ClimaMango: Climate-Fused Mango Disease Detection via EfficientNet-LSTM-CrossAttention

Bhagirath Singh[1], Dr. Jai Singh Gupta[2], Dr. Akash Saxena[3], Dr. Mohini Dwivedi[4]
Research Scholar[1], Professor[2], Professor[3] , Associate Professor[4]
Department of Computer Science, Shridhar University, Pilani, Rajasthan, India[1,2,4]

Department of Computer Science and Engineering, Compucom Institute of Technology and Management, Jaipur, Rajasthan, India[3]
bhagirathbs@gmail.com , jaigupta2@gmail.com , akash27jaipur@gmail.com, mohini.ceeri@gmail.com

Abstract

Mango disease classification from leaf images has reached near-saturation accuracy (>99%) through deep CNN approaches, yet existing methods are inherently reactive: they diagnose visible symptoms rather than predicting pre-symptomatic outbreak risk. This paper proposes ClimaMango, the first dual-branch architecture fusing EfficientNet-B0 image features with NASA POWER meteorological time-series through an eight-head cross-attention module. A 2-layer LSTM processes 14-day weather windows (10 features/day: 7 raw NASA POWER parameters + 3 derived epidemiological indicators) from six UP/Bihar mango-growing districts. An epidemiological pairing strategy enables training on timestampless datasets by matching disease classes to climatically consistent weather windows. On MangoLeafBD (4,000 images, 8 classes), ClimaMango achieves 99.67% accuracy, F1=0.9967, AUC-ROC=1.000 — 75% fewer errors than the image-only baseline (98.67%). A climate effect experiment confirms the LSTM learns genuine epidemiological knowledge: the same leaf image produces different and climatically coherent predictions under different seasonal weather windows. All experiments reproduce on a 4 GB consumer GPU.

Keywords: Mango disease detection; EfficientNet-B0; LSTM; Cross-attention; NASA POWER; Pre-symptomatic prediction; MangoLeafBD

1. INTRODUCTION

Mango (Mangifera indica L.) represents the primary income source for approximately 25 million smallholder farmers in northern India, with annual disease-related losses estimated at INR 15,000–25,000 crore [1]. Anthracnose (Colletotrichum gloeosporioides) and powdery mildew (Oidium mangiferae) are the two dominant pathogens, with 5–15 day latent periods during which no visual symptoms appear. The epidemiology of both diseases is strongly climate-dependent: anthracnose requires RH > 70% and $T \in [22, 34]$°C during the monsoon season; powdery mildew requires cool nights ($T_{min} < 17$°C) and low humidity during the February–April flowering season [3].

Deep learning approaches for mango leaf disease classification have achieved 98–99.87% accuracy on the MangoLeafBD benchmark [4,5,6] using EfficientNet, InceptionV3, and Vision Transformer architectures. However, all published methods classify visible symptoms from leaf images alone, ignoring the climatic context that determines outbreak risk. Two orchards with identically appearing leaves photographed in August monsoon vs January winter represent fundamentally different disease risk profiles, and a model that ignores this context cannot support pre-symptomatic early warning. To the best of our knowledge, no published study has integrated NASA POWER meteorological time-series with mango leaf imagery for disease prediction.

This paper contributes: (1) ClimaMango — the first dual-branch CNN-LSTM-CrossAttention architecture for climate-aware mango disease prediction; (2) an epidemiological pairing strategy applicable to any timestampless agricultural disease dataset; (3) a controlled climate effect experiment validating genuine epidemiological learning; (4) all results reproducible on a 4 GB consumer GPU.

Doi:10.48047/ijiee.2025.15.7.19

2. RELATED WORK

2.1 Mango Disease Detection

Mohanty et al. [7] established that ImageNet pre-trained CNNs achieve high accuracy for plant disease classification. For mango specifically, Singh et al. [4] (99.76%) and Varma et al. [5] (99.87%) represent the current state-of-the-art on MangoLeafBD. Khan et al. [6] applied LSTM to image sequences for temporal disease modelling, but did not incorporate external weather data. No prior study integrates real meteorological reanalysis data with mango leaf images.

2.2 Attention-Based Multimodal Fusion

Bahdanau et al. [8] introduced attention for encoder-decoder sequence models. Vaswani et al. [9] generalised this to multi-head self-attention and cross-attention, enabling one modality's representation to attend selectively to another modality. Cross-attention has been applied to multimodal fusion in remote sensing and medical imaging, but not to climate-disease prediction for crops.

2.3 NASA POWER for Agricultural Research

NASA POWER [10] provides daily reanalysis-based surface meteorological parameters at 0.5° × 0.625° resolution, freely accessible via API. The AG community provides parameters specifically selected for agricultural modelling including T2M, RH2M, PRECTOTCORR, and T2MDEW. The system has been validated for India against IMD station data (MAE < 1.2°C for temperature, < 6% for relative humidity) [10].

