Computers and Electronics in Agriculture 250 (2026) 111960

Contents lists available at ScienceDirect

Computers and Electronics in Agriculture

journal homepage: www.elsevier.com/locate/compag

Original papers

Haar wavelet-guided dynamic Feature Pyramid Network for an efficient paddy leaf disease classification

Kiriharan Tharmini [a,b], Thanikasalam Kokul [c,*], Amirthalingam Ramanan [c], Terensan Suvanthini [d], Subha Fernando [a]

[a] Department of Computational Mathematics, University of Moratuwa, Sri Lanka
[b] Department of Computer Science, Eastern University, Sri Lanka
[c] Department of Computer Science, University of Jaffna, Sri Lanka
[d] Department of Agricultural Biology, University of Jaffna, Sri Lanka

ARTICLE INFO

Dataset link: https://doi.org/10.5281/zenodo.18334503

Keywords:
Paddy leaf disease classification
Feature pyramid network
Knowledge distillation
Haar wavelet-guided deep learning

ABSTRACT

Early and accurate identification of paddy leaf diseases is crucial for timely intervention and effective crop management. Although deep learning based approaches have demonstrated strong performance in paddy leaf disease classification, they often struggle to capture the diverse multiscale symptom patterns of diseases, resulting in the misclassification of challenging samples. Moreover, the high computational cost of deep learning based models remains a significant limitation. To address these challenges, this study proposes a novel Haar wavelet-guided deep Feature Pyramid Network (FPN) model. In the proposed framework, deep features are extracted using a backbone network, and a FPN is employed to capture multiscale disease symptoms. In parallel, Haar wavelet features are extracted to provide complementary texture cues that are not explicitly modelled by deep features. Guided by these wavelet features, a gating-based dynamic feature selection module is then utilised to identify image-specific FPN features, which are subsequently used for classification. In the second phase of the study, a lightweight variant of the proposed model is developed using a knowledge distillation technique, achieving high classification accuracy while substantially reducing the number of parameters and FLOPs. In addition, a new benchmark dataset, named NP-LankaPaddy, is constructed to address the limited availability of datasets representing Sri Lankan paddy leaf disease conditions. Experimental results demonstrate that both the proposed approach and its lightweight variant outperform state-of-the-art methods across four benchmark datasets, achieving notably high accuracies of 98.04% and 97.84%, respectively, on the well known Paddy Doctor dataset. The source code and dataset supporting this study are openly available at https://doi.org/10.5281/zenodo.18334503.

1. Introduction

Rice, the staple food for more than half of the world's population, plays a vital role in global food security, with Asia remaining the dominant region for both its production and consumption (Food and Agriculture Organization of the United Nations (FAO), 2025). However, paddy diseases pose a significant threat to agricultural sustainability by reducing both yield quantity and grain quality. The occurrence and severity of these diseases are influenced by multiple factors, including climatic conditions, water availability, nutrient management, humidity, and farming practices (Ranasinghe et al., 2022). In addition, environmental stresses, particularly temperature variations such as cold stress during early growth stages, further complicate disease manifestation by affecting germination and seedling development (Pan et al., 2023; Xie et al., 2025). Consequently, accurate and timely disease identification is essential for effective crop management. Despite this importance, current practices still rely heavily on manual visual inspection by experts, which is time-consuming, subjective, and prone to inconsistency, thereby limiting timely intervention and large-scale monitoring.

Sri Lanka represents a critical case study, where paddy cultivation occupies approximately 29% of total cultivated land and supports millions of farming households (Department of Agriculture, Sri Lanka, 2024). Paddy farming has a long history in the country and is organised around the *Maha* and *Yala* cultivation seasons. Common diseases such as rice blast, brown spot, bacterial leaf streak, and sheath blight significantly affect production (Department of Agriculture, Sri Lanka, 2025).

* Corresponding author.
E-mail addresses: tharminik@esn.ac.lk (K. Tharmini), kokul@univ.jfn.ac.lk (T. Kokul), a.ramanan@univ.jfn.ac.lk (A. Ramanan), suvanthinit@univ.jfn.ac.lk (T. Suvanthini), subhaf@uom.lk (S. Fernando).

https://doi.org/10.1016/j.compag.2026.111960
Received 21 February 2026; Received in revised form 24 May 2026; Accepted 25 May 2026

Although numerous automated disease recognition systems have been developed in other regions, their direct applicability to Sri Lanka remains limited due to region-specific climatic conditions, environmental variability, and differences in disease manifestation (Yugander et al., 2015). Furthermore, real-world field conditions introduce additional challenges, including complex backgrounds, varying illumination, occlusions, and noise. Early-stage disease symptoms often appear as subtle variations in colour, texture, and intensity, making accurate detection particularly challenging.

Over the past decade, a wide range of automated computational methods have been developed for detecting paddy leaf diseases, as summarised in Table 1. Early approaches predominantly relied on handcrafted texture and statistical descriptors (Pothen and Pai, 2020; Singh et al., 2023; Kaviya and Selvakumar, 2023). In addition, features based on colour, shape, and lesion geometry have been explored (Chawathe and Sudarshan, 2020), often in conjunction with segmentation techniques to isolate regions of interest from complex backgrounds (Dubey and Choubey, 2025; Bharanidharan et al., 2023). Although these methods are computationally efficient and interpretable, they depend heavily on manual feature engineering and exhibit limited generalisation across varying field conditions.

With the advancement of deep learning, convolutional neural networks (CNNs) have emerged as the dominant approach for paddy disease classification due to their ability to automatically learn hierarchical feature representations from raw images (Mahadevan et al., 2024; Ritharson et al., 2024; Krishnamoorthy et al., 2021; Prottasha et al., 2022). To mitigate the challenge of limited labelled data, transfer learning has been widely adopted (Rahman et al., 2020; Petchiammal and Murugan, 2025; Padhi et al., 2024). In addition, attention mechanisms (Al-Gaashani et al., 2023) have been introduced to emphasise discriminative regions, while multimodal frameworks incorporating thermal imaging have been explored to improve robustness under varying environmental conditions (Padhi et al., 2025). Despite achieving high classification accuracy, these approaches exhibit several limitations. Deep architectures rely on repeated downsampling operations, which lead to the loss of fine-grained spatial details that are essential for detecting early-stage symptoms and subtle lesion patterns. Furthermore, pretrained models are typically optimised for general object recognition tasks and fail to capture domain-specific characteristics such as multiscale symptom variations and subtle texture and colour differences.

A few recent studies have explored hybrid and efficiency-oriented approaches to address the limitations of deep-learning and handcrafted feature methodologies. Hybrid models combined handcrafted and deep features (Vandana and Vinod, 2025; Uddin and Joolee, 2024) to improve feature representation; however, simple feature concatenation often introduces redundancy, increases computational cost, and fails to effectively model feature complementarity. Lightweight architectures (Subbarayudu and Kubendiran, 2025; Aggarwal et al., 2024; Padhi et al., 2024) have been proposed to facilitate deployment on resource-constrained devices; however, their reduced representational capacity constrains their ability to effectively capture complex multiscale disease patterns. Similarly, model compression techniques, including structured pruning (Zhu et al., 2022; Kumar et al., 2025) and knowledge distillation (Luo et al., 2023; Cai et al., 2024), improve efficiency but can discard critical features. In knowledge distillation, transferring knowledge between highly similar or identical teacher and student models is not optimal, as it limits the diversity and richness of the learned representations, whereas heterogeneous teacher–student configurations provide more informative supervisory signals. Moreover, many existing models are trained on controlled datasets with simple backgrounds, limiting their robustness in real-world field environments (Yugander et al., 2015).

Despite significant progress, existing approaches exhibit three key limitations: an inability to effectively model multiscale disease patterns, high computational complexity that hinders real-world deployment, and limited robustness under diverse field conditions. The primary limitation arises from the presence of multiscale symptoms, as illustrated in Fig. 1, where disease patterns vary substantially in size, shape, colour, and intensity across different disease types and stages of progression. Such variability makes it difficult to consistently distinguish between diseases, even for experienced experts. Although this challenge is fundamental, most existing approaches lack effective mechanisms to explicitly model multiscale characteristics. Evidence from related domains highlights the importance of multiscale and discriminative feature representations for capturing subtle patterns (Li et al., 2026; Jiang et al., 2025). In addition, existing approaches often prioritise classification accuracy without adequately addressing computational efficiency, which is critical for real-world deployment. Lightweight and efficient models are essential for implementation on resource-constrained devices commonly used in agricultural settings; however, many high-performing models remain computationally intensive. Furthermore, most existing methods are developed and evaluated using images captured under controlled conditions, such as detached leaves with uniform backgrounds, which do not accurately reflect real field environments. As a result, these models often exhibit degraded performance when applied to practical scenarios characterised by complex backgrounds, varying illumination, and occlusions.

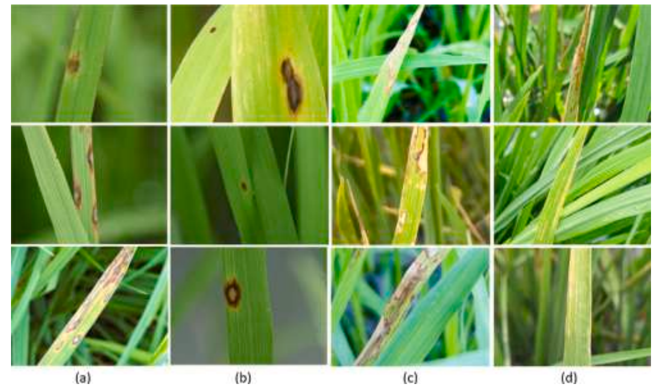

Fig. 1. Illustration of multiscale disease symptoms on rice leaves. From left column to right column: (a) rice blast, (b) brown spot, (c) sheath blight and (d) leaf streak.

This study aims to address these limitations by proposing a novel hybrid architectural framework. The proposed approach integrates a Feature Pyramid Network (FPN) within a HydraNet (Shazeer et al., 2018)-inspired architecture to facilitate accurate disease recognition in the presence of multiscale symptoms. The proposed Haar wavelet-guided FPN model enables multiple parallel branches to process multiscale features through adaptive feature selection, dynamically selecting and emphasising the most relevant feature representations based on the input, thereby allowing the model to specialise in recognising disease symptoms with varying sizes and visual characteristics. Unlike conventional gating strategies in existing HydraNet models, which primarily rely on spatial features, the proposed approach employs a wavelet-based gating mechanism. This mechanism facilitates the exploitation of complementary frequency-domain information, which is crucial for capturing subtle texture variations and complex multiscale disease patterns.

Although the proposed Haar wavelet-guided FPN model achieves outstanding disease recognition accuracy on publicly available datasets, its computational cost remains relatively high. To mitigate this limitation, this study develops a lightweight and computationally efficient model through a knowledge distillation framework while preserving comparable performance. This strategy achieves a favourable balance between accuracy and efficiency, enabling deployment on resource-constrained devices.

Table 1
Comparative analysis of paddy disease classification approaches over the past five years. The table is intended to provide a general summary of existing methods and reported performance trends in the literature rather than a strict benchmark comparison, as the studies were conducted under different experimental settings. H.C. denotes handcrafted features, Deep indicates deep learning based features, and Hybrid represents combined feature approaches. Pub refers to publicly available datasets, #cls indicates the number of disease classes, #img denotes the number of images, RLD indicates Rice Leaf Diseases, RLDD indicates Rice Leaf Diseases Dataset, RLDIS indicates Rice Leaf Disease Image Samples, RV3 indicates Rice Leaf 5 diseases, RDID indicates Rice Diseases Image Dataset and NG indicates information not reported in the original publication.

