



Original Research Paper

A Multi-Task Deep Learning Framework for Simultaneous Detection of Leaf Disease and Nutrient Deficiency in *Oryza sativa*

M. Kaliappan¹, S Vimal², Gaurav Dhiman^{3,4,5}, S. Pandiarajan⁶

¹Department of Artificial Intelligence and Data Science, Ramco Institute of Technology, Rajapalayam, Tamil Nadu 626117, India

²Department of Computer Science and Engineering, KIT- Kalaignar Karunanidhi Institute of Technology, Coimbatore, Tamil Nadu 641402 India

³Department of Computer Science and Engineering, Yuan Ze University, Tao Yuan, Taiwan

⁴University Centre for Research and Development, Chandigarh University, Mohali, India

⁵Division of Research and Development, Lovely Professional University, Phagwara, India

⁶Department of CSBS, KIT- Kalaignar Karunanidhi Institute of Technology, Coimbatore, Tamil Nadu 641402 India

Email: kaliappan@ritrjpm.ac.in, svimal@kitcbe.ac.in, gdhiman0001@gmail.com, kitpandiarajan@gmail.com

Abstract

Rice (*Oryza sativa*) constitutes the dietary staple for more than 3.5 billion people worldwide, yet annual crop losses attributable to foliar diseases and macro-nutrient deficiencies consistently erode global food security. Conventional diagnostic practice demands trained agronomists performing field visits, a resource-intensive and delay-prone process ill-suited to the temporal dynamics of infection progression and nutritional stress. This paper presents a novel multi-task deep learning framework that simultaneously (AgroMTL-Rice) identifies leaf diseases and nutrient deficiencies from a single RGB leaf image, thereby enabling dual-diagnosis within a single forward pass. The proposed architecture couples a pre-trained EfficientNetB0 feature extractor with a symptom-disentanglement-inspired dual-branch head that separates texture-oriented disease semantics from colour-physiology-oriented nutrient semantics. Two publicly available Kaggle datasets, the Rice Leaf Diseases dataset (bacterial leaf blight, brown spot, leaf smut) and the Nutrient Deficiency Symptoms in Rice dataset includes nitrogen, phosphorus, and potassium deficiency are integrated under a unified multi-task label construction scheme, enabling cross-domain knowledge transfer between the two diagnostic tasks. Experimental evaluation on a stratified hold-out test set yields a weighted-average F1-score of 0.98 for the disease branch and 0.92 for the nutrient branch, with macro-average AUC values reaching 1.00 for all disease classes and 0.98–1.00 for nutrient classes. Fine-tuning of the top 40 EfficientNetB0 layers further sharpens discriminative boundaries. Ablation studies confirm that the shared-projection layer reduces nutrient classification error by 14.3% relative to single-task baselines. AgroMTL-Rice demonstrates competitive or superior performance compared to recent state-of-the-art single-task convolutional and transformer-based methods while requiring a single model deployment, substantially reducing inference overhead in resource-constrained agricultural IoT scenarios.

Keywords: rice leaf disease detection, nutrient deficiency classification, multi-task learning, EfficientNet, transfer learning, precision agriculture, deep learning, convolutional neural networks.

*Corresponding Author:

Received: 25th July, 2026 | Revised: 26th August, 2026 | Accepted: 18th September, 2026 | Available Online: 24th September, 2026

(DOI): 10.70102/AEJ.2026.18.7s.49

1. Introduction

Rice is the most widely cultivated cereal crop in the world accounting for about 29% of the world cereal production and has been the major source of calories for the people of South Asia, Southeast Asia and Sub-Saharan Africa [1]. Biotic stressors, including fungal, bacterial, and viral pathogens, cause 10–40% of the rice harvest to be lost annually, according to the Food and Agriculture Organization, despite decades of agronomic advances [2]. At the same time, abiotic stressors, especially macro-nutrient deficiencies of nitrogen (N), phosphorus (P), and potassium (K), reduce both yield and grain quality, leading to estimated economic losses of more than USD 24 billion per year in Asian growing regions [3].

The visual symptoms of leaf diseases and nutrient deficiencies are expressed on foliar surfaces and thus are amenable to image-based automated diagnosis. However, the symptoms of some disease-deficiency co-occurrences are morphologically very similar. For example, nitrogen deficiency causes chlorosis that resembles the initial yellowing stage of bacterial leaf blight infection. This inter-class ambiguity leads to significant diagnostic burden even to experienced agronomists, motivating the need for computationally-rigorous, objective and scalable AI-assisted diagnostic tools.

Recently, deep convolutional neural networks (DCNNs) have been very successful in agricultural disease identification [4, 5]. The PlantVillage dataset [6] has generated a wave of transfer learning studies with 99%+ accuracy in controlled photographic conditions. Since then, researchers have explored lightweight architectures, attention mechanisms and self-supervised pre-training to improve performance under real-world field conditions [7, 8]. However, most published models treat disease identification and nutrient diagnosis as independent, siloed tasks, each requiring a dedicated model with separate data pipelines, training regimes, and deployment footprints.

Multi-task learning (MTL) is an attractive alternative paradigm. MTL models are able to leverage task correlations to improve generalisation, reduce overfitting on smaller task-specific datasets, and lower total computational cost at deployment by sharing representational backbone layers across related tasks [9]. In the context of rice foliar diagnostics, both disease and nutrient stress manifest in leaf tissue and share low-level visual features such as texture anomalies, pigment gradients and lesion morphology. The MTL framework is thus theoretically apt to extract both the signal types from a single image representation.

While promising, the development of an effective MTL framework for dual rice diagnosis faces several

non-trivial challenges. First, disease and nutrient datasets are often acquired from heterogeneous sources collected under various illumination conditions, background complexities and photographic protocols, leading to domain shift which can destabilize joint training. Second, when a given training image contains only disease labels or only nutrient labels (which is necessarily the case when fusing disjoint datasets), the model must learn to suppress spurious activations in the uninformative output head without being fooled by the implicit ‘none’ class. Third, task interference, where gradients from one task’s loss function adversely affect the parameter updates of another task, is still an open problem in the MTL literature [10, 11]. In the end, the combined dataset suffers from a serious class imbalance (images of the disease are more strongly labelled ‘none’ for the nutrient head, and vice versa), which requires careful weighting and augmentation strategies.

In this paper, we tackle the above challenges by designing, implementing and thoroughly evaluating the AgroMTL-Rice. The particular contribution of this work is the following:

1. We introduce a novel symptom disentanglement inspired dual-branch architecture with separate 256-dimensional feature projections for disease and nutrient tasks respectively, trained jointly on a unified multi-task label scheme.
2. We introduce a systematic dataset fusion protocol aligning heterogeneous image repositories from Kaggle under a consistent label hierarchy, allowing reproducible multi-task training without the need for manual image curation.
3. We conduct a comprehensive evaluation, including per-class F1 metrics, multi-class ROC-AUC analysis, visualization of confusion matrix, analysis of training curves, and ablation studies comparing single-task and multi-task configurations.
4. We evaluate AgroMTL-Rice against five recent state-of-the-art methods showing that the proposed model achieves competitive accuracy at significantly reduced per-task model footprint.
5. All training code is provided in a single-cell Google Colab notebook for full reproducibility.

The rest of the paper is structured as follows. Section 2 discusses related literature. Section 3 outlines the proposed methodology. 4. Results and discussion of the experiment Finally, Section 5 concludes the paper and presents some future directions.

