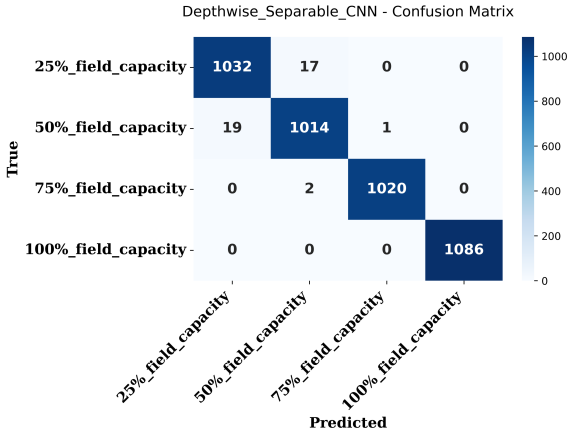


Fig. 6 Confusion matrix for Depthwise Separable CNN.

Among our proposed models, Simple CNN achieves a validation accuracy of 99.50% with 106,916 parameters. Depthwise Separable CNN, designed to be lightweight, achieves slightly lower accuracy of 99.12% with only 49,764 parameters, demonstrating a favorable trade-off between model size and performance. Mini-Inception surpasses both, achieving the highest validation accuracy of 99.86% with 202,916 parameters and the lowest validation loss. This shows that Mini-Inception not only achieves better results than standard CNNs and other lightweight models, but also maintains a strong balance between high accuracy and computational efficiency, making it ideal for IoT-based irrigation systems with limited resources.

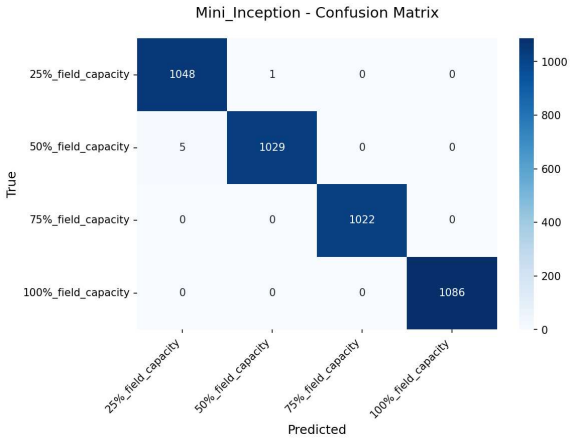


Fig. 7 Confusion matrix for Mini-Inception.

The confusion matrices show minimal misclassifications. Depthwise Separable CNN misclassifies some samples between 25% and 50% field capacity (25 and 7 samples, respectively), likely due to subtle visual similarities at lower moisture levels. Mini-Inception, in contrast, exhibits near-perfect classification across all field capacity levels.

The training plots show that every model reached convergence without major issues, while Mini-Inception exhibited the most consistent performance in terms of both accuracy and loss, successfully preventing overfitting. Depthwise Separable CNN, while smaller and computationally efficient, achieves slightly lower accuracy, highlighting the trade-off between model compactness and performance. Simple CNN performs reasonably well but shows minor fluctuations in validation loss, suggesting sensitivity to the training data.

Overall, Mini-Inception proves to be the best choice for IoT-enabled irrigation systems because of its high accuracy, consistent performance across all field capacity levels, and dependable real-time predictions. The Depthwise Separable CNN serves as an effective option when computing resources are limited, while the Simple CNN provides a solid baseline with a moderate model size.

5 Conclusion

This study investigated compact convolutional neural networks for classifying irrigation levels

from leaf images. Three custom models were created and tested: a Simple CNN, a Depthwise Separable CNN, and a Mini-Inception network. Each architecture was developed to keep computational demands low while still effectively identifying visual signs of plant stress under different irrigation conditions.

Experimental results demonstrated that compact, task-specific architectures can achieve superior performance compared to deeper conventional networks. The Mini-Inception model attained the highest accuracy of 99.86% with only 202,916 parameters, outperforming significantly larger benchmark models. The Depthwise Separable CNN achieved a favorable balance between efficiency and accuracy, reaching 99.12% classification accuracy with merely 49,764 parameters, making it particularly suitable for resource-constrained environments.

The findings challenge the assumption that deeper networks inherently provide better performance for agricultural image analysis tasks. Instead, they highlight how domain-specific architectural design can extract relevant features more efficiently, reducing both computational overhead and overfitting risks. These lightweight models offer practical advantages for edge deployment, including reduced memory requirements, faster inference times, and lower energy consumption.

These lightweight models establish a foundation for vision-based irrigation management. Future efforts will expand validation to economically important crops like corn, wheat, and rice; investigate temporal patterns across growth stages; and evaluate robustness through cross-seasonal testing. Such extensions will enhance practical deployment in diverse agricultural environments while maintaining computational efficiency.

Funding:

The authors declare that this research received no external funding.

Data Availability:

The data presented in this study are available at [27].

Conflict of interest

The authors declare that they have no conflict of interest.