Approach	Year	Feature type			Dataset				Parameters (M)	Accuracy (%)
		H.C	Deep	Hybrid	Pub	Name	#cls	#img		
Pothen and Pai (2020)	2020	✓	–	–	✓	RLD (UCI)	3	120	NG	94.60
Jiang et al. (2020)	2020	–	✓	–	×	NG	4	NG	0.6	96.80
Li et al. (2020)	2020	–	✓	–	×	Hunan Rice	3	5320	8.0	98.26
Rahman et al. (2020)	2020	–	✓	–	✓	NG	10	1426	0.8	93.33
Prottasha et al. (2022)	2021	–	✓	–	×	NG	11	1677	3.5	96.30
Bari et al. (2021)	2021	–	✓	–	×	NG	3	2400	NG	99.17
Swathika et al. (2021)	2021	–	✓	–	×	NG	1	3081	NG	69.90
Sowmyalakshmi et al. (2021)	2021	–	✓	–	✓	RLD (UCI)	3	120	44.7	94.20
Krishnamoorthy et al. (2021)	2021	–	✓	–	×	NG	3	5200	55.9	95.67
Ganesan and Chinnappan (2022)	2022	–	✓	–	✓	RLDIS (Mendely)	4	5932	33.0	90.00
Latif et al. (2022)	2022	–	✓	–	✓	RiceLeafsv3 (Kaggle)	5	2730	143.7	96.08
Lu et al. (2022)	2022	–	✓	–	×	NG	5	NG	1.0	99.35
Yakkundimath et al. (2022)	2022	–	✓	–	×	UAS Dharwad	24	1200	138.4	92.24
Al-Gaashani et al. (2023)	2023	–	✓	–	×	RDID (Kaggle)	4	400	28.6	98.71
Ahad et al. (2023)	2023	–	✓	–	×	NG	9	900	97.2	97.62
Zeng et al. (2023)	2023	–	✓	–	×	RLD (Kaggle)+ RLDIS (Mendely)	4	10 241	42.6	95.57
Benita et al. (2023)	2023	–	✓	–	✓	RLDIS (Mendely)	6	2710	NG	89.83
Singh et al. (2023)	2023	✓	–	–	×	NG	1	NG	NG	99.10
Jiang et al. (2023)	2023	–	✓	–	×	NG	4	805	86.0	92.00
Zhou et al. (2023)	2023	–	✓	–	×	NG	5	1353	20.8	99.40
Hasan et al. (2023)	2023	–	✓	–	×	NG	3	2700	3.9	97.90
					✓	Dhan-Shomadhan (Mendely)	4	1106	3.9	88.30
					✓	RLDIS (Mendely)	4	5932	3.9	92.70
Kaviya and Selvakumar (2023)	2023	✓	–	–	✓	NG	4	2092	NG	95.34
Kumar et al. (2023)	2023	–	✓	–	✓	RLDIS (Mendely)	4	5932	48.0	99.55
					✓	RLD (UCI)	3	120	48.0	97.05
Salamai et al. (2023)	2023	–	✓	–	✓	Paddy Doctor	10	10 407	60.0	98.74
Petchiammal et al. (2023)	2023	–	✓	–	✓	Paddy Doctor	13	16 225	21.8	97.50
Dubey and Choubey (2024)	2024	–	✓	–	×	Plant Village	4	NG	NG	94.23
Kaur et al. (2024)	2024	–	✓	–	✓	RV3	6	2710	138.4	97.50
Waghmare (2024)	2024	–	✓	–	✓	RV3	6	2710	60.4	95.03
Kumar et al. (2024)	2024	–	✓	–	✓	Paddy Doctor	13	16 225	60.0	94.90
Padhi et al. (2024)	2024	–	✓	–	✓	Paddy Doctor	10	23 916	19.5	96.91
Jithya et al. (2024)	2024	–	✓	–	✓	Paddy Doctor(Kaggle)	10	13 876	20.0	99.94
Ritharson et al. (2024)	2024	–	✓	–	×	RLDIS (Mendely)	9	5932	138.4	99.70
Alsakar et al. (2024)	2024	✓	–	–	✓	RLDIS (Mendely)	4	5932	NG	99.53
					×	NG	5	2800	NG	99.40
					✓	RLDD (Mendeley)	3	4684	NG	99.10
Jain and Kumar (2024)	2024	–	✓	–	×	Paddy Doctor (Kaggle)	10	NG	25.6	98.03
Pandi et al. (2024)	2024	–	✓	–	✓	RLD (UCI)	3	120	NG	96.70
Uddin and Joolee (2024)	2024	✓	✓	✓	✓	Paddy Doctor	13	16 225	55.9	98.52
Padhi et al. (2025)	2025	–	✓	–	✓	RLDIS (Mendely)	4	5932	41.0	99.43
Petchiammal and Murugan (2025)	2025	–	✓	–	✓	RLDIS (Mendely)	4	5932	8.1	97.60
Kumar and Sakthivel (2025)	2025	–	✓	–	✓	RLDD (Kaggle)	6	2627	143.7	97.60
Bala Ayyappan et al. (2025)	2025	–	✓	–	✓	Paddy Doctor (Kaggle)	10	10 407	8.1	97.50
Wu et al. (2025)	2025	–	✓	–	✓	Paddy Doctor (Kaggle)	10	13 876	0.4	97.98
Subbarayudu and Kubendiran (2025)	2025	–	✓	–	✓	Paddy Doctor (Kaggle)	10	13 876	8.1	98.50
Rakshit et al. (2025)	2025	–	✓	–	✓	RLDIS (Mendely)	4	5932	138.4	99.97
Navaneetha et al. (2025)	2025	–	✓	–	✓	RLDIS (Mendely)	4	5932	NG	99.84
Vijayakumar et al. (2025)	2025	–	✓	–	✓	Paddy Doctor	13	16 225	60	96.80

This study introduces NP-LankaPaddy, a novel benchmark dataset for paddy leaf disease recognition, constructed using image samples collected from the Northern region of Sri Lanka. Although several public datasets exist, their applicability to the Sri Lankan context is limited, as disease manifestations are always region-specific. In contrast to most existing datasets, NP-LankaPaddy comprises in-field images capturing paddy leaf disease symptoms across diverse growth stages and environmental conditions.

The key contributions of this work are summarised as follows:

1. A novel Haar wavelet-guided FPN architecture is proposed for paddy leaf disease recognition, enabling effective capture of disease characteristics in the presence of multiscale symptoms.
2. A new benchmark dataset is introduced for the recognition of four paddy leaf diseases. To the best of our knowledge, this represents the first publicly available paddy leaf disease classification dataset collected in the Sri Lankan context.
3. The proposed approach demonstrates state-of-the-art classification performance on four publicly available benchmark datasets, while exhibiting better computational efficiency compared to existing methods.

2. Materials and methods

This study proposes a hybrid feature based model for paddy disease classification. A novel Feature Pyramid Network based HydraNet architecture integrated with an adaptive wavelet gating framework is designed to address the challenge of multiscale lesion manifestations. As illustrated in Fig. 2, deep features are extracted at multiple spatial

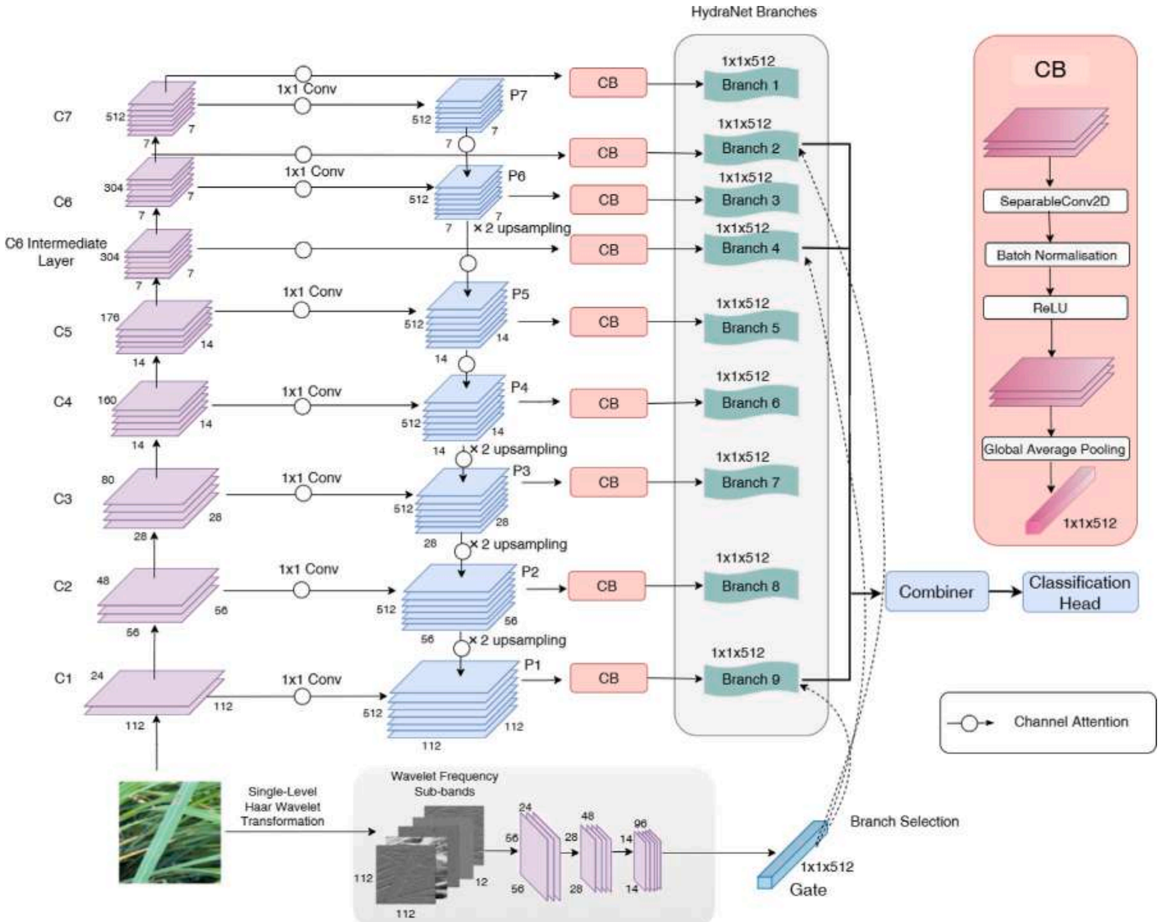

Fig. 2. Overall all architecture of the proposed model. Deep features are extracted at multiple spatial scales from different depths of the EfficientNetV2-M backbone and refined through nine parallel branches. In parallel, a single-level 2D discrete Haar wavelet transform decomposes the input image into multiple subbands, which are further processed by convolutional layers to obtain learnable, disease-aware gating features. These features dynamically guide the selection and fusion of branch outputs.

Fig. 3. Symptoms of four common paddy diseases (left to right): Brown spot — circular to oval brown lesions with dark brown to purplish black margins; Rice blast — spindle shaped leaf lesions with white/ashy centre with dark brown to reddish brown margin; Leaf streak — linear, light-brown lesions between veins; Sheath blight — elliptical or irregular with a white/grey centre and a distinct brown to dark brown margin.

scales from different depths of the EfficientNetV2-M backbone to capture disease symptoms of varying sizes. These features are subsequently refined by a FPN and distributed into multiple parallel branches, each representing a distinct scale of deep representation. In parallel, a single level two dimensional discrete Haar wavelet transform is applied to decompose the input image into multiple subbands, which are further processed by convolutional layers to obtain learnable disease aware gating features. As the Haar wavelet subbands preserve localised intensity variations and abrupt spatial transitions, which are critical cues for early stage disease symptoms, they provide complementary information to the deep features. Finally, these wavelet based gating features dynamically guide the selection and fusion of the branch outputs. In the second phase of the proposed methodology, a lightweight version of the proposed model is developed using a knowledge distillation framework.

The following subsections of the methodology describe the newly constructed NP-LankaPaddy dataset and other benchmark datasets, the proposed network architecture, the training procedures employed, the knowledge distillation framework, the implementation details, and the evaluation metrics and experimental protocols used to assess the performance of the proposed model.

2.1. Datasets

To comprehensively evaluate the classification performance and generalisation capability of the proposed framework, experiments were conducted using the NP-LankaPaddy dataset together with four publicly available benchmark datasets. These datasets differ substantially in terms of disease categories, acquisition environments, dataset size, and visual complexity, thereby enabling a rigorous evaluation under both controlled and real-world cultivation conditions. Consequently, the experiments effectively demonstrate the generalisation capability and robustness of the proposed approach.

2.1.1. NP-LankaPaddy dataset

Although several benchmark datasets exist for paddy leaf disease classification, they mainly comprise images from India and other South Asian countries and do not reflect Sri Lankan cultivation conditions.

Fig. 4. Diversity and generalisability of the NP-LankaPaddy dataset. Images were captured under varying weather conditions, from different viewing angles, and with diverse background settings.

Existing datasets (Tjikoa, 2022; Shah et al., 2017; Saputra, 2022) also predominantly rely on laboratory settings with white backgrounds, limiting their ability to capture real field visual variability. Since disease manifestations are strongly influenced by regional factors such as climate, soil, rice cultivar, and farming practices (de S. Seneviratne and Jeyanandarajah, 2004), a Sri Lankan paddy leaf disease dataset with natural backgrounds is essential for developing reliable and practical recognition systems.

In this study, a new benchmark dataset, termed the NP-LankaPaddy dataset, was constructed for the classification of four major paddy leaf diseases prevalent in Sri Lanka: Brown spot, rice blast, sheath blight, and leaf streak. There is limited availability of publicly accessible datasets representing paddy leaf diseases under Sri Lankan field conditions. The proposed dataset addresses this gap by including four major disease categories relevant to the local context. Consequently, it provides a useful resource for advancing research in Sri Lanka and other paddy growing regions with similar disease distributions.

As illustrated in Fig. 3, brown spot is characterised by circular to oval brown lesions on young leaves and coleoptiles, typically surrounded by dark brown to purplish black margins. Seedlings often exhibit reddish-brown discolouration of varying sizes, which may ultimately result in chalky and discoloured grains. Rice blast is identified by spindle shaped leaf lesions with water soaked lesions with white centres and green margin, which can progress to neck rot and white/ashy centre with dark brown to reddish brown margin. Leaf streak manifests as small, linear, light-brown lesions between veins, often appearing translucent under light and humid conditions, and can ultimately lead to leaf browning and plant death. Sheath blight initial symptoms appear as small, oval to irregular, white/grey or water-soaked lesions on the leaf sheath, as the disease progresses, the lesion centre gradually turns whitish to light grey, while the margins become irregular and brown.

In the NP-LankaPaddy dataset, images were collected across the two major paddy growing seasons in Sri Lanka, namely the *Maha* season and the *Yala* season. Paddy leaf samples were acquired directly from cultivation fields in four major districts of the Northern Province: Jaffna, Mannar, Kilinochchi, and Mullaitivu. The dataset comprises images from multiple paddy cultivars in these regions, including *Aattakkari*, Bg406, Bg360, and Bg300. All images were captured under natural field conditions using mobile phone cameras rather than professional imaging equipment, in order to reflect realistic and practical deployment scenarios. As illustrated in Fig. 4, the NP-LankaPaddy dataset is highly diverse, encompassing images captured from varied viewing angles under different lighting conditions, with diverse backgrounds and environmental variations, and without artificial filtering or alterations, thereby reflecting real-world field conditions.

As summarised in Table 2, a total of 2002 images were collected, of which 1400 images (70%) were allocated for training, while the remaining images were reserved for testing. All images were manually annotated by trained personnel with expertise in plant pathology from the Department of Agricultural Biology, Faculty of Agriculture, University of Jaffna, Sri Lanka. The dataset was annotated at the disease class level based on established diagnostic characteristics under field conditions and includes both clear and complex symptom variations with overlapping features.

Table 2
Composition and characteristics of the NP-LankaPaddy dataset.

Dataset characteristic	Description
Dataset name	NP-LankaPaddy
Total number of images	2002
Image resolution (pixels)	768 × 1024
Number of disease classes	4
Brown spot	575 images (Train: 402, Test: 173)
Leaf streak	610 images (Train: 427, Test: 183)
Rice blast	512 images (Train: 358, Test: 154)
Sheath blight	305 images (Train: 213, Test: 92)
Image acquisition setting	In-field (natural cultivation conditions)
Paddy growing seasons	*Maha* (1667 images), *Yala* (335 images)
Data collection regions	Jaffna, Mannar, Kilinochchi, Mullaitivu in Northern Province
Rice cultivars	*Aattakkari*, Bg406, Bg360, Bg300
Image variability factors	Illumination, viewing angles, and background clutter
Annotation methodology	Manual annotation under the supervision of agricultural officers
Preprocessing steps	Cropping, and resizing
Image capture devices	iPhone 14, Galaxy A54
Dataset link	https://doi.org/10.5281/zenodo.18334503

2.1.2. Publicly available benchmark datasets

In addition to the newly constructed NP-LankaPaddy dataset, the proposed framework was further evaluated on four publicly available benchmark datasets, namely the UCI-Rice Leaf Diseases dataset (Shah et al., 2017), Mendeley-Rice Leaf Disease Image Samples dataset (Sethy, 2024), RiceLeafsv3 dataset (Tjikoa, 2022), and Paddy Doctor dataset (Petchiammal et al., 2022).