2. Literature Review

The identification of plant diseases using deep convolutional neural networks has been rapidly developed, starting from the pioneering work of Mohanty et al. [6] in which a deep neural network was used to identify 26 diseases in 14 crop species with

99.35% accuracy in a controlled setting using the PlantVillage dataset. Later work has concentrated on improving generalization to field images which suffer from occlusion, variable illumination and background clutter [7].

AlexNet style models can achieve up to 99.53% accuracy. However the authors also highlight the laboratory bias of the dataset. Ferentinos [12] performed a systematic comparison of six CNN architectures on PlantVillage. [13] Brahim et al. applied deep learning to detect tomato disease and stressed the importance of data augmentation in solving the overfitting problem in small-sample cases, which are particularly relevant to agricultural applications. Islam et al. [14] studied transfer learning of VGG16 and InceptionV3 pre-trained models on ImageNet for rice blast detection and achieved an accuracy of 93.2% in the field images. Most recently, Chen et al. [15] investigated the pre-trained Vision Transformer (ViT) architectures with masked image modelling and showed that self-supervised pre-training greatly enhances the feature quality in the limited labelled data scenarios with a multi-class rice disease benchmark of 97.1% F1.

Attention mechanisms have been widely studied to guide the CNN feature extraction to the symptom carrying leaf areas and suppress the background noise. Zheng et al. [16] proposed channel-spatial attention modules in ResNet-50 and increased the accuracy of wheat disease classification from 91.3% to 96.7%. Similarly, Jiang et al. [17] proposed a dual-attention network for maize leaf disease identification, in which the spatial regions and spectral channels are highlighted respectively, achieving 98.1% macro-F1 on a six-class benchmark. Together these results demonstrate that architectural inductive biases towards symptom-region focus yield large accuracy gains.

Rice disease identification has attracted focused research activity, with potential consequences for global food security. Sethy et al. [18] benchmarked twelve CNN models on a custom rice disease dataset (bacterial leaf blight, brown spot and leaf smut) — the same classes as in the current study — and found EfficientNetB0 to be the most efficient model per parameter, achieving 97.8% accuracy with only 4.0M parameters, outperforming larger ResNet and VGG models. Prajapati et al. [19] developed MobileNetV2 based classifier which can be deployed in mobile devices with model size of 3.2 MB and 94.5% accuracy. This shows the feasibility of on-device rice disease identification. Ramesh et al. [20] extended this line to a federated learning framework that enables model training on distributed edge devices at agricultural stations without centralised data

aggregation, and reported 91.3% average accuracy with provable privacy guarantees.

Lu et al. [21] demonstrated that multi-scale feature aggregation using feature pyramid networks (FPN) can significantly improve the detection of small lesion regions in high-resolution leaf images, increasing the mean average precision (mAP) from 78.3% to 89.6% on a rice disease object detection benchmark. Guo et al. [22] pre-trained on contrastive learning on unlabelled rice field images and then fine-tuned on labelled disease data, leading to an 8.2% macro-F1 improvement of generalisation on unseen geographical growing regions. In conclusion, the current best-practice toolkit for rice disease CNNs consists of EfficientNet-family backbones, attention modules, multi-scale aggregation, and self-supervised pre-training, according to these studies.

In contrast to disease detection, automated detection of nutrient deficiency in crops using leaf images is a less explored area. Patel et al. [23] trained a ResNet-34 classifier to differentiate between N, P and K deficiency symptoms in rice, reporting 88.4% weighted-average accuracy on a benchmark of 1,200 images. They identified potassium deficiency as the most difficult class to identify due to inter-class visual similarity with iron deficiency and senescent yellowing. To detect nitrogen deficiency, Khan et al. [24] used hyperspectral imaging and a 1D-CNN spectral classifier with 95.1% accuracy. However, they stated that the hyperspectral acquisition hardware is too expensive for smallholder farmers to use.

A comprehensive review of image processing approaches to nutrient deficiency assessment has been undertaken by Barbedo [25] who concludes that RGB-based methods are preferable for large-scale deployment due to low hardware cost and that the principal limiting factor is the availability of large, geographically diverse and consistently labelled training datasets. Rumpf et al. [26] showed that early nutrient stress, before visible symptoms appear, can be detected by means of vegetation indices derived from multispectral UAV images and support vector machine classifiers, albeit with the need for specialised unmanned aerial vehicle platforms. Pandey et al. [27] proposed a lightweight ShuffleNetV2 model for on-device N-P-K deficiency diagnosis with 91.7% accuracy and only 1.4 MB model size, which is suitable to be deployed on sub-USD-50 IoT devices.

In agricultural image analysis, multi-task learning has been used for simultaneous pest and disease detection [28], prediction of crop type and health status [29] and combined disease severity grading and disease identification [30]. Nguyen et al. [28] trained a two-head CNN on a grape leaf dataset for the joint tasks of disease identity classification and

disease severity estimation. They demonstrated that the joint model improved the performance of 2.1 and 3.4 percentage points compared to independently trained single-task networks, respectively, which is consistent with the theoretical expectation that related tasks provide mutually beneficial regularization.

Liu et al. [29] used MTL to simultaneously identify crop species and classify stress types with satellite images. Their shared feature extraction layers reduced the total number of parameters by 37% compared to two separate models without loss of per-task accuracy. We propose a multi-task learning (MTL) optimization technique, named gradient surgery, which projects the conflicting inter-task gradients onto their orthogonal complement to remove the destructive interference, and apply it to the joint classification of plant disease severity and type, achieving a 1.9 percentage point increase in macro-F1 over naive gradient summation. [30] These results support the use of MTL for the dual diagnosis scenario addressed in this work.

Wang et al. [31] proposed a transformer-based MTL framework for agricultural image analysis by adding task-specific adapter modules into a shared ViT-B/16 backbone, achieving state-of-the-art performance for three concurrent plant health assessment tasks. Deng et al. [32] investigated strategies for task groupings in agricultural MTL and showed that task pairs with visual correlation (e.g., disease identity + lesion segmentation) exhibit larger accuracy gains over less correlated task pairs (e.g., disease identity + plant species). The result directly supports the design decision of combining disease identification and nutrient deficiency detection in AgroMTL-Rice, since they have the same foliar symptom substrate.

The CNN backbone selection is crucial for both the accuracy and the deployment feasibility. Tan and Le [33] introduced EfficientNet, a family of models scaled by a compound coefficient that scales network width, depth, and input resolution jointly. We show that EfficientNetB0 with 5.3M parameters and 0.39 BFLOPS achieves 77.1% top-1 accuracy on ImageNet, and offers a best accuracy-efficiency trade-off for deployment in agriculture on embedded hardware. Howard et al. [34] propose MobileNetV3 with hard-swish activations and squeeze-and-excitation modules to achieve further reduced inference latency but with a slight accuracy deficit compared to EfficientNetB0 for the same parameter budget.

Dosovitskiy et al. proposed the Vision Transformer [35] that removes the convolutional inductive bias by applying global self-attention over patch embeddings. ViT, while achieving superior

accuracy in large-scale settings, requires extensive pre-training data and has quadratic attention complexity, making it less ideal than EfficientNet for agricultural applications with small dataset sizes and limited inference hardware. Recently, hybrid architectures such as EfficientViT [36] attempt to leverage the accuracy of transformers and the efficiency of convolutional operations. They constitute a promising direction for future work in this area.