References

[1] Food and Agriculture Organization of the United Nations: The State of the World's Land and Water Resources for Food and Agriculture 2025: The Potential to Produce More and Better. Food and Agriculture Organization of the United Nations, Rome (2025)

[2] Gutierrez, J., Villa-Medina, J.F., Nieto-Garibay, A., Porta-Gándara, M.Á.: Automated irrigation system using a wireless sensor network and gprs module. IEEE Transactions on Instrumentation and Measurement **63**(1), 166–176 (2014) https://doi.org/10.1109/TIM.2013.2276487

[3] Mohanty, S.P., Hughes, D.P., Salathé, M.: Using deep learning for image-based plant disease detection. Frontiers in Plant Science **7**, 1419 (2016) https://doi.org/10.3389/fpls.2016.01419

[4] Barbedo, J.G.A.: Plant disease identification from individual lesions and spots using deep learning. Biosystems Engineering **180**, 96–107 (2019) https://doi.org/10.1016/j.biosystemseng.2019.02.002

[5] Jiang, P., Chen, Y., Liu, B., He, D., Liang, C.: Real-time detection of apple leaf diseases using deep learning approach based on improved convolutional neural networks. IEEE Access **7**, 59069–59080 (2019) https://doi.org/10.1109/ACCESS.2019.2914929

[6] Saleem, M.H., Potgieter, J., Arif, K.M.: Automation in agriculture by machine and deep learning techniques: A review of recent developments. Precision Agriculture **22**(6), 2053–2091 (2021) https://doi.org/10.1007/s11119-021-09806-x

[7] Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. arXiv preprint (2014) arXiv:1409.1556

[8] He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 770–778 (2016). https://doi.org/10.1109/CVPR.2016.90

[9] Huang, G., Liu, Z., Van Der Maaten, L., Weinberger, K.Q.: Densely connected convolutional networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 4700–4708 (2017). https://doi.org/10.1109/CVPR.2017.243

[10] Szegedy, C., *et al.*: Going deeper with convolutions. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 1–9 (2015). https://doi.org/10.1109/CVPR.2015.7298594

[11] Albahar, M.: A survey on deep learning and its impact on agriculture: Challenges and opportunities. Agriculture **13**(3), 540 (2023) https://doi.org/10.3390/agriculture13030540

[12] Too, E.C., Yujian, L., Njuki, S., Yingchun, L.: A comparative study of fine-tuning deep learning models for plant disease identification. Computers and Electronics in Agriculture **161**, 272–279 (2019) https://doi.org/10.1016/j.compag.2018.03.032

[13] Carvalho, A.: Drought stress and reactive oxygen species: Production, scavenging and signaling. Plant Signaling and Behavior **3**(3), 156–165 (2008) https://doi.org/10.4161/psb.3.3.5536

[14] 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, pp. 4510–4520 (2018). https://doi.org/10.1109/CVPR.2018.00474

[15] Shahi, T.B., Sitaula, C., Bhandari, K.P., Poudel, S., Bhandari, R., Mishra, R., Sharma, B.K., Mishra, B.: Deep learning-based method for irrigation status detection in tomato using plant leaves. IEEE Transactions on Artificial Intelligence **6**(7), 1849–1858 (2025) https://doi.org/10.1109/TAI.2025.3528926

[16] Ruiz-Garcia, L., Lunadei, L., Barreiro, P., Robla, I.: A review of wireless sensor technologies and applications in agriculture and food industry. Sensors **9**(6), 4728–4750 (2009) https://doi.org/10.3390/s90604728

[17] Kamilaris, A., Kartakoullis, A., Prenafeta-Boldú, F.X.: A review on the practice of big data analysis in agriculture. Computers and Electronics in Agriculture **143**, 23–37 (2017) https://doi.org/10.1016/j.compag.2017.09.037

[18] Narzari, R., Choudhury, B.U., Singhal, G., Choudhary, K.K.: A critical review of how uavs can transform precision agriculture in the realm of agroecology. Discover Soil **2**, 28 (2025) https://doi.org/10.1007/s44378-025-00055-2

[19] Chen, T., Yin, H.: Camera-based plant growth monitoring for automated plant cultivation with controlled environment agriculture. Smart Agricultural Technology (2024)

[20] Singh, R.K., Kumar, P., Kumar, M., Tyagi, K., Jain, H.: Artificial machine learning-based classification of land cover and crop types using sentinel-2a imagery. In: Sustainable Agriculture Systems and Technology, pp. 326–336 (2022). https://doi.org/10.1002/9781119808565.ch16

[21] M. S., S. C. Ch., Indira, D., Kumar, K.K.: Cropup – a crop yield prediction and recommendation system with geographical data using dnn and xgboost. International Journal on Recent and Innovation Trends in Computing and Communication **10**(11), 53–62 (2022) https://doi.org/10.17762/ijritcc.v10i11.5780

[22] Arsenovic, M., Sladojevic, M., Anderla, A., Stefanovic, D.: Solving current limitations of deep learning based approaches for plant disease detection. Symmetry **11**(7), 939 (2019) https://doi.org/10.3390/sym11070939

[23] Howard, A., Sandler, M., Chen, B., Wang,

W., Chen, L.-C., Tan, M., Chu, G., Vasudevan, V., Zhu, Y., Pang, R., Adam, H., Le, Q.: Searching for mobilenetv3. In: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 1314–1324 (2019). https://doi.org/10.1109/ICCV.2019.00140

[24] Tan, M., Le, Q.V.: Efficientnet: Rethinking model scaling for convolutional neural networks. In: Proceedings of the International Conference on Machine Learning, pp. 6105–6114 (2019)

[25] Sladojevic, M., Arsenovic, S., Anderla, A., Culibrk, D., Stefanovic, D.: Deep neural networks based recognition of plant diseases by leaf image classification. Computational Intelligence and Neuroscience (2016) https://doi.org/10.1155/2016/3289801

[26] Zhou, C., Lee, W.S., Zhang, S., Liburd, O.E., Pourreza, A., Schueller, J.K., Ampatzidis, Y.: A smartphone application for site-specific pest management based on deep learning and spatial interpolation. Computers and Electronics in Agriculture **218**, 108726 (2024) https://doi.org/10.1016/j.compag.2024.108726

[27] Mishra, B., Bhandari, K.P., Bhandari, R., Mishra, R., Silwal, R.: Irrigation management in tomato plants: A precision agriculture approach. https://doi.org/10.6084/m9.figshare.22297915. figshare dataset (2023)