UCI-Rice Leaf Diseases Dataset: The UCI rice leaf diseases dataset (Shah et al., 2017) is a small-scale multiclass classification benchmark comprising 120 images captured under controlled conditions against a uniform white background. It contains three disease categories: bacterial leaf blight, brown spot, and leaf smut. For model training and evaluation, the dataset is divided into 70% training and 30% testing subsets. To enhance generalisation and compensate for the limited dataset size, data augmentation techniques are applied during training.

Mendeley-Rice Leaf Disease Image Samples Dataset: The Mendeley Rice Leaf Disease Image Samples dataset (Sethy, 2024) is a widely used benchmark for high-precision paddy disease classification and has been adopted in numerous comparative studies. It comprises 5932 images spanning four disease categories: bacterial blight, blast, brown spot, and tungro. Although the images exhibit some background variation, the overall diversity remains limited compared with in-field datasets. For this study, the dataset is partitioned into 70% training and 30% testing subsets to evaluate classification performance.

RiceLeafsv3: The RiceLeafsv3 dataset (Tjikoa, 2022) is utilised to assess the generalisation capability of the proposed framework across multiple paddy leaf disease categories. It comprises 2710 images organised into predefined training and validation subsets, captured under controlled conditions against a uniform white background. The dataset includes six classes: bacterial leaf blight, brown spot, healthy, leaf blast, leaf scald, and narrow brown spot. In this study, the original split provided by the authors is adopted without modification, where the training subset is used for model learning and the validation subset for testing, ensuring a fair and consistent evaluation protocol aligned with prior work.

Paddy Doctor: The Paddy Doctor dataset (Petchiammal et al., 2022) constitutes the most challenging benchmark employed in this study, comprising 13 paddy disease categories collected under real-world in-field conditions. It includes a total of 16,225 labelled paddy leaf images spanning the following classes: bacterial Leaf Blight, bacterial Leaf Streak, bacterial panicle blight, black stem borer, blast, brown spot, downy mildew, hispa, leaf roller, tungro, white stem borer, yellow stem borer, and normal. These categories exhibit considerable variability in symptom manifestation, illumination conditions, and background clutter. This study utilises the paddy-doctor-diseases-small-split version of the dataset, which provides a predefined training-testing partition. The training set comprises 12,980 images (80%), while the remaining 3245 images (20%) are reserved for testing. As the dataset is pre-partitioned, the original split is directly adopted to ensure fair and reproducible evaluation and to enable consistent comparison with existing studies that employ the same dataset configuration.

2.2. Architecture of the proposed model

The proposed Haar wavelet-guided FPN paddy disease classification model as shown in Fig. 2, comprises three main modules: (i) a Feature Pyramid Network for extracting and fusing rich multiscale spatial representations; (ii) a Haar wavelet-based feature extraction module for capturing discriminative frequency-domain information; and (iii) a HydraNet-inspired gating mechanism that dynamically selects the most informative deep feature pathways based on the characteristics of the input images. Collectively, these components facilitate robust and context-aware feature fusion, effectively addressing challenges such as scale variation, texture diversity, and real in-field imaging conditions in paddy leaf disease recognition.

2.2.1. Feature pyramid network module

Conventional CNN-based paddy disease classification approaches typically rely on deep features extracted at a single spatial resolution, often from the final convolutional layer, which limits their ability to capture disease symptoms occurring at multiple scales. In paddy leaf images, visual manifestations such as lesions, streaks, and discolourations vary from fine-grained spots to large infected regions. Furthermore, successive downsampling operations in deep networks reduce spatial resolution, frequently causing the loss of critical fine-grained details that are essential for accurate and early disease recognition. To address these limitations, a Feature Pyramid Network (FPN) module is integrated into the proposed approach.

The proposed FPN module constructs a hierarchical multiscale feature representation that preserves fine-grained local details as well as broader contextual information, thereby enabling the network to jointly model small lesion patterns and large-scale disease regions through feature fusion across multiple spatial resolutions. As shown in Fig. 2, the proposed FPN module exploits low-level feature maps to capture fine-grained cues such as small lesions and sharp disease boundaries, while higher-level maps encode semantically rich representations that model broader contextual patterns, including extensive discolouration and advanced disease spread.

EfficientNetV2-M (Tan and Le, 2021) is selected as the backbone of the FPN module due to its favourable balance between computational efficiency and representational capacity. The proposed FPN module employs bottom up and top down pathways with lateral connections to facilitate effective multiscale feature integration. The bottom up pathway corresponds to the standard feed forward operations of the EfficientNetV2-M backbone, producing a hierarchy of feature maps with increasing semantic richness and decreasing spatial resolution. In this pathway, feature maps are extracted from key stages of the backbone, denoted as C_l to C_7 in Fig. 2, to capture a broad spectrum of information, ranging from low level texture and edge details in shallow layers to high level semantic representations in deeper layers. These multi depth features provide a comprehensive characterisation of paddy leaf disease patterns across different abstraction levels.

The top down pathway of the proposed FPN upsamples higher level, semantically rich feature maps and fuses them with corresponding lower level feature maps through lateral connections. This pathway is crucial for enhancing sensitivity to small lesions and subtle discolourations on paddy leaves. Prior to feature fusion, a 1×1 convolution with 512 filters is applied along the lateral (horizontal) connections to align channel dimensions, ensuring compatibility for element-wise addition while preserving both local and global contextual information.

In the EfficientNetV2-M backbone, feature maps extracted at different stages exhibit progressively increasing strides with respect to the input image resolution (H, W), resulting in feature representations of size $(H/s_l, W/s_l)$. To achieve spatial alignment during fusion, each coarser feature map is upsampled according to the stride ratio between adjacent pyramid levels. In addition, channel attention is consistently applied to both the lateral feature maps and the top-down upsampled features prior to fusion, thereby enabling adaptive recalibration of feature responses across both pathways. This dual-path attention mechanism strengthens the representation of fine-grained colour variations and small lesion patterns in paddy leaves, leading to more effective multiscale disease classification. The proposed top-down fusion process is therefore formulated as:

$$P_l = \mathcal{A}(C_l) + \mathcal{A}\Big(\text{Upsample}\big(P_{l+1}, \tfrac{s_{l+1}}{s_l}\big)\Big), \tag{1}$$

where $\mathcal{A}(\cdot)$ denotes channel attention, which emphasises informative channels and suppresses less relevant ones. C_l is the backbone feature at level l, P_l is the fused feature map, and s_l represents the stride. Upsample$(\cdot)$ resizes P_{l+1} by $\frac{s_{l+1}}{s_l}$ to match the spatial resolution of level l for element-wise fusion.

In the proposed FPN module, a Convolution Batch Normalisation (CB) block is applied after the feature fusion stage, as illustrated in Fig. 2. This block enhances feature discrimination by sequentially integrating separable convolution, batch normalisation, and ReLU activation. The separable convolution enables efficient spatial and channel-wise feature refinement of the fused multiscale representations while reducing computational complexity. Batch normalisation stabilises the fused feature distributions and facilitates reliable training, while the ReLU activation enhances discriminative disease-related patterns. At the end of the CB block, a Global Average Pooling layer is employed to aggregate spatial information across each feature map, compressing the refined multiscale features into a compact $1 \times 1 \times 512$ representation.

The proposed FPN model produces nine outputs, each with a dimensionality of $1 \times 1 \times 512$, representing multilevel, multiscale features of the input image derived from different pyramid levels.

2.2.2. Haar wavelet-based feature extraction module

In the proposed approach, Haar wavelet-based features are extracted and utilised as complementary frequency-domain cues to guide the FPN-based deep features through an adaptive HydraNet-based gating mechanism. These Haar wavelet-based features complement the deep FPN representations by capturing frequency-domain and localised texture information that is not explicitly modelled by conventional convolutional feature hierarchies. While FPN-based deep features primarily learn spatial-domain representations with increasing semantic abstraction across multiple scales, they may overlook subtle high-frequency patterns that are crucial for discriminating fine-grained visual variations in paddy leaf disease classification.

The proposed Haar wavelet-based feature extraction module operates directly on the raw input image and generates the gating vector used in the HydraNet-based gating mechanism. Wavelet features offer a joint spatial frequency representation that explicitly separates low- and high-frequency components while preserving spatial locality (Mallat, 1989). High-frequency sub-bands capture abrupt intensity changes, sharp edges, and fine-grained texture variations associated with small lesions and early-stage disease symptoms, whereas low-frequency components encode smooth intensity variations and coarse structural patterns corresponding to larger infected regions and advanced disease progression. This complementary frequency decomposition renders wavelet features well suited for guiding branch selection in multiscale architectures.

In the proposed module, the input image is decomposed using a single-level two-dimensional discrete Haar wavelet transform. For each colour channel, this decomposition produces four sub-bands as shown in Fig. 5 one low-frequency approximation (LL) and three high-frequency detail components (LH, HL, and HH). In a two-dimensional

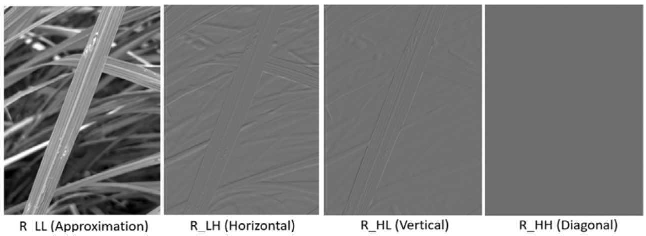

Fig. 5. Illustration of the single-level 2D discrete Haar wavelet decomposition showing four sub-bands for Red colour channel: LL (approximation), and LH, HL, HH (horizontal, vertical, and diagonal detail components).

discrete wavelet transform, the LL sub-band captures low-frequency approximation information representing global structure, while the LH, HL, and HH sub-bands encode high-frequency directional details corresponding to horizontal, vertical, and diagonal intensity variations, respectively. In the proposed module, these sub-bands are stacked to form a frequency-aware representation of size $\frac{H}{2} \times \frac{W}{2} \times 12$, where H and W denote the original image dimensions.

As described in Mallat (1989), the one-level two-dimensional discrete Haar wavelet decomposition of an image channel $I(x, y)$ can be expressed as follows:

$$\begin{aligned} LL &= (I * \phi_x * \phi_y) \downarrow 2, \\ LH &= (I * \phi_x * \psi_y) \downarrow 2, \\ HL &= (I * \psi_x * \phi_y) \downarrow 2, \\ HH &= (I * \psi_x * \psi_y) \downarrow 2, \end{aligned} \tag{2}$$

where $*$ denotes convolution, ϕ and ψ are the Haar scaling (low-pass) and wavelet (high-pass) functions, respectively, and $\downarrow 2$ represents downsampling by a factor of two along each spatial dimension.

In the implementation, wavelet decomposition is integrated as a learnable operation within the network rather than using a fixed configuration. Symmetric padding is applied to preserve edge continuity and reduce boundary artefacts, while a stride of two maintains multiscale feature representation through downsampling. The computational overhead introduced by the proposed learnable wavelet decomposition is negligible, since the decomposition operation is lightweight and implemented using efficient convolutional operations within the network. Rather than directly utilising the raw wavelet coefficients as gating features, the proposed module learns disease-aware representations through trainable transformations. To further enhance adaptability, two trainable scalar parameters, α and β, are introduced to modulate the contribution of different frequency components as follows:

$$\tilde{W}_k = \begin{cases} \alpha \cdot W_k, & k = \mathrm{LL}, \\ \beta \cdot W_k, & k \in \{\mathrm{LH}, \mathrm{HL}, \mathrm{HH}\}, \end{cases} \tag{3}$$

where W_k denotes the Haar wavelet sub-bands and $\tilde{W}_k$ represents the modulated coefficients. The parameter α controls the contribution of low-frequency structural information, while β regulates high-frequency detail responses associated with lesion boundaries and fine texture variations. Both parameters are initialised to 1.0 and optimised through backpropagation, enabling the gating mechanism to dynamically adapt to disease characteristics of varying scale.

In place of directly utilising the extracted features, the proposed framework further refines the frequency-aware representations through a lightweight convolutional sub-network. This sub-network transforms the fixed wavelet responses into learnable, disease-aware features. Specifically, the stacked wavelet sub-bands of dimension $\frac{H}{2} \times \frac{W}{2} \times 12$ are processed by successive convolution batch normalisation ReLU blocks, interleaved with max-pooling operations. These stages progressively reduce the spatial resolution while increasing the channel dimensionality, enabling the aggregation of local frequency responses into compact and discriminative global representations. Finally, global average pooling is applied to condense the refined wavelet feature maps into a compact $1 \times 1 \times 512$ gating vector.

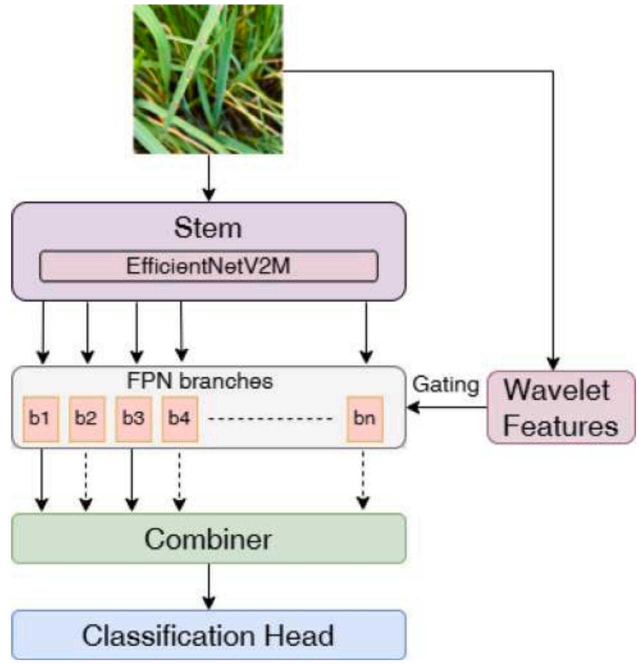

Fig. 6. Illustration of the HydraNet-inspired dynamic feature selection module. EfficientNetV2-M serves as the stem, while multilevel, multiscale features generated by the FPN constitute the branches. Wavelet features are employed to guide the gating mechanism for adaptive branch selection.