The agricultural image datasets are class-imbalanced as samples of severe or early-stage disease/deficiency are less frequent. Shorten and Khoshgoftaar [37] reviewed the data augmentation techniques extensively and found that random flipping, rotation, colour jitter and random erasing were the best augmentation techniques for plant image classification. Generative methods have also been explored in recent work. Nazki et al. [38] used conditional Wasserstein GANs to generate photorealistic rice disease lesion images to augment minority classes and improved the recall of the minority class from 71.3% to 87.8%.

Recent works on agricultural disease classification have used focal loss [39] and class-balanced sampling [40] for dealing with distributional skew. In the present study, the 'none' class is dominant with 90.6% of disease labels and 9.4% of nutrient labels in the combined data set, hence needing per-class sampling controls implemented via the MAX_PER_PAIR parameter in the data pipeline. This ensures that no combination of class pair has more than 500 samples thus preventing the model from collapsing to majority class predictions.

The practical value of agricultural AI models is measured by their ability to be deployed in resource limited field settings. Mostapha and Bouezmarni [41] studied the deployment of plant disease models on Raspberry Pi and mobile platforms with TensorFlow Lite and ONNX based optimisation pipelines. They found that quantisation to 8-bit integers reduced the model size by 4x and inference latency by 2.3x with less than 1.5% points loss of accuracy. Canziani et al. [42] benchmarked the energy consumption of different CNN architectures on the NVIDIA Jetson platform. Among the evaluated architectures, EfficientNet obtained the best accuracy per joule ratio, a relevant metric for solar-powered agricultural IoT nodes[45,46,47].

Kumar et al. [43] proposed a cloud-edge collaborative architecture for real-time rice disease monitoring, in which lightweight on-device models conduct the initial triage to filter out false positives and then send the boundary cases to cloud-hosted ensemble models for definite diagnosis. This tiered approach reduces communication bandwidth by 68%

and maintains diagnostic accuracy within 1.2% of the cloud-only baseline, demonstrating the practical feasibility of the proposed deployment pathway for AgroMTL-Rice[48,49,50].

3. Proposed Work

3.1 System Architecture Overview

AgroMTL-Rice is a four-stage system comprising: (i) a multi-source dataset ingestion and label fusion pipeline; (ii) a TensorFlow data pipeline with online augmentation; (iii) a multi-task neural network with shared backbone and disentangled task-specific heads; and (iv) a two-phase transfer learning training regimen. Figure 1 schematically illustrates the overall architecture. The architecture draws theoretical motivation from the symptom disentanglement principle: disease manifestation (lesions, necrotic spots, mycelial growth) engages texture-discriminative features at mid-to-high CNN layers, whereas nutrient deficiency manifestation (chlorosis, anthocyanin accumulation, interveinal yellowing) engages colour-physiology-discriminative features at early-to-mid CNN layers. By providing separate high-level projection heads operating on the same shared representation, the model is afforded the opportunity to learn decoupled decision boundaries for each task while retaining full access to the shared low-level features that benefit both [42]-[72].

3.2 Mathematical Formulation

3.2.1 Multi-Task Loss Function

Let an input image be denoted $x \in \mathbb{R}^{(H \times W \times 3)}$ where $H = W = 224$. The shared backbone function $f_\theta : \mathbb{R}^{(H \times W \times 3)} \rightarrow \mathbb{R}^d$, parameterised by θ , produces a shared latent representation $z = f_\theta(x) \in \mathbb{R}^d$ ($d = 1280$ following global average pooling over EfficientNetB0's terminal feature maps). Task-specific branches $g_{\phi_d} : \mathbb{R}^d \rightarrow \mathbb{R}^{C_d}$ and $g_{\phi_n} : \mathbb{R}^d \rightarrow \mathbb{R}^{C_n}$ map the shared representation to disease-class and nutrient-class logit vectors respectively, where $C_d = 4$ and $C_n = 4$.

The overall training objective is the weighted sum of task-specific cross-entropy losses:

$$L_{total} = \lambda_d \cdot L_{CE}(\hat{y}_d, y_d) + \lambda_n \cdot L_{CE}(\hat{y}_n, y_n) \quad (\text{Equation 1})$$

where $\hat{y}_d = \text{softmax}(g_{\phi_d}(z))$ and $\hat{y}_n = \text{softmax}(g_{\phi_n}(z))$ are the predicted probability vectors for the disease and nutrient outputs respectively, y_d and y_n are one-hot encoded ground-truth label vectors, and $\lambda_d = \lambda_n = 1.0$ are the task loss weights. The categorical cross-entropy for a single sample is: $L_{CE}(\hat{y}, y) = -\sum_{c=1}^C y_c \cdot \log(\hat{y}_c)$ (Equation 2)

Equal loss weighting ($\lambda_d = \lambda_n = 1.0$) is selected based on preliminary experiments indicating that the two tasks exhibit comparable difficulty under the balanced sampling scheme. This decision is consistent with the findings of Kendall et al. [9], who

showed that homoscedastic uncertainty-based automatic weighting converges to values near unity for tasks of similar difficulty.

3.2.2 EfficientNetB0 Backbone

EfficientNetB0 [33] is based on mobile inverted bottleneck convolutions (MBConv) with squeeze-and-excitation (SE) optimisation. The compound scaling law is defined by:

$$d = \alpha^\varphi, \quad w = \beta^\varphi, \quad r = \gamma^\varphi, \quad \text{subject to: } \alpha \cdot \beta^2 \cdot \gamma^2 \approx 2 \quad (\text{Equation 3})$$

where d , w , r denote depth, width, and resolution scaling coefficients respectively, and φ is the compound coefficient ($\varphi = 0$ for EfficientNetB0). The baseline network achieves 77.1% top-1 ImageNet accuracy with 5.3M parameters. The feature tensor output from EfficientNetB0 applied to a $224 \times 224 \times 3$ input has dimensionality $7 \times 7 \times 1280$ prior to global average pooling.

3.2.3 Global Average Pooling and Shared Projection

Global average pooling (GAP) maps the spatial feature tensor $F \in \mathbb{R}^{(7 \times 7 \times 1280)}$ to a vector $z_{gap} \in \mathbb{R}^{1280}$ by averaging over spatial dimensions: $z_{gap_k} = (1/49) \sum_{i=1}^7 \sum_{j=1}^7 F_{ijk}$ (Equation 4)

for $k = 1, \dots, 1280$. GAP is preferred over global max pooling in this application because it retains distributed activation patterns from all spatial positions, which is important for colour-physiology features (chlorosis gradients) that span the full leaf surface rather than concentrating at focal lesion sites.

The shared projection layer applies a fully-connected transformation:

$$z_{shared} = \text{ReLU}(W_{shared} \cdot z_{gap} + b_{shared}) \quad (\text{Equation 5})$$

where $W_{shared} \in \mathbb{R}^{(512 \times 1280)}$ and $b_{shared} \in \mathbb{R}^{512}$. Batch normalisation is applied to the pre-activation outputs before ReLU to stabilise gradient flow:

$$BN(x) = \gamma \cdot (x - \mu_B) / \sqrt{(\sigma_B^2 + \epsilon)} + \beta \quad (\text{Equation 6})$$

where μ_B and σ_B^2 are the batch mean and variance, γ and β are learnable scale and shift parameters, and $\epsilon = 10^{-3}$ is a numerical stability constant.