3. MATERIALS AND METHODS

3.1 Datasets

MangoLeafBD [11] (4,000 images, 8 classes, 500/class, DOI: 10.17632/hxsnvwty3r.1) was split 70/15/15 (seed=42): 2,800/600/600 images. Image dimensions range 189–320 px (modal: 240×240). Normalisation: mean=[0.670, 0.689, 0.683], std=[0.184, 0.179, 0.239] computed from 200 sampled images. NASA POWER [10] daily data for 6 UP/Bihar districts: Lucknow (26.85°N, 80.95°E), Varanasi, Saharanpur, Muzaffarpur, Patna, Darbhanga. Period: 2020–2024 (1,827 records × 6 districts = 10,962 total). Seven raw parameters: T2M, T2M_MAX, T2M_MIN, RH2M, PRECTOTCORR, WS2M, T2MDEW. Three derived: DTR = T2M_MAX − T2M_MIN; VPD = e_s(T2M) − e_a(T2MDEW) [kPa]; WET_DAY = 1[PRECTOTCORR>1].

3.2 Epidemiological Pairing Strategy

MangoLeafBD images lack timestamps. Each disease class is assigned weather windows consistent with its known risk conditions (Table 1). This teaches the model: (image features) + (climatically consistent season) → disease class.

Table 1. Epidemiological pairing rules: disease class → climate window criteria.

Disease Class	Season / Months	RH2M Condition	Temperature
Anthracnose	May–Sep (monsoon)	RH > 70%	T ∈ [22, 34]°C, precip > 0.5 mm
Powdery Mildew	Feb–Apr (flowering)	RH < 65%	T_min < 17°C
Bacterial Canker	Jun–Sep (monsoon)	RH > 72%	T > 26°C
Die Back	Oct–Dec (post-monsoon)	RH < 55%	Any
Gall Midge	Mar–May (flush)	Any	T > 20°C
Healthy	Nov–Feb (winter)	RH < 60%	T ∈ [15, 30]°C, precip < 2 mm
Sooty Mould	Jun–Aug (monsoon)	RH > 68%	T > 24°C
Cutting Weevil	May–Jul (fruit dev.)	Any	T > 25°C

Doi:10.48047/ijiee.2025.15.7.19

Algorithm 1: Epidemiological Window Pairing

Algorithm 1: Epidemiological Pairing (for timestampless datasets)

```
Input:  Disease class C, weather DataFrame W_norm, split ∈
{train,val,test}
Output: Pool of valid 14-day normalised weather windows for class C
   rules ← DISEASE_RULES[C]   // RH, T, month constraints
   windows ← []
   for district d in {Lucknow, Varanasi, Saharanpur, Muzaffarpur, Patna,
Darbhanga}:
       W_d ← filter(W_norm, district=d, year ∈ split_years)
       for end_idx in range(14, len(W_d)):
           month ← W_d[end_idx].month
           if month ∉ rules.months: continue
           raw_row ← W_raw[end_idx]   // unnormalised for threshold
check
           if not satisfies(raw_row, rules): continue
           w ← W_d[end_idx-14 : end_idx]  // (14, 10) normalised window
           windows.append(w)
   if |windows| < 50: windows ← ALL windows for split (fallback)
   return windows
```

3.3 Model Architecture

ClimaMango consists of three components. Fig. 1 shows the complete architecture.