2.2.3. Gating based dynamic feature selection module

In the proposed model, FPN features are extracted from multiple depths and spatial scales to capture a broad range of semantic and structural cues. However, the relevance of these multiscale features is not uniform across all paddy leaf images, as disease symptoms vary significantly in size, texture, severity, and spatial distribution. For example, early-stage diseases may manifest as small, fine-grained lesions, whereas advanced infections often appear as large discoloured regions. To address this challenge, a gating-based dynamic feature selection mechanism is introduced to adaptively determine the importance of each FPN branch on a per-image basis. Haar wavelet-based features, extracted in parallel, serve as guidance signals for this gating process by providing explicit frequency-domain information that reflects the dominant texture and scale characteristics of the input image. By utilising these complementary cues, the gating mechanism selectively emphasises the most relevant FPN branch features while suppressing less informative ones, thereby enhancing discrimination of diverse disease patterns and improving robustness under varying field conditions.

The proposed gating-based feature selection module is inspired by the conceptual design of HydraNet, while adopting a distinct architectural formulation. Similar to HydraNet, the proposed module comprises four conceptual components: a stem, branches, a gate, and a combiner as shown in Fig. 6. In this architecture, EfficientNetV2-M serves as the stem, providing a shared backbone for feature extraction. Unlike the parallel network pathways employed in conventional HydraNet, the branches in the proposed module correspond to multilevel, multiscale feature representations generated by the FPN. Specifically, features extracted at nine pyramid levels are treated as independent feature branches.

The gate unit of the proposed dynamic feature selection module first takes the learnable Haar wavelet features as a feature vector, which can be represented as follows:

$$G \in \mathbb{R}^{1 \times 512}. \tag{4}$$

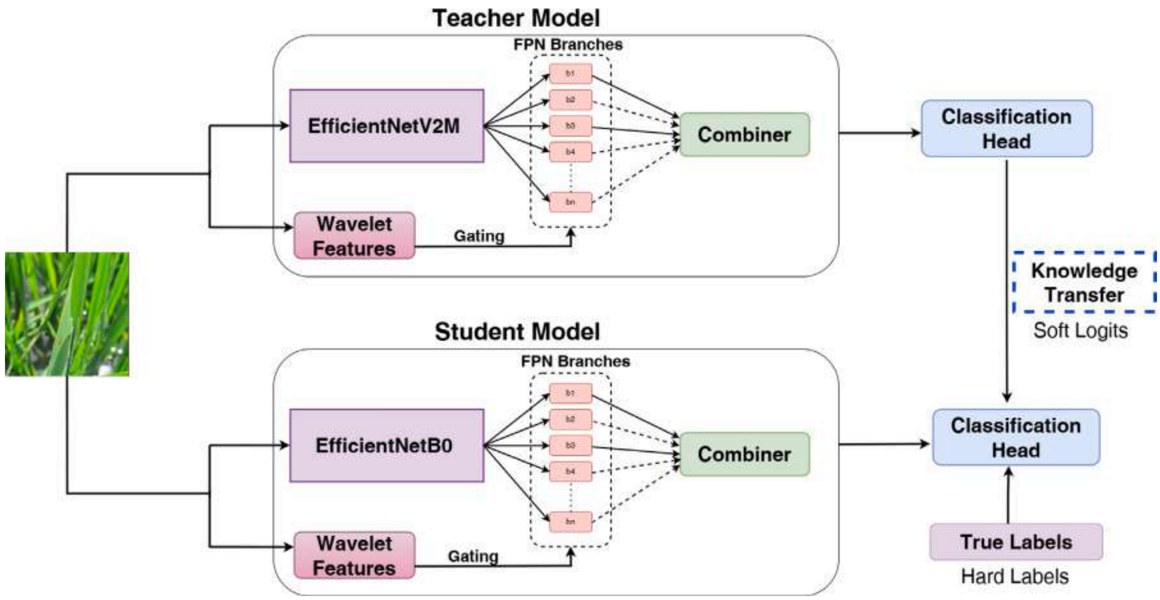

Fig. 7. Overview of the proposed Haar wavelet-guided FPN knowledge distillation framework. EfficientNetV2-M based Haar wavelet-guided FPN is employed as the teacher model, while EfficientNet-B0 based Haar wavelet-guided FPN serves as the student model.

This vector represents the global spatial frequency characteristics of the input image, capturing the relative presence of coarse structural patterns and fine-grained texture variations associated with multiscale disease symptoms. The gating vector G is subsequently passed through a fully connected layer with softmax activation to compute normalised importance scores for each FPN branch as follows:

$$\boldsymbol{\alpha} = \mathrm{softmax}(W_g G + b_g), \quad \boldsymbol{\alpha} \in \mathbb{R}^9, \tag{5}$$

where W_g and b_g denote the learnable weight matrix and bias term, respectively. Each element α_i represents the relative importance assigned to the ith FPN branch for a given input image.

To reduce redundancy and enhance discriminative focus, only the top k branches with the highest importance scores are retained as follows:

$$\{i_1, i_2, \dots, i_k\} = \mathtt{top\_k}(\boldsymbol{\alpha}, k), \tag{6}$$

where k is a predefined constant specifying the number of dominant branches selected for feature fusion and is determined experimentally.

Finally, the combiner of the proposed module aggregates the selected branch features through weighted summation as follows:

$$\hat{F} = \sum_{j=1}^{k} \alpha_{i_j} \cdot F_{i_j}, \tag{7}$$

where F_{i_j} denotes the feature vector corresponding to the jth selected FPN branch. This adaptive fusion mechanism dynamically emphasises features from the most relevant spatial scales based on the wavelet-informed gating scores. By integrating wavelet-informed gating with dynamic branch selection, the proposed architecture establishes a context-aware feature fusion pipeline in which spatial frequency information explicitly guides the utilisation of multilevel spatial representations. This design enables robust discrimination of paddy leaf disease patterns across multiscale symptoms under real in-field conditions.

After the combiner, the aggregated adaptively fused features $\hat{F}$ are forwarded to the classification head, which consists of three fully connected layers with 256, 128, and 4 neurons, respectively, interleaved with ReLU activation and dropout layers.

2.2.4. Loss function and training strategy

In this study, paddy disease recognition is formulated as a multiclass image classification task. To balance adaptation and stability, the final convolutional layers of the backbone are fine tuned while the early layers are frozen. The proposed model is trained in an end to end manner using the sparse categorical cross entropy loss (L_{SCE}), defined as follows:

$$\mathcal{L}_{\mathrm{SCE}} = -\frac{1}{N} \sum_{i=1}^{N} \log \left(p_{i,y_i} \right), \tag{8}$$

where N denotes the number of training samples, y_i represents the ground-truth class label of the ith sample, and p_{i,y_i} corresponds to the predicted probability of the correct class. To address class imbalance and enhance generalisation performance, the training dataset is balanced using class-wise data augmentation. Augmentation techniques, including rotation, width and height shifts, zooming, and horizontal and vertical flipping, introduce realistic variations in scale, orientation, and appearance, thereby improving classification accuracy and reducing overfitting. An early stopping strategy based on validation loss is employed to ensure stable training, and the model parameters corresponding to the highest validation accuracy are retained for evaluation.

2.3. Lightweight model development using knowledge distillation

Although the proposed Haar wavelet-guided FPN model achieves outstanding accuracy, it incurs high computational complexity due to the use of the EfficientNetV2-M backbone. Therefore, in the subsequent phase of the methodology, a lightweight model is developed using a knowledge distillation technique (Song et al., 2023).

In the proposed approach, knowledge is transferred from a computationally intensive teacher model employing the EfficientNetV2-M backbone to a lightweight student model employing the EfficientNetB0 backbone. As illustrated in Fig. 7, both the teacher and student models have similar architectures, comprising an FPN feature extraction module with multiple branches, a Haar wavelet based feature extraction module, a gating based dynamic feature selection module, a feature combiner, and a classification head. The teacher model is first trained on the training dataset, and the same dataset is subsequently used to train the student model during knowledge distillation.

The proposed knowledge transfer mechanism utilises softened probability distributions obtained from the teacher model's classification layer, which encode rich inter-class relationships that are crucial for discriminating subtle variations in disease symptoms. For example, when processing a brown spot image, the teacher model produces a probability distribution $\mathbf{P}_T^{\tau} = [0.5, 0.2, 0.1, 0.2]$ across the four disease classes: brown spot, sheath blight, leaf streak, and rice blast, respectively. These softened probabilities capture semantic similarities and ambiguities among disease classes that cannot be represented by conventional one-hot encoding schemes.

Let $\mathbf{z}_T$ and $\mathbf{z}_S$ denote the logits produced by the teacher and student models, respectively, and let $\mathbf{y}$ represent the ground-truth labels. Let N_{class} denote the number of classes, and τ be the temperature parameter that controls the softness of the predicted probability distributions. The temperature-scaled (softened) class probabilities are computed as

$$\mathbf{P}_k^{\tau}(i) = \text{softmax}\left(\frac{\mathbf{z}_k}{\tau}\right)_i = \frac{\exp(z_{k,i}/\tau)}{\sum_{j=1}^{N_{\text{class}}} \exp(z_{k,j}/\tau)}, \quad k \in \{T, S\}, \tag{9}$$

where i denotes the class index, and $i \in \{1, 2, \dots, N_{\text{class}}\}$.

During the distillation process, the input paddy leaf images $\mathbf{x}$ are simultaneously fed into the fine-tuned teacher model and the student model, producing the softened probability distributions $\mathbf{P}_T^{\tau}$ and $\mathbf{P}_S^{\tau}$, respectively. The student model is optimised using a composite knowledge distillation loss (L_{KD}) that combines the cross-entropy loss (L_{CE}) with the ground-truth labels and the Kullback–Leibler (KL) divergence between the teacher and student predictions as follows:

$$\mathcal{L}_{KD} = (1-\alpha)\,\mathcal{L}_{CE}(\mathbf{P}_S, \mathbf{y}) + \alpha\,\tau^2\,\mathcal{L}_{KL}(\mathbf{P}_T^{\tau} \parallel \mathbf{P}_S^{\tau}), \tag{10}$$

$$\mathcal{L}_{CE}(\mathbf{P}_S, \mathbf{y}) = -\sum_{i=1}^{N_{\text{class}}} y_i \log(P_{S,i}), \tag{11}$$

where y_i denotes the one-hot encoded ground-truth label for the ith class. The Kullback–Leibler divergence term is expressed as:

$$\mathcal{L}_{KL}(\mathbf{P}_T^{\tau} \parallel \mathbf{P}_S^{\tau}) = \sum_{i=1}^{N_{\text{class}}} P_{T,i}^{\tau} \log\left(\frac{P_{T,i}^{\tau}}{P_{S,i}^{\tau}}\right). \tag{12}$$

The temperature parameter τ controls the smoothness of the predicted probability distributions, thereby facilitating the transfer of richer inter-class relationships, while the balancing coefficient α regulates the relative contribution of the soft teacher supervision during training. Based on empirical evaluation, the balancing coefficient α is set to 0.3 and the temperature parameter τ is set to 4, as these values provide an effective trade-off between knowledge transfer efficiency and classification performance.

2.4. Implementation details

The proposed approach was implemented in Python using the TensorFlow deep learning library and evaluated on an NVIDIA Tesla T4 GPU with 16 GB of dedicated memory. During training, the early layers of the backbone network are frozen to retain low-level visual features, while the final 15 layers are unfrozen to adapt higher-level semantic representations to the target domain. The teacher model is fine-tuned using the Adam optimiser with a learning rate of 5×10^{-5}. In contrast, the student model is trained exclusively under the knowledge distillation framework with a higher learning rate of 1×10^{-4} to promote efficient transfer of softened knowledge from the teacher. All input images are resized to 224 × 224 during both the training and testing phases. During the knowledge distillation process, the top three branches, as defined in Eq. (6), are selected based on experimental results. Training is performed with a batch size of 32 for up to 100 epochs using sparse categorical cross-entropy in conjunction with the distillation loss, and early stopping is employed to mitigate overfitting by retaining the model weights corresponding to the best validation accuracy. The code of this work is publicly available at https://github.com/tharmini/paddy-leaf-diseases-classification.git.

2.5. Evaluation metrics and experimental protocols

The multiclass classification performance of the proposed paddy leaf disease recognition approach is measured using the following evaluation metrics:

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN} \tag{13}$$

$$\text{Precision} = \frac{TP}{TP + FP} \tag{14}$$

$$\text{Recall} = \frac{TP}{TP + FN} \tag{15}$$

$$F_1\ score = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} \tag{16}$$

where TP, TN, FP, and FN represent the number of true positives, true negatives, false positives, and false negatives, respectively.

The performance of the proposed approach is compared with state-of-the-art deep feature based (Ritharson et al., 2024; Padhi et al., 2025; Sowmyalakshmi et al., 2021; Latif et al., 2022; Kumar et al., 2023; Petchiammal et al., 2023; Kaur et al., 2024; Waghmare, 2024; Kumar et al., 2024; Pandi et al., 2024; Petchiammal and Murugan, 2025; Navaneetha et al., 2025; Vijayakumar et al., 2025; Aggarwal et al., 2024; Nanda et al., 2025) and hand-crafted feature based (Pothen and Pai, 2020; Benita et al., 2023; Alsakar et al., 2024; Mekha and Teeyasuksaet, 2021) methods using their reported results. Although several other approaches have also been applied to these datasets, they are excluded from this comparison because they either do not consider all disease classes within a dataset or are not evaluated on the complete test set. In addition, to evaluate the generalisation capability of the proposed model, the model trained on the NP-LankaPaddy dataset was tested on the Paddy Doctor dataset, which contains three common disease classes. Moreover, all ablation experiments were conducted by training the models on the training set of the NP-LankaPaddy dataset and evaluating them on the corresponding test set.

3. Experimental results

3.1. Performance comparison with state-of-the-art methods

The classification performance of the proposed Haar wavelet-guided FPN model and its lightweight variant, referred to as the Haar wavelet-guided FPN architecture with Knowledge Distillation (KD), is evaluated on the test sets of four benchmark datasets as well as our newly construed NP-LankaPaddy dataset.