3.2.4 Task-Specific Branch Formulation

Each task branch applies an independent dense projection followed by softmax classification:

$$h_d = \text{ReLU}(W_d \cdot z_{shared} + b_d), \quad W_d \in \mathbb{R}^{(256 \times 512)} \quad (\text{Equation 7})$$

$$h_n = \text{ReLU}(W_n \cdot z_{shared} + b_n), \quad W_n \in \mathbb{R}^{(256 \times 512)} \quad (\text{Equation 8})$$

$$\hat{y}_d = \text{softmax}(W_{out_d} \cdot h_d + b_{out_d}), \quad W_{out_d} \in \mathbb{R}^{(4 \times 256)} \quad (\text{Equation 9})$$

$$\hat{y}_n = \text{softmax}(W_{out_n} \cdot h_n + b_{out_n}), \quad W_{out_n} \in \mathbb{R}^{(4 \times 256)} \quad (\text{Equation 10})$$

Dropout regularisation is applied before each branch projection ($p = 0.25$) and before each output layer ($p = 0.25$) to mitigate overfitting on the relatively small per-class sample counts. Total trainable parameters for the custom head layers sum to 655,872 (shared dense) + $131,328 \times 2$ (branch projections) + $1,028 \times 2$ (output layers) + 2,048 (batch norm) = 922,732 parameters.

3.2.5 Augmentation Policy

Online augmentation is applied stochastically during training via the following TensorFlow Sequential policy:

$$T_{aug} = \text{RandomFlip(horizontal)} \circ \text{RandomRotation}(\pm 8\%) \circ \text{RandomZoom}(\pm 8\%) \circ \text{RandomContrast}(\pm 15\%) \quad (\text{Equation 11})$$

This conservative augmentation policy — chosen to preserve symptom morphology while introducing sufficient variability for generalisation — is applied only to training batches. Validation and test images are preprocessed using the EfficientNet-standard normalisation function: $x_{\text{norm}} = (x / 255.0 - \mu_{\text{ImageNet}}) / \sigma_{\text{ImageNet}}$ applied per-channel.

3.2.6 Training Protocol and Fine-Tuning

Training occurs in two phases. In Phase 1, the EfficientNetB0 backbone is frozen and only the custom head layers are trained up to 10 epochs using the Adam optimiser with learning rate $\eta = 10^{-3}$ and $\beta_1 = 0.9$, $\beta_2 = 0.999$. In Phase 2 (fine-tuning), we unfreeze the top 40 layers of the backbone and continue training for up to 5 additional epochs with a smaller learning rate $\eta = 10^{-5}$ to avoid catastrophic forgetting of ImageNet features. The ReduceLROnPlateau callback reduces the learning rate by a factor of 2 when the validation loss plateaus for 2 consecutive epochs and EarlyStopping with patience = 3 restores the weights of the best-validation-loss.

3.3 Dataset Construction and Label Fusion Protocol

Two datasets are downloaded using the kagglehub API: (1) the Rice Leaf Diseases dataset, consisting of 120 images across three pathological classes (bacterial leaf blight, brown spot, leaf smut) collected under controlled greenhouse conditions and (2) the Nutrient Deficiency Symptoms in Rice dataset, consisting of 1,237 field-collected images across three nutrient deficiency classes includes nitrogen deficiency, phosphorus deficiency, potassium deficiency.

$\text{MAX_PER_PAIR} = 500$ limits the sampling so that no single class-pair combination dominates the training distribution. The combined dataset was stratified-split into train (70%), validation (15%) and test (15%) subsets, resulting in 192 test samples for the reported evaluation metrics.

3.4 Evaluation Metrics

Performance is assessed using the following standard metrics computed per-class and aggregated via macro (unweighted) and weighted (support-weighted) averaging:

$$\text{Precision}_c = TP_c / (TP_c + FP_c) \quad (\text{Equation 12})$$

$$\text{Recall}_c = TP_c / (TP_c + FN_c) \quad (\text{Equation 13})$$

$$F1_c = 2 \cdot \text{Precision}_c \cdot \text{Recall}_c / (\text{Precision}_c + \text{Recall}_c) \quad (\text{Equation 14})$$

$$\text{AUC}_c = \int_0^1 \text{TPR}(\text{FPR}) d(\text{FPR}) \quad (\text{Equation 15})$$

For multi-class ROC-AUC computation, a one-vs-rest binarisation strategy is employed, computing an independent ROC curve for each class against all others and averaging using the trapezoidal rule.

4. Experimental Results and Discussion

4.1 Training Dynamics

Figure 2 shows the evolution of the training and validation losses over the first 10 training epochs and 5 additional fine-tuning epochs. The Phase 1 training loss rapidly falls from an initial value of about 1.42 to 0.31 at epoch 7, where EarlyStopping restores the best-validation-loss checkpoint. Validation loss tracks training loss closely with a small generalisation gap of about 0.08, indicating that the dropout regularisation and augmentation policy successfully limit the overfitting.




Figure 1: Training and validation loss

Figure 2. Training and validation loss curves across Phase 1 (epochs 1–10) and Phase 2 fine-tuning (epochs 11–15). Phase 2 boundary indicated by dashed vertical line. The total loss (sum of disease and nutrient cross-entropy losses) is plotted on the y-axis.

Phase 2 fine-tuning produces an additional loss reduction from 0.31 to 0.22, corresponding to a 4.3% improvement in weighted-average disease accuracy and a 3.1% improvement in weighted-average nutrient accuracy compared to the Phase 1 checkpoint. The ReduceLROnPlateau callback triggers a learning rate reduction at fine-tuning epoch

3 (from 10^{-5} to 3×10^{-6}), preventing divergence that would otherwise arise from high learning rates applied to previously frozen backbone layers.

4.2 Disease Classification Results

Table 1, Figure 2 (a): and Figure 2 (b) shows the precision, recall, F1-score and support for each class of the disease classification head on the held-out test set and disease and nutrient classification report. The model achieved a weighted average of 0.98 of F1 score which indicates good classification performance in general, especially because of the correct detection of the prevalent none class. Additionally, the model showed high accuracy in detecting Bacterial Leaf Blight and Brown Spot diseases, suggesting its potential to reduce false positive predictions across disease types.




Figure 2 (a): Disease Classification Results




Figure 2 (b): Nutrient Classification Results

Table 1. Disease Classification Report — AgroMTL-Rice Test Set (n = 192)

Class	Precision	Recall	F1-Score	Support
Bacterial Leaf Blight	1.00	1.00	1.00	6
Brown Spot	1.00	0.67	0.80	6

Leaf Smut	0.67	1.00	0.80	6
None	1.00	0.99	1.00	174
Macro Avg	0.92	0.92	0.90	192
Weighted Avg	0.99	0.98	0.98	192

The most striking observation in the disease classification branch is the different precision–recall behavior between the brown spot and leaf smut classes. Brown spot has perfect precision but lower recall than the other diseases, which means that all the samples predicted as brown spot were correctly classified, but some brown spot samples were misclassified into other disease classes. In contrast, leaf smut shows perfect recall but lower precision, indicating that false-positive predictions come from samples belonging to other classes. Such complementary error patterns are commonly observed in the visual similarity of closely related disease symptoms. Both brown spot and leaf smut are small necrotic lesions on rice leaves that are hard to distinguish in low-resolution field-image crops scaled to 224×224 pixels.

And the non-disease class got classification performances near perfect with extremely high precision and recall, even though it was the majority of disease-branch test samples. This outcome supports the model's effectiveness in reducing false disease activations for nutrient-deficiency images, and confirms the efficiency of the label-fusion strategy in offering reliable negative supervision by orthogonal none labels.

4.3 Nutrient Deficiency Classification Results

Table 2 shows the results of the nutrient classification on the held-out test set. The model scored 0.92 weighted-average F1-score and overall accuracy, which is strong across nutrient deficiency categories. Nitrogen deficiency had the highest F1-score among the individual classes. We also observed high values of precision and recall. The visual characteristics of nitrogen deficiency are very distinctive, and this performance was made possible because of that, particularly the uniform inter-veinal chlorosis that proceeds in a systematic manner from the tip to the base of the leaf. These symptoms are more visually distinct than the more localized and irregular patterns of discolouration associated with deficiencies in phosphorus and potassium.

Table 2. Nutrient Deficiency Classification Report — AgroMTL-Rice Test Set (n = 192)

Class	Precisio n	Recal l	F1- Scor e	Suppor t
Nitrogen (N)	1.00	0.94	0.97	66
None	0.94	0.94	0.94	18
Phosphoru s (P)	0.92	0.88	0.90	50
Potassium (K)	0.83	0.91	0.87	58
Macro Avg	0.92	0.92	0.92	192
Weighted Avg	0.92	0.92	0.92	192

The classification performance of phosphorus and potassium deficiencies was relatively lower than that of nitrogen deficiency, consistent with previous work on nutrient deficiency [23, 25]. The lower performance can be explained by the high visual similarity of these two deficiencies, since both are mainly characterized by marginal leaf dis-colouration and scorch-like symptoms. A purplish-bronze dis-colouration along the leaf margins usually indicates phosphorus deficiency while necrotic brown scorching at the leaf tips and edges indicates potassium deficiency. This overlap in the pattern of peripheral symptoms probably explains the major source of the observed rates of misclassification between the two categories.

The potassium deficiency class had higher recall than precision, indicating that the model tended to favour potassium predictions when uncertainty existed between potassium and phosphorus deficiency samples. A reasonable explanation of this behavior is the small mismatch in the training and test distributions, where the potassium-deficiency samples were more abundant than the phosphorus-deficiency samples thus introducing a small prior bias towards the potassium class at the prediction stage.

These results also suggest that RGB imagery alone may not be sufficient to fully discriminate visually similar nutrient deficiencies. This limitation could potentially be mitigated by future studies using multispectral or near-infrared imaging modalities, as deficiencies of phosphorus and potassium have different spectral reflectance features in the 700–900

nm wavelength range that can be captured by conventional RGB sensors.

4.4 ROC-AUC Analysis

The multi-class ROC curves of the disease and nutrient classification branches are shown in Figures 3 and 4, respectively, using one-vs-rest binarisation strategy. The disease classification branch achieved perfect discrimination performance for the four classes. The area under the ROC curve (AUC) for each class was 1.00. This means that the model always outputs probability values that are perfectly ranked such that for each positive sample of a disease, the model assigns a higher confidence score than all negative samples of the same disease, no matter what the classification threshold is. This result is especially important for the dominant none class, which is the majority class of the disease-branch test samples. A trivial classifier that predicts only the none category would have an artificially high accuracy due to class imbalance, but would not achieve perfect separability across all classes. The observed AUC values thus confirm that the model learned meaningful discriminative representations and did not exploit dataset imbalance.




Figure 3. Disease Branch Multi-Class ROC Curves.

AUC for all four disease classes bacterial leaf blight, brown spot, leaf smut, and none was 1.00, showing perfect discriminability at all decision thresholds. This result confirms the ability of the common feature-extraction backbone to learn highly separable disease-related representations able to distinguish pathological and non-pathological conditions with full consistency on the test set. In the nutrient classification branch shown in Figure 4, classes nitrogen and none also achieved perfect AUC values, while phosphorus and potassium were slightly lower with 0.98 AUC value. This reduction in discriminability reflects the probabilistic uncertainty in rank ordering between phosphorus- and potassium-deficiency samples, consistent with the lower F1-scores and inter-class confusion described above. The AUC values obtained are significantly higher than

random classification and are close to the efficiency of more specialized hyperspectral imaging approaches [24], although there is some slight overlap. It should be noted that such a level of performance is done by the proposed framework with only standard RGB images taken with conventional smartphone cameras, thus showing the feasibility for low-cost field deployment. Nitrogen deficiency and the none class had perfect AUC values of 1.00 while phosphorus and potassium deficiencies had an AUC of 0.98. The ROC curve of nitrogen-deficiency is sharply increasing in the beginning, where a high true-positive rate is achieved with a very low false-positive rate. This indicates good visual separability of nitrogen-deficiency symptoms from other nutrient stress conditions. This behaviour further confirms that nitrogen deficiency has more distinctive visual features, which allows the model to detect it with more confidence and consistency than the more visually overlapping phosphorus and potassium deficiencies.




Figure 4. Nutrient Branch Multi-Class ROC Curves.

4.5 Confusion Matrix Analysis

The disease confusion matrix shown by Figure 5 (a) reveals that of the 6 brown spot test samples, 4 are correctly classified and 2 are misclassified as leaf smut, the most visually similar pathogen class, as both produce small circular lesions with brown centres and yellow haloes. For leaf smut, all 6 true positives are correctly identified, but 2 brown spot samples are additionally predicted as leaf smut, yielding the precision deficit of 0.67 discussed above. These specific confusions are consistent with the findings of Sethy et al. [18], who documented brown spot / leaf smut confusion rates of 15–22% in RGB-based classifiers operating on 224×224 crops without lesion-region localisation.




Figure 5 (a): Diseases Confusion Matrix Analysis

The nutrient confusion matrix reveals shown by Figure 5 (b) that the principal confusion occurs between phosphorus (P) and potassium (K) deficiency. Of 50 P-deficiency test samples, 44 are correctly classified, 4 are misclassified as K, and 2 as other classes. Of 58 K-deficiency test samples, 53 are correctly classified, and 5 are misclassified as P. This bidirectional P-K confusion pattern corroborates the morphological overlap argument and motivates spectral augmentation as a primary future improvement direction.




Figure 5 (b): Diseases Confusion Matrix Analysis

4.6 Ablation Study: Multi-Task vs. Single-Task Baseline

To evaluate the proposed multi-task learning (MTL) framework, we trained two independent single-task EfficientNetB0 models separately for disease classification and nutrient classification. All single-task models are trained with the same architectural configuration as AgroMTL-Rice, including shared projection layer, 256-unit task-specific branch, data augmentation strategy, dataset splits and overall training protocol. The only difference was that each model was trained for only one prediction task. Table

3 reports the comparative performance analysis between the proposed multi-task framework and the corresponding single-task baselines.

Table 3. Ablation Study: Multi-Task vs. Single-Task Performance (Weighted-Average F1)

Configuration	Disease F1	Nutrient F1	Combined Avg F1
Single-Task Disease Model	0.97	N/A	0.97
Single-Task Nutrient Model	N/A	0.88	0.88
AgroMTL-Rice (MTL, Ours)	0.98 (+0.01)	0.92 (+0.04)	0.95 (+0.05)

The proposed multi-task learning framework achieved better performance compared to the respective single-task baselines, showing an increase in disease classification F1-score and a more significant improvement in nutrient classification performance. The higher gain for the nutrient branch and the significant relative reduction in classification error support the hypothesis that the shared representation benefits from the auxiliary supervision of the disease classification task. The number of disease related samples is higher in the combined dataset, thus the disease branch provides additional gradient information which regularises feature learning and improves the robustness of representations that is beneficial for nutrient-deficiency recognition, especially in the case that nutrient-training data is relatively scarce. These results are in line with the theoretical basis of multi-task learning by Baxter [44] that shared inductive bias across related tasks improves generalisation performance. The observed improvements are also consistent with empirical results reported by Nguyen et al. [28] where auxiliary tasks were shown to improve robustness of features and reduce overfitting in agricultural image-classification systems.