As summarised in Table 3, the proposed Haar wavelet-guided FPN model achieves superior or competitive classification accuracy across all evaluated datasets. Furthermore, its knowledge-distilled variant demonstrates outstanding classification performance while utilising significantly fewer parameters. Notably, on controlled environment datasets such as the UCI-Rice Leaf Diseases dataset and RiceLeafsv3, where images are captured against uniform white backgrounds, the knowledge distillation variant of the proposed model achieves accuracies of 97.22% and 97.24%, respectively, demonstrating robust discriminative capability under controlled image acquisition conditions. In addition, both the proposed model and its KD variant achieve 100% classification accuracy on the Mendeley-Rice Leaf Disease Image dataset. Furthermore, on the more challenging Paddy Doctor dataset, the proposed wavelet-guided FPN model and its KD variant attain accuracies of 98.04% and 97.84%, respectively.

In addition to its strong classification performance, the knowledge distillation variant of the proposed Haar wavelet-guided FPN model

Table 3
Classification and computational efficiency comparison on benchmark datasets. The reported results of previous approaches are adopted for this comparison. #Params denotes the number of model parameters, and N/M indicates values not mentioned in the original publications.

Dataset	Approach	Year	#Params (M)	FLOPs (G)	Accuracy (%)
UCI-Rice leaf diseases (Shah et al., 2017)	SVM + HOG (Pothen and Pai, 2020)	2020	N/M	N/M	94.60
	Pretrained DenseNet (Chen et al., 2020)	2020	25	8.7	92.22
	Random Forest (Mekha and Teeyasuksaet, 2021)	2021	N/M	N/M	69.44
	ResNetV2 (Sowmyalakshmi et al., 2021)	2021	44.7	16.7	94.20
	Vision Transformer - RDTNet (Kumar et al., 2023)	2023	N/M	N/M	97.05
	CNN + Decision Tree (Pandi et al., 2024)	2024	N/M	N/M	96.70
	Proposed Haar wavelet-guided FPN model	2026	56.5	16.2	**97.22**
	Proposed Haar wavelet-guided FPN with KD model	2026	**7.0**	**5.3**	**97.22**
Mendeley-Rice leaf disease image (Sethy, 2024)	Vision Transformer - RDTNet (Kumar et al., 2023)	2023	N/M	N/M	99.54
	Custom VGG16 (Ritharson et al., 2024)	2024	138.4	30.9	99.70
	LBP + Colour Correlogram + SVM (Alsakar et al., 2024)	2024	N/M	N/M	99.14
	EfficientNetB3 (Aggarwal et al., 2024)	2024	12	**2**	99.00
	ResNet50 (Rakshit et al., 2025)	2025	138.4	30.9	99.97
	KLD-CNN (Navaneetha et al., 2025)	2025	N/M	N/M	99.84
	Darknet53 + SVM (Padhi et al., 2025)	2025	N/M	N/M	99.43
	DenseNet121 (Petchiammal and Murugan, 2025)	2025	8.1	5.7	97.60
	Proposed Haar wavelet-guided FPN model	2026	56.5	16.2	**100.00**
	Proposed Haar wavelet-guided FPN with KD model	2026	**7.0**	5.3	**100.00**
RiceLeafsv3 (Tjikoa, 2022)	VGG19 (Latif et al., 2022)	2022	143.7	39.2	96.08
	Random Forest (Benita et al., 2023)	2023	N/M	N/M	89.83
	VGG16 (Kaur et al., 2024)	2024	138.4	30.9	97.50
	ResNet152V2 (Waghmare, 2024)	2024	60.4	22.4	95.03
	Proposed Haar wavelet-guided FPN model	2026	56.5	16.2	**97.61**
	Proposed Haar wavelet-guided FPN with KD model	2026	**7.0**	**5.3**	97.24
Paddy Doctor (Petchiammal et al., 2022)	ResNet34 (Petchiammal et al., 2023)	2023	21.8	7.4	97.50
	AlexNet (Kumar et al., 2024)	2024	60	1.4	94.90
	MobileNetV4 (Nanda et al., 2025)	2025	**3.8**	**0.4**	97.83
	ResNet152 (Vijayakumar et al., 2025)	2025	60	22.6	96.80
	Proposed Haar wavelet-guided FPN model	2026	56.5	16.2	**98.04**
	Proposed Haar wavelet-guided FPN with KD model	2026	7.0	5.3	97.84
NP-Lanka Paddy (Ours) (Tharmini et al., 2026)	**Proposed Haar wavelet-guided FPN model**	2026	56.5	16.2	**95.18**
	Proposed Haar wavelet-guided FPN with KD model	2026	**7.0**	**5.3**	95.02

demonstrates competitive computational efficiency. As shown in Table 3, the model comprises 7 million parameters and requires 5.3 GFLOPs. Furthermore, the model achieves an average inference time of 82.60 ms per image (12.1 FPS). These results indicate a favourable trade-off between model complexity and classification performance.

The class-wise and average performance of the proposed Haar wavelet-guided FPN model and its knowledge distillation variant on the NP-LankaPaddy dataset are reported in Table 4. The results indicate that the KD variant achieves comparable performance across all classes while maintaining lower computational cost. The dataset exhibits moderate class imbalance; however, this is mitigated through data augmentation during training. Despite this imbalance, the class-wise results in Table 4 show consistent performance across all categories, with minority classes achieving comparable F1-scores. This indicates that the model does not exhibit bias towards majority classes and maintains balanced predictive capability.

Furthermore, during cross-dataset validation, the model achieved an accuracy of 98.05% on the three shared classes without any fine-tuning, demonstrating strong generalisation capability across datasets despite variations in data distribution. Although a comprehensive evaluation across all datasets is constrained by differences in class categories, these results highlight the robustness and generalisability of the proposed approach.

3.2. Ablation studies

In this study, a comprehensive ablation analysis was conducted to justify the design choices of the proposed model and to validate its parameters.

3.2.1. Component analysis

The proposed Haar wavelet-guided FPN model comprises several key components, including an EfficientNetV2-M backbone, a Haar wavelet based feature extraction module, a Feature Pyramid Network (FPN), a Convolution Batch Normalisation (CB) block, and a wavelet gating based dynamic feature selection module. To validate the necessity and effectiveness of these components, an ablation study was conducted, and the corresponding results are summarised in Table 5.

The EfficientNetV2-M model demonstrates strong baseline performance, whereas the Haar wavelet feature-based model alone is less effective, indicating that wavelet features in isolation are insufficient for reliable classification. Incorporating the FPN into the baseline enhances performance, highlighting the importance of multi-scale feature aggregation in capturing disease patterns of varying sizes. Further improvements are achieved by integrating the proposed wavelet-guided gating mechanism with the FPN, which emphasises the most relevant multiscale features through adaptive selection. Finally, the inclusion of the CB block leads to the best overall performance 95.18%, indicating that additional feature refinement enhances the discriminative capability of the model. Overall, these results demonstrate that each component contributes progressively, with the wavelet-guided gating mechanism playing a key role in improving classification effectiveness.

3.2.2. Effectiveness of the proposed wavelet-guided gating mechanism

To evaluate the effectiveness of the proposed wavelet-guided gating mechanism, a comparative analysis was conducted against standard attention-based feature weighting methods, including the Squeeze-and-Excitation (SE) Block and CBAM (Convolutional Block Attention Module), within the same FPN-based framework. As presented in Table 6, the proposed wavelet-guided gating mechanism achieves superior performance compared to these conventional attention modules. Unlike SE and CBAM, which derive attention weights solely from deep

Table 4
Class-wise performance of the proposed model on the NP-LankaPaddy dataset. Abbreviations are defined as P: Precision, R: Recall, and F1: F1-Score.

Disease Class	Model	P (%)	R (%)	F1 (%)
Rice blast	Haar wavelet-guided FPN	96.84	99.35	98.08
	Haar wavelet-guided FPN with KD model	98.03	96.75	97.39
Leaf streak	Haar wavelet-guided FPN model	94.57	95.08	94.82
	Haar wavelet-guided FPN with KD model	95.03	93.99	94.51
Brown spot	Haar wavelet-guided FPN model	92.40	91.33	91.86
	Haar wavelet-guided FPN with KD model	90.11	94.74	92.37
Sheath blight	Haar wavelet-guided FPN model	98.88	95.65	97.24
	Haar wavelet-guided FPN with KD model	100	94.57	97.21
Average	Haar wavelet-guided FPN model	95.67	95.35	95.50
	Haar wavelet-guided FPN with KD model	95.79	95.01	95.37

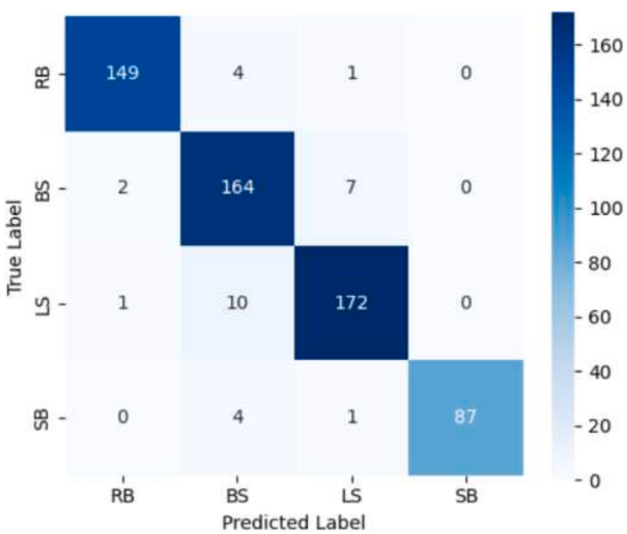

(a) Confusion matrix.

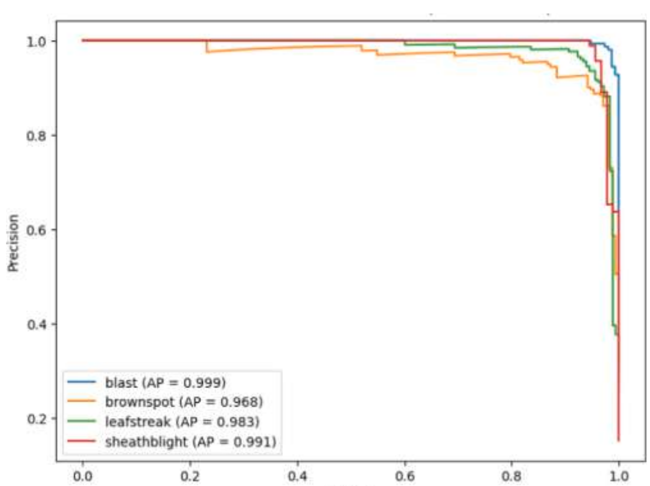

(b) Precision–recall curves.

Fig. 8. Confusion matrix (top) and class-wise precision–recall curves (bottom) of the proposed Haar wavelet-guided FPN student model with knowledge distillation evaluated on the NP-LankaPaddy dataset. Here, RB, BS, LS, and SB denote Rice blast, Brown spot, Leaf streak, and Sheath blight, respectively.

Table 5
Component-wise evaluation of the proposed architecture under various configurations, comparing the contributions of the baseline model (EfficientNetV2-M), the wavelet-based model, the Feature Pyramid Network (FPN), the wavelet gating module, and the Convolution Batch Normalisation (CB) block.

Model variant	Baseline	Wavelet	FPN	Wavelet gating	CB block	Accuracy (%)
1	✓	×	×	×	×	92.19
2	×	✓	×	×	×	82.56
3	✓	×	✓	×	×	92.69
4	✓	✓	✓	✓	×	94.68
5	✓	✓	✓	✓	✓	**95.18**

Table 6
Comparison of different feature attention and gating strategies integrated within the proposed FPN-based framework for multiscale feature selection and fusion.

Feature enhancement strategy	Accuracy (%)
Direct concatenation (No Enhancement)	92.69
SE Block	92.29
CBAM	94.02
Wavelet-guided Gating (Proposed)	**94.68**

feature representations, the proposed approach utilises wavelet-based frequency-domain features to guide the gating process. This enables explicit modelling of disease-relevant texture patterns and multiscale lesion characteristics, thereby facilitating more informative and context-aware branch selection. The results demonstrate that incorporating frequency-aware guidance improves the selection and emphasis of the most discriminative multiscale features, leading to enhanced classification performance over conventional attention-based approaches.

3.2.3. Placement of channel attention in FPN

In the proposed FPN module, channel attention is applied between feature maps to enhance feature representations. To validate the effectiveness of channel attention placement within the FPN module, an ablation study was conducted, and the corresponding results are presented in Table 7. Based on the results, it can be clearly observed that incorporating channel attention within the FPN improves the overall performance. Furthermore, applying channel attention to both the lateral and top-down pathways jointly yields better performance than applying attention to either pathway individually, achieving an overall accuracy of 95.18%. These results confirm that dual-path channel attention is most effective for capturing fine-grained lesion patterns and colour variations, and is therefore adopted in the proposed model.

3.2.4. Evaluation of wavelet decomposition configurations

The impact of Haar wavelet decomposition settings, including padding strategy, stride configuration, and learnable parameters, is analysed in Table 8. The results show that symmetric padding combined with stride two and trainable parameters (α, β) achieves the

Table 7
Effect of channel attention placement within the top-down and lateral pathways of the Feature Pyramid Network.

Attention configuration	Accuracy (%)
Without attention	94.52
Top-down attention only	94.68
Lateral attention only	94.85
Combined attention (Lateral + Top-down)	**95.18**

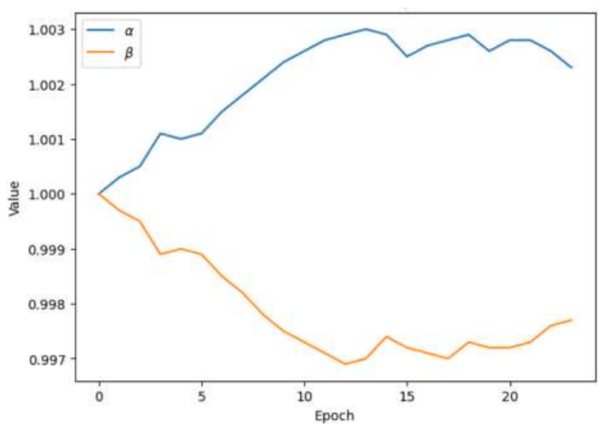

Fig. 9. Evolution of the learnable wavelet parameters α (LL) and β (LH, HL, HH) across training epochs, demonstrating their dynamic adaptation during optimisation.