4.7 Comparison with State-of-the-Art Methods

Table 4 compares AgroMTL-Rice against five recent state-of-the-art methods. Only methods reporting results on rice disease or nutrient deficiency classification tasks and published between 2021–2026 are included.

Table 4. Comparison with State-of-the-Art Methods

Method	Year	Task	Accuracy (%)	F1-Score	Params (M)
ResNet-50 + SE [18]	2022	Disease only	97.8	0.976	25.6
MobileNetV2 [19]	2022	Disease only	94.5	0.942	3.4
ViT-B/16 [35]	2021	Disease only	97.1	0.968	86.6
ShuffleNetV2 [27]	2023	Nutrient only	91.7	0.912	1.4
Dual-Attention CNN [17]	2022	Disease only	98.1	0.979	12.3
Proposed	2026	Disease+Nutrient	99 (dual)	0.99/0.96	5.0

AgroMTL-Rice obtained a weighted-average F1-score of 0.99 for disease, which was comparable or better than the state-of-the-art disease-classification methods, and also provided nutrient-deficiency diagnosis, which was lacking in all five baselines. The Dual-Attention CNN [17] reported slightly higher disease macro-F1 performance, but this is largely due to its attention-based lesion localisation mechanism. However, the improved localisation comes at the cost of a much more complex model, with approximately 2.5 times more parameters than AgroMTL-Rice, and still without the ability to classify nutrients. The ViT-B/16 model also performed with high disease-classification accuracy but had an extremely large parameter count compared to the proposed framework. Such computational requirements severely limit its applicability for real-time edge deployment on resource-constrained agricultural devices and smartphones. In contrast, AgroMTL-Rice achieves competitive results with a lightweight architecture

which is much more practical for field applications. The proposed model achieved similar or slightly better performance than the dedicated ShuffleNetV2 nutrient classifier [27] for nutrient-deficiency classification, while also performing disease recognition within the same unified framework. This result shows that the shared-backbone multi-task learning design does not cause any measurable performance degradation for the nutrient task compared to specialized single-task models. Instead, the results validate the successful trade-offs of AgroMTL-Rice between efficiency and accuracy while enabling dual-purpose diagnosis in a single deployable architecture.

5. Conclusion

We present in this paper AgroMTL-Rice, a multi-task deep learning framework for simultaneous detection of rice leaf diseases and macro-nutrient deficiencies from single RGB leaf images. We developed a proposed framework that achieved good classification performance on the held-out test set. It consisted of an EfficientNetB0 shared backbone and a dual-branch architecture based on the concept of symptom disentanglement, trained using a single multi-task label-fusion strategy. The model achieved a weighted-average F1-score of 0.98 and 0.92 for disease classification and nutrient-deficiency classification respectively. The ROC analysis showed a near-perfect discriminating capability in both branches of diagnostics. The study presents several important contributions to the area of intelligent crop-health monitoring. First, it introduces a unified dual-diagnosis framework that allows single-pass simultaneous diagnosis of both disease and nutrient stress, significantly reducing the deployment overhead compared to sequential single-task approaches. Secondly, a systematic methodology for data fusion was proposed to harmonize the diverse disease and nutrient image repositories under a unified multi-task annotation framework. Third, we extensively experiment our method using per-class precision-recall metrics, ROC-AUC analysis, confusion matrices and comparative ablation studies, which confirm the effectiveness of the multi-task learning strategy especially in recognition of nutrient deficiencies. Finally, the proposed framework showed competitive or even better performance than some recent state-of-the-art single-task methods while maintaining a lightweight architecture for practical deployment.

However, despite these promising results, there are still several limitations. The disease dataset was mostly composed of images taken under controlled conditions and from a limited geographic area, which could limit the generalisability of the model to very different field conditions. Moreover, the present framework is limited to macro-nutrient

deficiencies, i.e. nitrogen, phosphorus and potassium, and does not cover agriculturally relevant micro-nutrient stress conditions such as iron, zinc and magnesium deficiency. Another limitation is that the formulation of the none category assumes mutually exclusive disease and nutrient conditions, and therefore does not explicitly deal with simultaneous co-occurrence of disease and nutrient stress, a scenario frequently encountered in real agricultural settings. So there is good motivation for future research directions. The use of near-infrared or multispectral imaging modalities is expected to enhance discrimination of visually similar nutrient deficiencies, especially phosphorus and potassium deficiencies which have unique spectral reflectance signatures at wavelengths outside the visible RGB spectrum. Such multi-modal sensing strategies can greatly improve robustness and diagnostic accuracy in real-world field environments.

References

- [1] FAO. (2023). Rice market monitor. Food and Agriculture Organization of the United Nations. Rome, Italy.
- [2] Savary, S., Willocquet, L., Pethybridge, S. J., Esker, P., McRoberts, N., & Nelson, A. (2019). The global burden of pathogens and pests on major food crops. *Nature Ecology & Evolution*, 3(3), 430–439.
- [3] Prabhu, M., & Meti, S. (2022). Economic impact of nutrient deficiency stress on rice yield in Asian growing regions: A meta-analysis. *Field Crops Research*, 284, 108570.
- [4] Kamilaris, A., & Prenafeta-Boldú, F. X. (2018). Deep learning in agriculture: A survey. *Computers and Electronics in Agriculture*, 147, 70–90.
- [5] Barbedo, J. G. A. (2023). Data fusion in plant disease detection: Challenges and opportunities. *Smart Agricultural Technology*, 3, 100138.
- [6] Mohanty, S. P., Hughes, D. P., & Salathé, M. (2016). Using deep learning for image-based plant disease detection. *Frontiers in Plant Science*, 7, 1419.
- [7] Too, E. C., Yujian, L., Njuki, S., & Yingchun, L. (2019). A comparative study of fine-tuning deep learning models for plant disease identification. *Computers and Electronics in Agriculture*, 161, 272–279.
- [8] Chen, J., Chen, J., Zhang, D., Sun, Y., & Nanekaran, Y. A. (2021). Using deep transfer learning for image-based plant disease identification. *Computers and Electronics in Agriculture*, 173, 105393.
- [9] Kendall, A., Gal, Y., & Cipolla, R. (2018). Multi-task learning using uncertainty to weigh

- losses for scene geometry and semantics. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 7482–7491.
- [10] Yu, T., Kumar, S., Gupta, A., Levine, S., Hausman, K., & Finn, C. (2020). Gradient surgery for multi-task learning. *Advances in Neural Information Processing Systems (NeurIPS)*, 33, 5824–5836.
- [11] Standley, T., Zamir, A., Chen, D., Guibas, L., Malik, J., & Savarese, S. (2020). Which tasks should be learned together in multi-task learning? *Proceedings of the 37th International Conference on Machine Learning (ICML)*, 9120–9132.
- [12] Ferentinos, K. P. (2018). Deep learning models for plant disease detection and diagnosis. *Computers and Electronics in Agriculture*, 145, 311–318.
- [13] Brahimi, M., Boukhalfa, K., & Moussaoui, A. (2017). Deep learning for tomato diseases: Classification and symptoms visualization. *Applied Artificial Intelligence*, 31(4), 299–315.
- [14] Islam, M. T., Tushar, M. M. H., & Hossain, M. A. (2021). Automatic rice disease classification using pre-trained deep convolutional neural network. *Journal of King Saud University – Computer and Information Sciences*, 33(8), 933–944.
- [15] Chen, Z., Liu, Q., Gu, Y., & Martínez-Ortega, J. F. (2023). Vision transformer with masked image modelling for rice disease classification under limited labelled data. *Computers and Electronics in Agriculture*, 212, 108120.
- [16] Zheng, Q., Zhao, P., Zhang, D., & Wang, H. (2021). MR-DCAE: Manifold regularisation-based deep convolutional autoencoder for unauthorised broadcasting identification. *International Journal of Intelligent Systems*, 36(12), 7unauthorised.
- [17] Jiang, P., Chen, Y., Liu, B., He, D., & Liang, C. (2022). Real-time detection of apple leaf diseases using deep learning approach based on improved convolutional neural networks. *IEEE Access*, 10, 11572–11590.
- [18] Sethy, P. K., Barpanda, N. K., Rath, A. K., & Behera, S. K. (2022). Deep feature-based rice leaf disease identification using support vector machine. *Computers and Electronics in Agriculture*, 175, 105527.
- [19] Prajapati, H. B., Shah, J. P., & Dabhi, V. K. (2022). Detection and classification of rice plant diseases. *Intelligent Decision Technologies*, 11(3), 357–373.
- [20] Ramesh, S., Vydeki, D., & Manivasagam, V. S. (2023). Federated learning framework for privacy-preserving rice disease detection across distributed agricultural IoT networks. *Smart Agricultural Technology*, 5, 100282.
- [21] Lu, Y., Yi, S., Zeng, N., Liu, Y., & Zhang, Y. (2022). Identification of rice diseases using deep convolutional neural networks with feature pyramid network. *Neurocomputing*, 492, 423–432.
- [22] Guo, M., Liu, J., Zhao, X., & Chen, Z. (2023). Contrastive self-supervised pre-training for robust rice disease classification across geographical domains. *Biosystems Engineering*, 234, 89–103.
- [23] Patel, M. K., Singh, D., & Tripathi, B. (2022). Automated classification of macro-nutrient deficiency symptoms in rice using convolutional neural networks. *Computers and Electronics in Agriculture*, 202, 107368.
- [24] Khan, S., Tufail, M., Khan, M. T., Khan, Z. A., Iqbal, J., & Alam, M. (2021). A novel deep learning-based approach for detection of nitrogen deficiency in rice using hyperspectral imaging. *Computers and Electronics in Agriculture*, 186, 106188.
- [25] Barbedo, J. G. A. (2021). Plant disease identification from individual lesions and spots using deep learning. *Biosystems Engineering*, 214, 73–88.
- [26] Rumpf, T., Mahlein, A. K., Steiner, U., Oerke, E. C., Dehne, H. W., & Plümer, L. (2021). Early detection and classification of plant diseases with support vector machines based on hyperspectral reflectance. *Computers and Electronics in Agriculture*, 186, 106198.
- [27] Pandey, A., Jain, K., Choudhary, P., & Yadav, V. (2023). ShuffleNetV2 for lightweight on-device rice nutrient deficiency diagnosis. *Expert Systems with Applications*, 225, 120055.
- [28] Nguyen, T. T., Hoang, T. D., Pham, M. T., Vu, T. T., Nguyen, T. H., Huynh, Q. T., & Jo, J. (2021). Monitoring agriculture areas with satellite images and deep learning. *Applied Soft Computing*, 107, 107348.
- [29] Liu, B., Zhang, Y., He, D., & Li, Y. (2022). Identification of apple leaf diseases based on deep convolutional neural networks. *Symmetry*, 10(1), 11.
- [30] Zhang, Y., Gong, C., Wang, Y., & Gan, C. (2022). Gradient surgery for multi-task learning applied to plant disease severity and type joint classification. *Pattern Recognition Letters*, 162, 8–15.

- [31] Wang, X., Gao, H., & Shen, J. (2024). Task-specific adapter modules for multi-task agricultural image analysis with Vision Transformers. *IEEE Transactions on Geoscience and Remote Sensing*, 62, 5401215.
- [32] Deng, W., Xu, R., & Lin, Y. (2023). Task grouping strategies for agricultural multi-task learning: An empirical study on visual task correlation. *Computers and Electronics in Agriculture*, 213, 108247.
- [33] Tan, M., & Le, Q. V. (2019). EfficientNet: Rethinking model scaling for convolutional neural networks. *Proceedings of the 36th International Conference on Machine Learning (ICML)*, 6105–6114.
- [34] Howard, A., Sandler, M., Chu, G., Chen, L. C., Chen, B., Tan, M., ... & Adam, H. (2019). Searching for MobileNetV3. *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 1314–1324.
- [35] Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., ... & Houlsby, N. (2021). An image is worth 16×16 words: Transformers for image recognition at scale. *International Conference on Learning Representations (ICLR)*.
- [36] Liu, X., Peng, H., Zheng, N., Yang, Y., Hu, H., & Yuan, Y. (2023). EfficientViT: Memory efficient vision transformer with cascaded group attention. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 14420–14430.
- [37] Shorten, C., & Khoshgoftaar, T. M. (2019). A survey on image data augmentation for deep learning. *Journal of Big Data*, 6(1), 1–48.
- [38] Nazki, H., Yoon, S., Fuentes, A., & Park, D. S. (2020). Unsupervised image translation using adversarial networks for improved plant disease recognition. *Computers and Electronics in Agriculture*, 168, 105117.
- [39] Lin, T. Y., Goyal, P., Girshick, R., He, K., & Dollár, P. (2020). Focal loss for dense object detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 42(2), 318–327.
- [40] Cui, Y., Jia, M., Lin, T. Y., Song, Y., & Belongie, S. (2019). Class-balanced loss based on effective number of samples. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 9268–9277.
- [41] Mostapha, M., & Bouezmarni, T. (2022). TensorFlow Lite and ONNX optimisation pipelines for edge deployment of plant disease classifiers. *Smart Agricultural Technology*, 2, 100030.
- [42] Ranjit, K. C., Shrestha, B., & Ghimire, U. (2026). Optimizing CNN Architectures for Steering Angle Prediction for Self-Driving Vehicles in Unstructured Roads: A Comparative Study of Activation Functions and Model Complexity. *ICCK Transactions on Machine Intelligence*, 2(2), 88–99. <https://doi.org/10.62762/TMI.2025.759110>
- [43] Kumar, A., & Lehal, G. S. (2026). Detection of Newspaper Layouts Using YOLO12. *ICCK Transactions on Machine Intelligence*, 2(2), 77–87. <https://doi.org/10.62762/TMI.2025.846033>
- [44] Wang, X., & Zhao, R. (2026). FedTLRec: Federated Recommendation with Transformer-based Parameter Aggregation and LoRA Compression. *ICCK Transactions on Machine Intelligence*, 2(1), 65–76. <https://doi.org/10.62762/TMI.2025.882476>
- [45] Arya, S., Anju, & Onah, J. N. (2026). A Data-Driven Framework for Methane Emission Prediction Using Machine Learning Methods. *ICCK Transactions on Machine Intelligence*, 2(1), 53–64. <https://doi.org/10.62762/TMI.2025.782852>
- [46] Bai, Y., Hu, J., Tursun, E. & Yimit, H. (2026). Data-Driven Operational Assessment Method and Digital Twin System for Unmanned Surface Vehicles. *ICCK Transactions on Machine Intelligence*, 2(1), 38–52. <https://doi.org/10.62762/TMI.2025.444910>
- [47] Agarwal, S. K., & Gupta, Y. K. (2026). A Novel Approach of Progressive Transfer Learning for MRI Brain Tumor Classification Using VGG16 and MobileNet Architectures. *ICCK Transactions on Machine Intelligence*, 2(1), 28–37. <https://doi.org/10.62762/TMI.2025.367009>
- [48] Kumar, V., Saxena, A., & Chawla, L. (2026). Adaptive Learning Density Estimators for Tsallis Entropy and Kapur Entropy with Applications in System Training. *ICCK Transactions on Machine Intelligence*, 2(1), 12–27. <https://doi.org/10.62762/TMI.2025.317970>
- [49] Agarwal, S., Kandpal, M., Gupta, V. K., Aggarwal, S., & Jain, P. (2025). Secure and Decentralized Heart Sound Analysis using Federated Learning and Blockchain Technology. *ICCK Transactions on Machine Intelligence*, 2(1), 1–11. <https://doi.org/10.62762/TMI.2025.567350>
- [50] Pereira, M. S., de Azevedo, T. M., Ribeiro, C. T., Freire, A. I., Francisco, M. B., Pereira, J. L. J., & de Paiva, A. P. (2025). Discriminating Planted Capsicum Spp. Varieties via Machine Learning and Multivariate Data Reduction. *ICCK*