Table 8
Effect of wavelet decomposition settings on model performance. Sym denotes symmetric padding, while Zero denotes zero padding; S represents the stride value; and Train indicates whether the parameters (α, β) are trainable.

Sym	Zero	S = 1	S = 2	Train	Accuracy (%)
✓	×	×	✓	×	93.52
✓	×	×	✓	✓	**95.18**
×	✓	×	✓	✓	94.51
✓	×	✓	×	✓	93.52

Table 9
Performance comparison of EfficientNet backbone models. #Params denotes the number of model parameters.

Backbone	#Params (M)	Accuracy (%)
EfficientNetB0	7.0	94.02
EfficientNetB1	9.5	93.20
EfficientNetV2B1	10.0	93.36
EfficientNetV2B2	11.8	93.02
EfficientNetV2B3	16.1	92.85
EfficientNetV2-M	56.5	**95.18**
EfficientNetV2-L	121.2	92.70

highest accuracy of 95.18%, indicating effective boundary preservation and multiscale feature representation. In contrast, zero padding and the exclusion of learnable parameters or stride two reduce the performance to 94.51% and 93.52%, respectively. As illustrated in Fig. 9, α and β remain close to unity while adapting dynamically during training, where α exhibits a slight increasing trend that emphasises low-frequency global structural information, whereas β shows a slight decreasing trend, indicating comparatively lower emphasis on high-frequency texture details and lesion boundaries. The complementary adaptation of α and β suggests that the model effectively balances global and fine-grained information during optimisation, thereby supporting robust multiscale feature representation.

Table 10
Performance comparison of the proposed model under different dynamic branch selection settings.

Branch Configuration	Accuracy (%)
Single branch (K = 1)	92.69
Dual branches (K = 2)	92.73
Triple branches (K = 3)	**95.18**
Quad branches (K = 4)	91.19
Five branches (K = 5)	92.52
Six branches (K = 6)	94.35

Table 11
Hyperparameter tuning results for learning rate and batch size.

Parameter	Value	Accuracy (%)
Learning rate	0.0001	94.18
	0.0005	**95.18**
	0.00001	93.36
Batch size	16	94.02
	32	**95.18**
	64	92.52

3.2.5. Comparison of backbone models

To achieve an optimal balance between classification accuracy and computational efficiency, several EfficientNet backbone models were evaluated within the same experimental setup, and the results are presented in Table 9. The findings show that lighter models such as EfficientNetB0 achieve competitive performance (94.02%) with relatively fewer parameters, while increasing model complexity does not consistently lead to improved accuracy. Among all variants, EfficientNetV2-M achieves the highest accuracy of 95.18%, demonstrating its superior capability in capturing complex disease patterns. In contrast, larger models such as EfficientNetV2-L exhibit reduced performance despite significantly higher parameter counts, indicating potential overfitting and reduced generalisation. Based on these results, EfficientNetV2-M is selected as the backbone for the proposed framework due to its favourable trade-off between accuracy and model complexity.

3.2.6. Selection of dynamic branches

The performance of the proposed approach also depends on the selection of the top-K feature pyramid branches. To analyse the effectiveness of the proposed approach under different K values, an additional experiment was conducted, and the results are presented in Table 10. The results show that a single branch achieves an accuracy of 92.69%, whereas the highest accuracy of 95.18% is obtained when three branches ($k = 3$) are selected, indicating an effective balance between multiscale feature representation and discriminative capability. Configurations with a larger number of branches result in reduced performance, which can be attributed to increased feature redundancy and diminished selection efficiency. Although a slight improvement is observed at $k = 6$, the performance remains lower than that of the $k = 3$ configuration.

3.2.7. Effect of hyperparameters on model performance

To evaluate the sensitivity of the proposed approach, an additional ablation study was conducted by varying the learning rates and batch sizes, and the corresponding results are presented in Table 11. The experimental results indicate that a moderate learning rate of 0.0005 achieves the most stable and effective optimisation performance by avoiding both slow convergence and unstable parameter updates. Similarly, a batch size of 32 provides the best balance between generalisation capability and training stability, whereas smaller and larger batch sizes result in comparatively lower performance.

Table 12
Comparison of lightweight backbone models for knowledge distillation based model development. The accuracies reported under Haar Wavelet-FPN correspond to the performance of each backbone when used in the proposed model, while those under Haar Wavelet-FPN w/KD represent the performance of the same backbones when trained as student models during knowledge distillation and w/o KD indicates without knowledge distillation.

Backbone	Params (M)	Accuracy (%) of Haar Wavelet-FPN	
		w/o KD	w/ KD
EfficientNetB1	9.5	93.20	93.52
EfficientNetV2B1	10.0	93.36	94.02
EfficientNetB0	7.0	94.02	**95.02**

3.2.8. Compression of backbone models for knowledge distillation

In the second phase of the methodology, a lightweight variant of the proposed model was developed using the knowledge distillation technique (Beyer et al., 2022). For the KD variant, several lightweight backbone architectures were trained and evaluated, and the experimental results are presented in Table 12. During KD training of this ablation experiment, EfficientNetV2-M serves as the teacher model to guide the learning process. Based on the experimental results, EfficientNetB0 was selected as the backbone for the KD variant, achieving an accuracy of 95.02% while maintaining a lightweight design with only 7 million parameters. From Table 12, it is also observed that knowledge distillation improves the classification performance of all three models compared to their corresponding baseline versions within the proposed Haar wavelet-guided FPN model architecture.

3.2.9. Selection of knowledge distillation hyperparameters

The effectiveness of the knowledge distillation process is strongly influenced by the selection of key hyperparameters: the temperature parameter (τ) and the loss-balancing coefficient (α), as defined in Eq. (10). An ablation study was conducted to analyse their impact, and the results are presented as tuning curves in Fig. 10. To reduce computational complexity, hyperparameters were tuned sequentially by first optimising τ while fixing α, followed by tuning α with the optimal τ.

The results show that classification accuracy increases with τ up to a moderate value and then declines, with the best performance achieved at $\tau = 4$. Lower temperature values result in overly sharp probability distributions, limiting the transfer of soft knowledge, while higher values introduce excessive smoothing, reducing discriminative capability. Similarly, the parameter α controls the balance between ground-truth supervision and distillation guidance. The results indicate that $\alpha = 0.3$ achieves the highest accuracy, while larger values lead to performance degradation due to over-reliance on the distillation loss. Based on these observations, $\tau = 4$ and $\alpha = 0.3$ provide an optimal trade-off, resulting in the best classification performance.

3.2.10. Effectiveness of knowledge distillation technique

In this study, a lightweight variant of the Haar wavelet-guided FPN model was developed using the knowledge distillation technique (Beyer et al., 2022). During training, an EfficientNetV2-M based model served as the teacher to guide the learning process of an EfficientNetB0 based student model. To demonstrate the effectiveness of the knowledge distillation technique, an ablation study was conducted. First, the proposed Haar wavelet-guided FPN model was trained using EfficientNetB0 as the baseline. Subsequently, the EfficientNetB0 based model was trained as a student model using knowledge distillation, guided by an EfficientNetV2-M based teacher model. The classification performance of these two models was compared across each benchmark dataset. Based on the experimental results presented in Table 13, it is evident that training the baseline model using the knowledge distillation technique leads to improved performance.

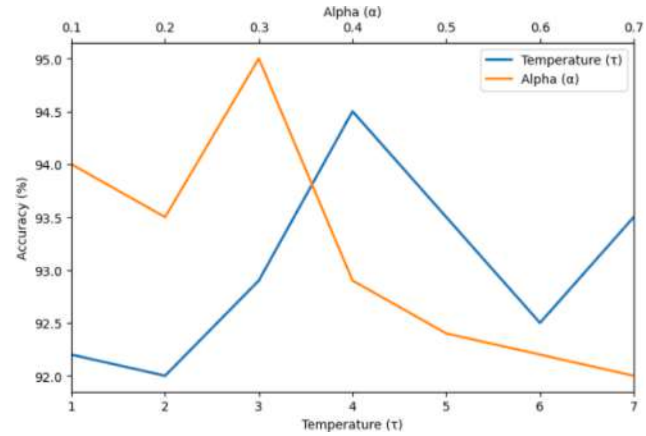

Fig. 10. Hyperparameter tuning results illustrating the impact of temperature (τ) and the loss-balancing coefficient (α) on accuracy. The bottom x-axis represents variations in temperature (τ) while keeping $\alpha = 0.5$ fixed, whereas the top x-axis represents variations in α with the temperature fixed at $\tau = 4$.

Table 13
Performance Comparison of EfficientNetB0 student backbone with and without Knowledge Distillation in the Haar wavelet-guided FPN framework across multiple datasets. # Classes denotes number of classes in the datase and #Images denotes total number of images in the dataset.

Dataset	#Classes	#Images	Accuracy (%)	
			w/o KD	w/KD
NP-LankaPaddy	4	2002	94.02	**95.02**
RiceLeafsv3	6	2710	95.21	**97.24**
Paddy Doctor	13	16 225	96.76	**97.84**
UCI-Rice leaf diseases	3	120	94.44	**97.22**
Mendeley-Rice leaf disease image samples	4	5932	**100**	**100**

4. Discussion

This study presents a novel model for paddy leaf disease classification. The proposed Haar wavelet-guided FPN architecture addresses the challenge of capturing multiscale disease symptoms. Unlike conventional deep learning approaches that rely on a single inference pathway, the proposed architecture dynamically selects and routes features based on the characteristics of each input image, enabling more effective handling of disease symptoms that vary in scale, texture, and severity. As shown in Fig. 11, the proposed Haar wavelet-guided FPN model successfully captures challenging disease patterns through the image-specific dynamic feature selection mechanism, whereas conventional CNNs fail to predict the correct classes.

As illustrated in Fig. 12, the feature visualisations provide further evidence of the effectiveness of the proposed multiscale feature fusion strategy. The different FPN levels capture complementary information across multiple scales through the top-down pathway. High-level features capture broader semantic regions corresponding to disease-affected areas, while mid-level features emphasise localised lesion patterns. In contrast, lower-level features retain fine-grained textures and subtle discolourations that are critical for early disease detection. This hierarchical representation demonstrates how the proposed framework effectively integrates global contextual information with detailed local features, thereby improving its capability to distinguish visually similar disease patterns.

In addition to achieving outstanding accuracy, the computational efficiency of the proposed model is further enhanced by developing a lightweight variant using a knowledge distillation strategy. This approach reduces the parameter count by approximately 89% and lowers

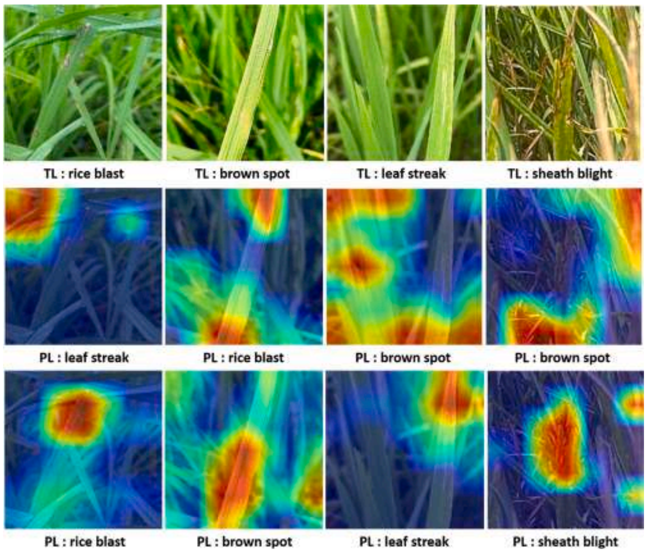

Fig. 11. Comparison of predicted labels and class activation heatmaps generated by the conventional EfficientNetB0 and the proposed Haar wavelet-guided FPN model (which also uses EfficientNetB0 as the backbone). The first row shows the raw input images with ground-truth labels. The second row presents the predicted labels and corresponding heatmaps produced by EfficientNetB0. The last row illustrates the predicted labels and corresponding heatmaps obtained using the proposed Haar wavelet-guided FPN model.

the GFLOPs by nearly 68%, resulting in a substantially more efficient model.

The development of a benchmark dataset collected within the Sri Lankan context represents another notable contribution of this work, as no prior initiative has specifically addressed the challenges faced by Sri Lankan farmers. Since the dataset and the code associated with this work are publicly available, they are expected to encourage future researchers working in this domain.

Based on the experimental results, brown spot disease is more challenging to classify than the other classes due to its high intra-class variability in lesion shape, size, and spatial distribution across different paddy varieties and growth conditions. The confusion matrix and precision–recall curves presented in Fig. 8 show that brown spot and leaf streak exhibit comparatively lower accuracy and average precision due to minor misclassifications and reduced precision consistency across varying recall levels. These observations highlight the difficulty in discriminating diseases with overlapping visual patterns. The Grad-CAM visualisations in Fig. 13 suggest that these errors arise from overlapping visual characteristics. While leaf streak typically presents elongated yellowish lesions, the model sometimes focuses on localised high-intensity brown regions instead of capturing the full elongated structure, leading to confusion with brown spot. This indicates that capturing global structural information is essential for distinguishing diseases with subtle visual differences. However, the heatmaps reveal that the model does not always effectively capture this global structure, suggesting that further improvement in global feature representation is required for more robust classification.

Despite its strong performance, the proposed framework has a few limitations. The integration of wavelet-guided gating and multiscale feature fusion introduces additional computational overhead compared to conventional lightweight CNN architectures. Although this overhead is substantially mitigated through model compression, further optimisation is required for deployment on ultra-low-power edge devices. In addition, the current framework assumes a single disease per image and does not explicitly address scenarios where multiple diseases co-occur on the same leaf, which is common in real-world agricultural environments. Furthermore, the high intra-class variability observed in certain

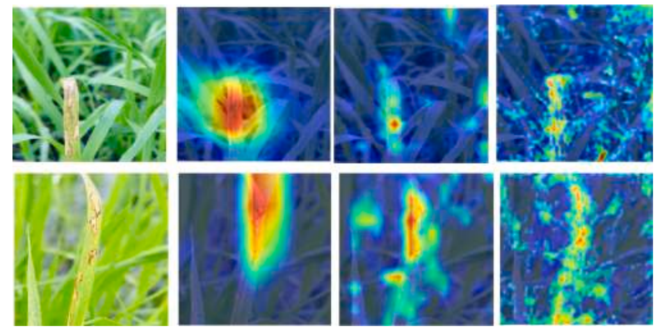

Fig. 12. Visualisation of multi-scale feature representations extracted from different levels of the FPN. Each row shows the input image (left), followed by the feature responses from the high-level (P7, as denoted in Fig. 2), mid-level (P4), and low-level (P2) layers, illustrating how features at different levels capture distinct disease-related patterns.