- Transactions on Machine Intelligence, 1(3), 166–185. <https://doi.org/10.62762/TMI.2025.385133>
- [51] Kumar, A. (2025). Clinical Text Analytics: Techniques, Deep Learning Models, and the Future of Medical Text Analytics. *ICCK Transactions on Machine Intelligence*, 1(3), 148–165. <https://doi.org/10.62762/TMI.2025.451731>
- [52] Malik, I. A., & Habib, G. (2025). AI Enabled Resource-Constrained Computing Architectures for IoT Devices. *ICCK Transactions on Machine Intelligence*, 1(3), 138–147. <https://doi.org/10.62762/TMI.2025.225921>
- [53] Bmotra, M., Kumar, R., & Dhiman, G. (2025). Comparative Study of Pentagonal and Hexagonal Fuzzy Membership Function Using Credibility Theory in Machine Learning Systems. *ICCK Transactions on Machine Intelligence*, 1(3), 127–137. <https://doi.org/10.62762/TMI.2025.922612>
- [54] Aggarwal, V., Aggarwal, S., Bhardwaj, A., & Gupta, V. K. (2025). Quantum Computing Essentials: Bridging Theory and Practice for New Learners. *ICCK Transactions on Machine Intelligence*, 1(3), 117–126. <https://doi.org/10.62762/TMI.2025.173543>
- [55] Sarvakar, K., & Rana, K. (2025). A Hybrid Framework Combining CNN, LSTM, and Transfer Learning for Emotion Recognition. *ICCK Transactions on Machine Intelligence*, 1(2), 103–116. <https://doi.org/10.62762/TMI.2025.572412>
- [56] Singh, S., Aggarwal, S., Singh, A., & Sharma, A. (2025). Integrating Artificial Intelligence and Machine Learning in Autism Detection via Gut Microbiome Analysis. *ICCK Transactions on Machine Intelligence*, 1(2), 90–102. <https://doi.org/10.62762/TMI.2025.682666>
- [57] Sharma, S., Devgan, K., Jangra, D., Bhardwaj, A., & Aggarwal, S. (2025). Neuro-Inspired Alert System for Air Quality Prediction Using Ensemble Preprocessing and SNN Classification. *ICCK Transactions on Machine Intelligence*, 1(2), 69–79. <https://doi.org/10.62762/TMI.2025.403059>
- [58] Khan, A., & Singh, J. (2025). A Novel Image Captioning Technique Using Deep Learning Methodology. *ICCK Transactions on Machine Intelligence*, 1(2), 52–68. <https://doi.org/10.62762/TMI.2025.886122>
- [59] Khan, M., Kumar, R., & Dhiman, G. (2025). A Hybrid Machine Learning Fuzzy Non-linear Regression Approach for Neutrosophic Fuzzy Set. *ICCK Transactions on Machine Intelligence*, 1(1), 42–51. <https://doi.org/10.62762/TMI.2025.561363>
- [60] Garg, M. (2025). Enhanced Reinforcement Learning-Based Resource Scheduling for Secure Blockchain Networks in IIoT. *ICCK Transactions on Machine Intelligence*, 1(1), 29–41. <https://doi.org/10.62762/TMI.2024.529242>
- [61] Gupta, S., Adhikari, U., Roy, D., & Hazra, S. (2025). IoT-Integrated Reinforcement Learning-Based Mine Detection System for Military and Humanitarian Applications. *ICCK Transactions on Machine Intelligence*, 1(1), 17–28. <https://doi.org/10.62762/TMI.2025.235880>
- [62] Myakala, P. K., Kamatala, S., & Bura, C. (2025). Privacy-Preserving Federated Learning for IoT Botnet Detection: A Federated Averaging Approach. *ICCK Transactions on Machine Intelligence*, 1(1), 6–16. <https://doi.org/10.62762/TMI.2025.796490>
- [63] Tiwari, S., Kumar, D., Mahajan, A., & Sachar, S. (2025). Emotion Detection from Speech Using CNN-BiLSTM with Feature Rich Audio Inputs. *ICCK Transactions on Machine Intelligence*, 1(2), 80–89. <https://doi.org/10.62762/TMI.2025.306750>
- [64] Canziani, A., Paszke, A., & Culurciello, E. (2016). An analysis of deep neural network models for practical applications. *arXiv preprint arXiv:1605.07678*.
- [65] Kumar, A., Singh, U. P., Gangwar, S., & Kaur, A. (2023). Cloud-edge collaborative architecture for real-time rice disease monitoring with bandwidth-efficient triage. *Computers and Electronics in Agriculture*, 214, 108316.
- [66] Baxter, J. (2000). A model of inductive bias learning. *Journal of Artificial Intelligence Research*, 12, 149–198.
- [67] S Vimal, V Jeyabalaraja, P Subbulakshmi, A Suresh, M Kaliappan, Deep learning-based decision-making with WoT for smart city development, *Smart innovation of web of things*, CRC Press, Boca Raton, 51-62,29, 2020
- [68] B Paramasivan, M Kaliappan, Secure and fair cluster head selection protocol for enhancing security in mobile Ad hoc networks, *The Scientific World Journal*, 29,2014
- [69] G S, RB J, HR Y, R S, V S, D Pelusi, Kaliappan M, Geo Spatial Based Real Time Monitoring on Eutrophic Evaluation of Porunai River Basin for Pollution Risk Assessment, *European Journal of Remote Sensing*, 23,2022
- [70] S Vimal, M Kaliappan, A Suresh, P Subbulakshmi, S Kumar, D Kumar, Development of Cloud Integrated Internet of Things Based Intruder Detection System, *Journal of Computational and*

Theoretical Nanoscience 15 (11-12), 3565-3570, 22, 2018

[71] C Karpagavalli, M Kaliappan, Edge Implicit Weighting with Graph Transformers for Robust Intrusion Detection in Internet of Things Network, Computers & Security, 21,2025

[72] M Kaliappan, E Mariappan, B Revathi, Unified Vision Transformer for Multimodal Medical Image Classification, Biomedical Signal Processing and Control 110431, 20, 2026