Fig. 13. Misclassified samples with corresponding Grad-CAM heatmaps. The first row shows a brown spot sample misclassified as leaf streak, while the second row shows a leaf streak sample misclassified as brown spot . The heatmaps indicate that the model focuses on localised lesion regions with overlapping visual characteristics.

diseases, such as brown spot across different paddy varieties, remains a challenging factor. Future work will focus on lightweight architectural refinement, multilabel disease recognition, and improved robustness to environmental variations. In particular, lightweight backbone networks such as MobileNetV4 will be explored within the proposed framework to further investigate their effectiveness for efficient paddy disease recognition, especially in terms of classification performance, computational complexity, and real-time deployment capability.

5. Conclusion

This study presents a novel adaptation of the Haar wavelet guided FPN framework that integrates Haar wavelet-guided feature gating and knowledge distillation to achieve an effective balance between classification accuracy and computational efficiency for paddy leaf disease classification. The proposed architecture demonstrates strong and consistent performance across multiple benchmarks, including the newly introduced NP-LankaPaddy dataset and established public datasets, highlighting its robustness under both controlled and real world field conditions. Adaptive feature selection through wavelet-guided gating enables dynamic prioritisation of multilevel features, while knowledge distillation facilitates an approximately 89% reduction in model complexity without compromising the performance. These results confirm the framework's suitability for deployment on resource constrained agricultural platforms. Future work will explore further architectural optimisation using lightweight components and extend the framework to multi disease and multilabel classification scenarios, addressing real world cases where multiple diseases may co-occur on a single leaf through the use of advanced multilabel learning strategies and enriched annotated datasets.

CRediT authorship contribution statement

Kiriharan Tharmini: Writing – original draft, Visualization, Validation, Methodology, Investigation, Data curation, Conceptualization. **Thanikasalam Kokul:** Writing – original draft, Supervision, Investigation, Conceptualization. **Amirthalingam Ramanan:** Writing – review & editing, Supervision, Investigation, Data curation, Conceptualization. **Terensan Suvanthini:** Writing – review & editing, Data curation. **Subha Fernando:** Writing – review & editing, Supervision, Methodology, Investigation, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

Part of the NP-LankaPaddy dataset collection was funded under research grant No. AHEAD/RA3/UBL/JFN/OVAA/44 obtained from the UBL, University of Jaffna, Sri Lanka.

Data availability

The code and dataset associated with this study are publicly available at: https://doi.org/10.5281/zenodo.18334503.

References

Aggarwal, Meenakshi, Khullar, Vikas, Goyal, Nitin, Prola, Thomas André, 2024. Resource-efficient federated learning over IoAT for rice leaf disease classification. Comput. Electron. Agric. (ISSN: 0168-1699) 221, 109001. http://dx.doi.org/10.1016/j.compag.2024.109001.

Ahad, Md. Taimur, Li, Yan, Song, Bo, Bhuiyan, Touhid, 2023. Comparison of CNN-based deep learning architectures for rice diseases classification. Artif. Intell. Agric. (ISSN: 2589-7217) 9, 22–35. http://dx.doi.org/10.1016/j.aiia.2023.07.001.

Al-Gaashani, Mehdhar S.A.M., Samee, Nagwan Abdel, Alnashwan, Rana, Khayyat, Mashael, Muthanna, Mohammed Saleh Ali, 2023. Using a ResNet50 with a kernel attention mechanism for rice disease diagnosis. Life 13 (6), 1277. http://dx.doi.org/10.3390/life13061277.

Alsakar, Yasmin M., Sakr, Nehal A., Elmogy, Mohammed, 2024. An enhanced classification system of various rice plant diseases based on multi-level handcrafted feature extraction technique. Sci. Rep. (ISSN: 2045-2322) 14 (1), 30601. http://dx.doi.org/10.1038/s41598-024-81143-1.

Bala Ayyappan, A., Gobinath, T., Kumar, M., Sivaramakrishnan, A., 2025. Rice plant disease detection using convolutional neural networks. Discov. Artif. Intell. 5 (1), 50. http://dx.doi.org/10.1007/s44163-025-00277-x.

Bari, Bifta Sama, Islam, Md. Nahidul, Rashid, Mamunur, Hasan, Md. Jahid, Razman, Mohd Azraai Mohd, Musa, Rabiu Muazu, Nasir, Ahmad Fakhri Ab, Abdul Majeed, Anwar P.P., 2021. A real-time approach of diagnosing rice leaf disease using deep learning-based faster R-CNN framework. PeerJ Comput. Sci. 7, e432. http://dx.doi.org/10.7717/peerj-cs.432.

Benita, D. Shirley, Anitha, J., Immanuel Alex, S., 2023. Investigation on leaf disease diagnosis in rice plant using machine learning approaches. In: 2023 International Conference on Sustainable Computing and Smart Systems. ICSCSS, pp. 350–355. http://dx.doi.org/10.1109/ICSCSS57650.2023.10169616.

Beyer, Lucas, Zhai, Xiaohua, Royer, Amélie, Markeeva, Larisa, Anil, Rohan, Kolesnikov, Alexander, 2022. Knowledge distillation: A good teacher is patient and consistent. In: 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition. CVPR, pp. 10915–10924. http://dx.doi.org/10.1109/CVPR52688.2022.01065.

Bharanidharan, N., Sannasi Chakravarthy, S.R., Rajaguru, Harikumar, Vinoth Kumar, V., Mahesh, T.R., Guluwadi, Suresh, 2023. Multiclass paddy disease detection using filter-based feature transformation technique. IEEE Access 11, 109477–109487. http://dx.doi.org/10.1109/ACCESS.2023.3322587.

Cai, Xiongjiang, Zhu, Yun, Liu, Shuwen, Yu, Zhiyue, Xu, Youyun, 2024. FastSegFormer: A knowledge distillation-based method for real-time semantic segmentation of surface defects in navel oranges. Comput. Electron. Agric. (ISSN: 0168-1699) 217, 108604. http://dx.doi.org/10.1016/j.compag.2023.108604.

Chawathe, Sudarshan, S., 2020. Rice disease detection by image analysis. In: 2020 10th Annual Computing and Communication Workshop and Conference. CCWC, pp. 0524–0530. http://dx.doi.org/10.1109/CCWC47524.2020.9031140.

Chen, Junde, Zhang, Defu, Nanehkaran, Yaser A., Li, Dele, 2020. Detection of rice plant diseases based on deep transfer learning. J. Sci. Food Agric. 100 (7), 3246–3256. http://dx.doi.org/10.1002/jsfa.10365.

de S. Seneviratne, S.N., Jeyanandarajah, Premala, 2004. Rice diseases - problems and progress. Trop. Agric. Res. Ext. 7, 29–48. http://dx.doi.org/10.4038/tare.v7i0.5416.

Department of Agriculture, Sri Lanka, 2024. Introduction to rice. Available at: https://doa.gov.lk/rrdi_rice_introduction-2. (Accessed 19 February 2025).

Department of Agriculture, Sri Lanka, 2025. Common rice diseases. Available at: https://doa.gov.lk/rrdi_commonricediseases. (Accessed 19 February 2025).

Dubey, Ratnesh Kumar, Choubey, Dilip Kumar, 2024. Adaptive feature selection with deep learning MBi-LSTM model based paddy plant leaf disease classification. Multimedia Tools Appl. 83 (9), 25543–25571. http://dx.doi.org/10.1007/s11042-023-16475-7.

Dubey, Ratnesh Kumar, Choubey, Dilip Kumar, 2025. Classification of paddy plant leaf diseases using optimized support vector machine. Procedia Comput. Sci. (ISSN: 1877-0509) 258, 1619–1628. http://dx.doi.org/10.1016/j.procs.2025.04.393.

Food and Agriculture Organization of the United Nations (FAO), 2025. FAO food outlook: Global output of key food commodity crops on course for new records. Available at: https://www.fao.org/home/en. (Accessed 19 February 2025).

Ganesan, Gangadevi, Chinnappan, Jayakumar, 2022. Hybridization of resnet with yolo classifier for automated paddy leaf disease recognition: An optimized model. J. Field Robot. 39 (7), 1085–1109. http://dx.doi.org/10.1002/rob.22089.

Hasan, Md. Mehedi, Rahman, Touficur, Shahab Uddin, A.F.M., Galib, Syed Md., Akhond, Mostafijur Rahman, Uddin, Md. Jashim, Hossain, Md. Alam, 2023. Enhancing rice crop management: Disease classification using convolutional neural networks and mobile application integration. Agriculture 13 (8), 1549. http://dx.doi.org/10.3390/agriculture13081549.

Jain, Shaveta, Kumar, Rajneesh, 2024. An efficient deep learning model for detection of rice leaf diseases using paddy doctor dataset. In: 2024 First International Conference on Software, Systems and Information Technology. SSITCON, IEEE, pp. 1–7. http://dx.doi.org/10.1109/SSITCON62437.2024.10797004.

Jiang, Minlan, Feng, Changguang, Fang, Xiaosheng, Huang, Qi, Zhang, Changjiang, Shi, Xiaowei, 2023. Rice disease identification method based on attention mechanism and deep dense network. Electronics 12 (3), http://dx.doi.org/10.3390/electronics12030508.

Jiang, Wenbo, Liu, Lirui, Bu, Yun, 2025. Application of deep learning in the phase processing of digital holographic microscopy. In: Photonics. Vol. 12, http://dx.doi.org/10.3390/photonics12080810.

Jiang, Feng, Lu, Yang, Chen, Yu, Cai, Di, Li, Gongfa, 2020. Image recognition of four rice leaf diseases based on deep learning and support vector machine. Comput. Electron. Agric. (ISSN: 0168-1699) 179, 105824. http://dx.doi.org/10.1016/j.compag.2020.105824.

Jithya, K.V., Delita, Tauro, Pruthvi, M.R., Carol, D. Souza, Brtiney, Correia, 2024. Paddy care: Paddy disease identification and classification using deep DenseNet network. In: Proceedings of the 2024 IEEE International Conference on Distributed Computing, VLSI, Electrical Circuits and Robotics. DISCOVER, IEEE, pp. 377–382. http://dx.doi.org/10.1109/DISCOVER62353.2024.10750707.

Kaur, Arshleen, Kukreja, Vinay, Tiwari, Pallavi, Manwal, Manika, Sharma, Rishabh, 2024. An efficient deep learning-based VGG19 approach for rice leaf disease classification. In: 2024 IEEE 9th International Conference for Convergence in Technology. I2CT, pp. 1–6. http://dx.doi.org/10.1109/I2CT61223.2024.10544229.

Kaviya, P., Selvakumar, B., 2023. Multimodal paddy leaf diseases detection using feature extraction and machine learning techniques. In: Advances in Distributed Computing and Machine Learning. Springer, pp. 499–508. http://dx.doi.org/10.1007/978-981-99-1203-2_42.

Krishnamoorthy, N., Narasimha Prasad, L.V., Pavan Kumar, C.S., Subedi, Bharat, Abraha, Haftom Baraki, Sathishkumar, V.E., 2021. Rice leaf diseases prediction using deep neural networks with transfer learning. Environ. Res. (ISSN: 0013-9351) 198, 111275. http://dx.doi.org/10.1016/j.envres.2021.111275.

Kumar, B. Naresh, Sakthivel, S., 2025. Rice leaf disease classification using a fusion vision approach. Sci. Rep. (ISSN: 2045-2322) 15 (1), 8692. http://dx.doi.org/10.1038/s41598-025-87800-3.

Kumar, Abhishek, Shibu, Athul, Lee, Dong-Gyu, 2025. Automatic network compression via evolutionary search-guided channel pruning of deep convolutional neural networks. Appl. Soft Comput. 182, 113521. http://dx.doi.org/10.1016/j.asoc.2025.113521.

Kumar, D. Girish., Subbarao, M. Venkata., Sai Pratibha, M., Swarna, M., Varshini, K., Prasanthi, N. Divya., 2024. Comparative analysis of deep learning architectures and optimizers for paddy leaf disease classification. In: 2024 International Conference on Integrated Circuits and Communication Systems. ICICACS, pp. 1–5. http://dx.doi.org/10.1109/ICICACS60521.2024.10498645.

Kumar, Ajitesh, Yadav, Dhirendra Prasad, Kumar, Deepak, Pant, Manu, Pant, Gaurav, 2023. Multi-scale feature fusion-based lightweight dual stream transformer for detection of paddy leaf disease. Environ. Monit. Assess. 195 (9), 1020. http://dx.doi.org/10.1007/s10661-023-11628-5.

Latif, Ghazanfar, Abdelhamid, Sherif E., Mallouhy, Roxane Elias, Alghazo, Jaafar, Kazimi, Zafar Abbas, 2022. Deep learning utilization in agriculture: Detection of rice plant diseases using an improved CNN model. Plants 11 (17), 2230. http://dx.doi.org/10.3390/plants11172230.

Li, Lingxiao, Liu, Linlin, Cheng, Fuhao, He, Yunan, Zhong, Zhuqiang, 2026. Cn-unet: Convnext unet with slicing-aided hyper segmentation for infrared small target detection. IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens. 19, 84–98. http://dx.doi.org/10.1109/JSTARS.2025.3631384.

Li, Dengshan, Wang, Rujing, Xie, Chengjun, Liu, Liu, Zhang, Jie, Li, Rui, Wang, Fangyuan, Zhou, Man, Liu, Wancai, 2020. A recognition method for rice plant diseases and pests video detection based on deep convolutional neural network. Sensors 20 (3), http://dx.doi.org/10.3390/s20030578.

Lu, Yang, Zhang, Xinmeng, Zeng, Nianyin, Liu, Wanting, Shang, Rou, 2022. Image classification and identification for rice leaf diseases based on improved WOACW_SimpleNet. Front. Plant Sci. 13, 1008819. http://dx.doi.org/10.3389/fpls.2022.1008819.

Luo, Yuexuan, Cai, Xiang, Qi, Jiandong, Guo, Dongdong, Che, Wenqing, 2023. FPGA–accelerated CNN for real-time plant disease identification. Comput. Electron. Agric. (ISSN: 0168-1699) 207, 107715. http://dx.doi.org/10.1016/j.compag.2023.107715.

Mahadevan, K., Punitha, A., Suresh, J., 2024. Automatic recognition of rice plant leaf diseases detection using deep neural network with improved threshold neural network. e-Prime-Adv. Electr. Eng. Electron. Energy 8, 100534. http://dx.doi.org/10.1016/j.prime.2024.100534.

Mallat, Stéphane, 1989. A theory for multiresolution signal decomposition: the wavelet representation. IEEE Trans. Pattern Anal. Mach. Intell. 11 (7), 674–693. http://dx.doi.org/10.1109/34.192463.

Mekha, Panuwat, Teeyasuksaet, Nutnicha, 2021. Image classification of rice leaf diseases using random forest algorithm. In: 2021 Joint International Conference on Digital Arts, Media and Technology with ECTI Northern Section Conference on Electrical, Electronics, Computer and Telecommunication Engineering. IEEE, pp. 165–169. http://dx.doi.org/10.1109/ECTIDAMTNCON51128.2021.9425696.

Nanda, Tanmay Rai, Shukla, Ananya, Srinivasa, Tanay Raghunandan, Bhargava, Jia, Chauhan, Sunita, 2025. Advancing real-time crop disease detection on edge computing devices using lightweight convolutional neural networks. In: Arai, Kohei (Ed.), Intelligent Systems and Applications. Springer Nature Switzerland, ISBN: 978-3-032-00071-2, pp. 537–552. http://dx.doi.org/10.1007/978-3-032-00071-2_33.

Navaneetha, D., Jadala, Vijaya Chandra, Murugamani, C., 2025. Paddy leaf disease classification using Kullback Leiberg Divergence loss based convolutional neural network. In: 2025 3rd International Conference on Data Science and Network Security. ICDSNS, pp. 1–7. http://dx.doi.org/10.1109/ICDSNS65743.2025.11168778.

Padhi, Jagamohan, Korada, Laxminarayana, Dash, Ashis, Sethy, Prabira, Behera, Santi Kumari, Nanthaamornphong, Aziz, 2024. Paddy leaf disease classification using EfficientNet B4 with compound scaling and swish activation: A deep learning approach. IEEE Access 12, 126426–126437. http://dx.doi.org/10.1109/ACCESS.2024.3451557.

Padhi, Jagamohan, Mishra, Kunal, Ratha, Ashoka Kumar, Behera, Santi Kumari, Sethy, Prabira Kumar, Nanthaamornphong, Aziz, 2025. Enhancing paddy leaf disease diagnosis -a hybrid CNN model using simulated thermal imaging. Smart Agric. Technol. (ISSN: 2772-3755) 10, 100814. http://dx.doi.org/10.1016/j.atech.2025.100814.

Pan, Ying-Hua, Nong, Bao-Xuan, Chen, Lei, Yang, Xing-Hai, Xia, Xiu-Zhong, Zhang, Zong-Qiong, Qing, Dong-Jin, Gao, Ju, Huang, Cheng-Cui, Li, Dan-Ting, Deng, Guo-Fu, 2023. Qtl mapping and identification of candidate genes for cold tolerance at the germination stage in wild rice. Genes Genom. 45 (7), 867–885. http://dx.doi.org/10.1007/s13258-023-01394-6.

Pandi, S. Senthil., Deepak Kumar, K., Raja, K., Senthilselvi, A., 2024. Rice plant leaf disease classification using machine learning algorithm. In: 2024 2nd International Conference on Advancement in Computation and Computer Technologies. InCACCT, pp. 59–63. http://dx.doi.org/10.1109/InCACCT61598.2024.10551249.

Petchiammal, A., Briskline, S., Kiruba, Murugan, D., Arjunan, Pandarasamy, 2022. Paddy doctor: A visual image dataset for automated paddy disease classification and benchmarking.

Petchiammal, A., Kiruba, Briskline, Murugan, D., Arjunan, Pandarasamy, 2023. Paddy doctor: A visual image dataset for automated paddy disease classification and benchmarking. In: Proceedings of the 6th Joint International Conference on Data Science & Management of Data. Association for Computing Machinery, New York, NY, USA, pp. 203–207. http://dx.doi.org/10.1145/3570991.3570994.

Petchiammal, A., Murugan, D., 2025. Automated paddy leaf disease identification using visual leaf images based on nine pre-trained models approach. Procedia Comput. Sci. (ISSN: 1877-0509) 252, 118–126. http://dx.doi.org/10.1016/j.procs.2024.12.013.

Pothen, Minu Eliz, Pai, Maya L., 2020. Detection of rice leaf diseases using image processing. In: 2020 Fourth International Conference on Computing Methodologies and Communication. ICCMC, IEEE, pp. 424–430. http://dx.doi.org/10.1109/ICCMC48092.2020.ICCMC-00080.

Prottasha, Md. Sazzadul Islam, Hossain, A., Rahman, Md., Reza, S.M. Salim, Hossain, Dilshad, 2022. Identification of various rice plant diseases using optimized convolutional neural network. IJCDS J. 12, 1539–1551. http://dx.doi.org/10.12785/ijcds/1201124.

Rahman, Chowdhury R., Arko, Preetom S., Ali, Mohammed E., Iqbal Khan, Mohammad A., Apon, Sajid H., Nowrin, Farzana, Wasif, Abu, 2020. Identification and recognition of rice diseases and pests using convolutional neural networks. Biosyst. Eng. 194, 112–120. http://dx.doi.org/10.1016/j.biosystemseng.2020.03.020.

Rakshit, Sudipta Vaskar, Muinul Islam Chowdhury, A.M.M., Chakma, Ayesh, Hasan, Imran, Roy, Swagata, Rahman, Taohidur, 2025. DeepLeafDiagnosis: A two-phase transfer learning strategy for perfect accuracy in rice leaf disease classification. In: 2025 International Conference on Quantum Photonics, Artificial Intelligence, and Networking. QPAIN, pp. 1–6. http://dx.doi.org/10.1109/QPAIN66474.2025.11172163.

Ranasinghe, N., Ramanan, Amirthalingam, Fernando, S., Hameed, P., Herath, Dinesh, Malepathirana, Tamasha, Suganthan, Ponnuthurai, Niranjan, M., Halgamuge, S., 2022. Interpretability and accessibility of machine learning in selected food processing, agriculture and health applications. J. Natl. Sci. Found. Sri Lanka 263–276. http://dx.doi.org/10.4038/jnsfsr.v50i0.11249.

Ritharson, P. Isaac, Raimond, Kumudha, Anitha Mary, X., Robert, Jennifer Eunice, Andrew, J., 2024. Deeprice: A deep learning and deep feature based classification of rice leaf disease subtypes. Artif. Intell. Agric. (ISSN: 2589-7217) 11, 34–49. http://dx.doi.org/10.1016/j.aiia.2023.11.001.

Salamai, Abdullah Ali, Ajabnoor, Nouran, Khalid, Waleed E., Ali, Mohammed Maqsood, Murayr, Abdulaziz Ali, 2023. Lesion-aware visual transformer network for paddy diseases detection in precision agriculture. Eur. J. Agron. 148, 126884. http://dx.doi.org/10.1016/j.eja.2023.126884.

Saputra, Dede Ikhsan Dwi, 2022. Rice Leafs Disease Dataset. Kaggle, Available at: https://www.kaggle.com/datasets/dedeikhsandwisaputra/rice-leafs-disease-dataset, (Accessed 19 February 2025).

Sethy, Prabira Kumar, 2024. Rice Leaf Disease Image Samples (Version 2). Mendeley Data, (Accessed 19 February 2025).

Shah, Jitesh, Prajapati, Harshadkumar, Dabhi, Vipul, 2017. Rice Leaf Diseases. UCI Machine Learning Repository, Available at: https://doi.org/10.24432/C5R013.

Shazeer, Noam, Fatahalian, Kayvon, Mark, William R., Mullapudi, Ravi Teja, 2018. Hydranets: Specialized dynamic architectures for efficient inference. In: 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8080–8089. http://dx.doi.org/10.1109/CVPR.2018.00843.

Singh, Amandeep, Singh, Kuldeep, Kaur, Jaspreet, Singh, Maninder Lal, 2023. Smart agriculture framework for automated detection of leaf blast disease in paddy crop using colour slicing and GLCM features based random forest approach. Wirel. Pers. Commun. (ISSN: 1572-834X) 131 (4), 2445–2462. http://dx.doi.org/10.1007/s11277-023-10545-7.

Song, Liangchen, Gong, Xuan, Zhou, Helong, Chen, Jiajie, Zhang, Qian, Doermann, David, Yuan, Junsong, 2023. Exploring the knowledge transferred by response-based teacher-student distillation. In: Proceedings of the ACM International Conference on Multimedia. Association for Computing Machinery, ISBN: 9798400701085, pp. 2704–2713. http://dx.doi.org/10.1145/3581783.3612162.

Sowmyalakshmi, R., Jayasankar, T., Pillai, V. Ayyem, Subramaniyan, Kamalraj, Irina, V., Pustokhina, Denis, Pustokhin, A., Shankar, K., 2021. An optimal classification model for rice plant disease detection. Comput. Mater. Contin. (ISSN: 1546-2226) 68 (2), 1751–1767. http://dx.doi.org/10.32604/cmc.2021.016825.

Subbarayudu, Chatla, Kubendiran, Mohan, 2025. An automated hybrid deep learning framework for paddy leaf disease identification and classification. Sci. Rep. 15 (1), 26873. http://dx.doi.org/10.1038/s41598-025-08071-6.

Swathika, R., Srinidhi, S., Radha, N., Sowmya, K., 2021. Disease identification in paddy leaves using CNN based deep learning. In: 2021 Third International Conference on Intelligent Communication Technologies and Virtual Mobile Networks. ICICV, pp. 1004–1008. http://dx.doi.org/10.1109/ICICV50876.2021.9388557.

Tan, Mingxing, Le, Quoc, 2021. EfficientNetv2:smaller models and faster training. In: Proceedings of the 38th International Conference on Machine Learning. pp. 10096–10106. http://dx.doi.org/10.48550/arXiv.2104.00298.

Tharmini, K., Ramanan, A., Suvanthini, T., Kokul, T., Fernando, S., 2026. Np-Lankapaddy Dataset (Version 1.0). Zenodo Repository.

Tjikoa, Ade Fiqri, 2022. Rice Leaf 5 Diseases. Kaggle, (Accessed 20 February 2026).

Uddin, Md. Azher, Joolee, Joolekha Bibi, 2024. Fusion of deep learned and handcrafted features for paddy disease recognition. In: 2024 International Conference on Artificial Intelligence, Metaverse and Cybersecurity. ICAMAC, pp. 1–6. http://dx.doi.org/10.1109/ICAMAC62387.2024.10829316.

Vandana, V. Mary Rajam, Vinod, Viji, 2025. A novel hybrid approach for paddy leaf disease severity classification using spatial bidirectional convolutional neural network classifier. Int. J. Innov. Stud. 9 (2), Available at: https://iji-studies.com/index.php/IJIS/article/view/419.

Vijayakumar, Senthilkumar, Kane, Shaunak Pai, Vaiappuri, Selvakkumar, Louis, Filious, Anantha Prabha, P., 2025. Quantum-enhanced hybrid CNN–QNNframework for intelligent rice disease diagnosis and real-time field recommendations on edge AI devices. In: 2025 IEEE 7th Symposium on Computers & Informatics. ISCI, pp. 358–366. http://dx.doi.org/10.1109/ISCI65687.2025.11167361.

Waghmare, Surendra, 2024. An advanced cnn based system for rice plant diseases detection and classification in precision agriculture. J. Electr. Syst. 20, 71–80. http://dx.doi.org/10.52783/jes.1823.

Wu, Yonghui, Long, Wenhan, Liu, Hao, Wang, Hanming, Li, Songquan, Zhu, Rongbo, 2025. FSLNet: Filter sensitivity-based lightweight network for rice leaf disease recognition. Comput. Electron. Agric. (ISSN: 0168-1699) 237, 110751. http://dx.doi.org/10.1016/j.compag.2025.110751.

Xie, Xiaoliang, Hu, Yihong, Li, Xialian, Li, Saijia, Li, Xiaoyu, Li, Ying, 2025. Measuring and enhancing food security resilience in china under climate change. Systems (ISSN: 2079-8954) 13 (12), http://dx.doi.org/10.3390/systems13121054.

Yakkundimath, Rajesh, Saunshi, Girish, Anami, Basavaraj, Palaiah, Surendra, 2022. Classification of rice diseases using convolutional neural network models. J. Inst. Eng. (India): Ser. B 103 (4), 1047–1059. http://dx.doi.org/10.1007/s40031-021-00704-4.

Yugander, Arra, Ladhalakshmi, Duraisamy, Prakasham, Vellaichamy, Mangrauthia, Satendra K., Prasad, Madamsetty S., Krishnaveni, Donempudi, Madhav, Maganti Sheshu, Sundaram, Raman M., Laha, Gouri Sankar, 2015. Pathogenic and genetic variation among the isolates of Rhizoctonia solani (AG 1-IA), the rice sheath blight pathogen. J. Phytopath. 163 (6), 465–474. http://dx.doi.org/10.1111/jph.12343.

Zeng, Nanxin, Gong, Gufeng, Zhou, Guoxiong, Hu, Can, 2023. An accurate classification of rice diseases based on ICAI-V4. Plants 12 (11), http://dx.doi.org/10.3390/plants12112225.

Zhou, Changjian, Zhong, Yujie, Zhou, Sihan, Song, Jia, Xiang, Wensheng, 2023. Rice leaf disease identification by residual-distilled transformer. Eng. Appl. Artif. Intell. (ISSN: 0952-1976) 121, 106020. http://dx.doi.org/10.1016/j.engappai.2023.106020.

Zhu, Dongqin, Feng, Quan, Zhang, Jianhua, Yang, Wanxia, 2022. Cotton disease identification method based on pruning. Front. Plant Sci. 13–2022, http://dx.doi.org/10.3389/fpls.2022.1038791.