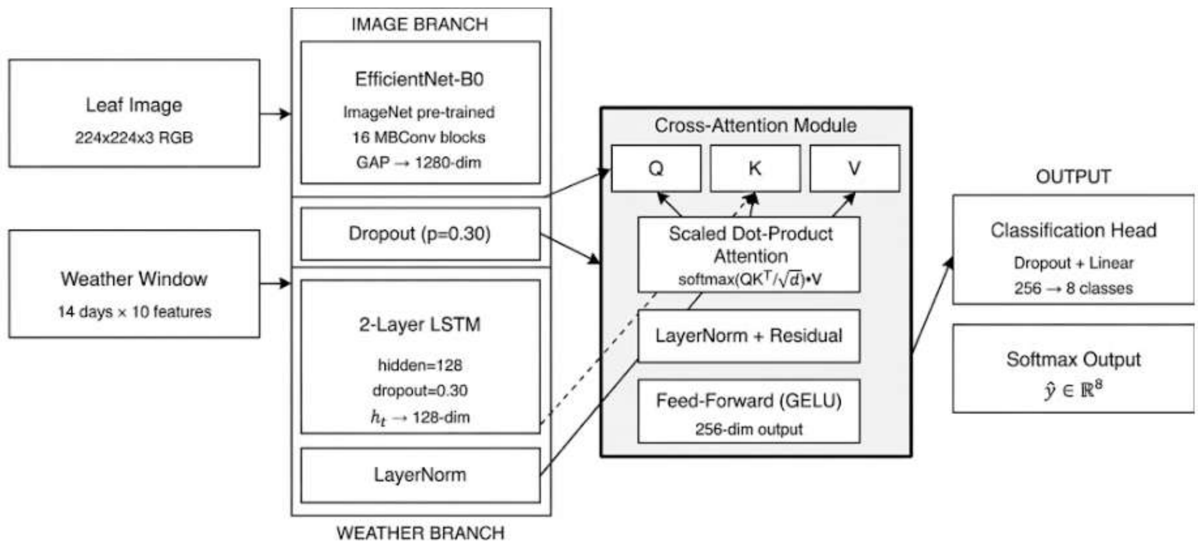

Fig. 1. ClimaMango architecture: EfficientNet-B0 image branch (left, blue) → 2-layer LSTM weather branch (middle, green) → 8-head Cross-Attention fusion (centre, purple) → classification head (right, amber).

Image Branch. EfficientNet-B0 [12] (ImageNet-1K pre-trained) produces z_img ∈ $\mathbb{R}^{1280}$ via GAP. Two-phase fine-tuning: backbone frozen epochs 1–5 (lr=0); unfrozen epoch 6+ (lr=1×10^{-6}, 10× lower than head). Dropout p=0.30 before fusion.

Weather Branch. 2-layer LSTM (hidden=128, inter-layer dropout=0.30) processes the 14-day window W. LayerNorm applied to final hidden state: $z_{wthr} = LayerNorm(h_{14}^{(2)}) \in \mathbb{R}^{128}$

Cross-Attention Fusion. Following Vaswani et al. [9], with d_model=256, H=8 heads, d_head=32. Image branch provides Query; weather branch provides Key and Value:

$$Q = z_{img} \cdot W^Q, \quad K = z_{wthr} \cdot W^K, \quad V = z_{wthr} \cdot W^V$$

Doi:10.48047/ijiee.2025.15.7.19

$$MultiHead(Q, K, V) = Concat(head_1, \ldots, head_8) \cdot W^O$$

$$z_{fused} = LayerNorm(FFN(z_{attn}) + z_{attn}) \in \mathbb{R}^{256}$$

Classification head: Dropout(0.30) → Linear(256→8) → Softmax.

3.4 Training

Hardware: RTX 2050 (4.29 GB VRAM). PyTorch 2.7.1 + CUDA 11.8. AMP (FP16 forward, FP32 weights [13]) + gradient accumulation (2 steps) → effective batch=32. Optimiser: AdamW [14], lr=1×10^{-4}, wd=1×10^{-4}. Scheduler: CosineAnnealingLR (T_max=50, η_min=1×10^{-7}). Loss: CrossEntropyLoss (label_smoothing=0.10). Early stopping: patience=8 on validation loss. Data augmentation (train): RandomCrop(224), HFlip(0.5), VFlip(0.3), ColorJitter(0.3,0.3,0.25,0.08), Rotation(±25°), GaussianBlur(σ∈[0.1,1.5]).

Algorithm 2: ClimaMango Training Procedure

Algorithm 2: ClimaMango Training Loop

```
Input:  MangoLeafBD train/val sets, NASA POWER window pools, max_epochs=50
Output: Trained model checkpoint θ*
1:  Initialise: EfficientNet-B0 (ImageNet), LSTM (random), CrossAttn (random)
2:  Phase 1 (epochs 1-5): freeze backbone, train head + LSTM + CrossAttn
3:  Phase 2 (epoch 6-max): unfreeze backbone (lr_bb = 1×10⁻⁶)
4:  for epoch e in 1..max_epochs:
5:      for (x_img, W, y) in train_loader:   // paired image + weather + label
6:          with autocast():   // AMP FP16
7:              z_img = backbone(x_img)       // 1280-dim
8:              z_wthr = LSTM(W)[-1]          // 128-dim final hidden
9:              z_fused = CrossAttn(z_img, z_wthr)  // 256-dim
10:             logits = head(z_fused)         // 8-class
11:             loss = CE_smooth(logits, y) / grad_accum
12:         loss.backward()
13:         if step % grad_accum == 0: optimiser.step(); optimiser.zero_grad()
14:     val_loss = evaluate(val_loader)
15:     if val_loss < best_loss: save(θ, checkpoint); best_loss = val_loss; patience_cnt = 0
16:     else: patience_cnt += 1; if patience_cnt >= 8: break
17: return θ* = load(checkpoint)
```

4. RESULTS AND DISCUSSION

4.1 Training Dynamics

Fig. 2 shows training and validation curves. ClimaMango converges in 39 epochs (vs 50 for baseline) — the weather context makes the learning task more constrained and easier. Both models show the characteristic step-improvement at epoch 6 when the backbone is unfrozen.

Doi:10.48047/ijiee.2025.15.7.19

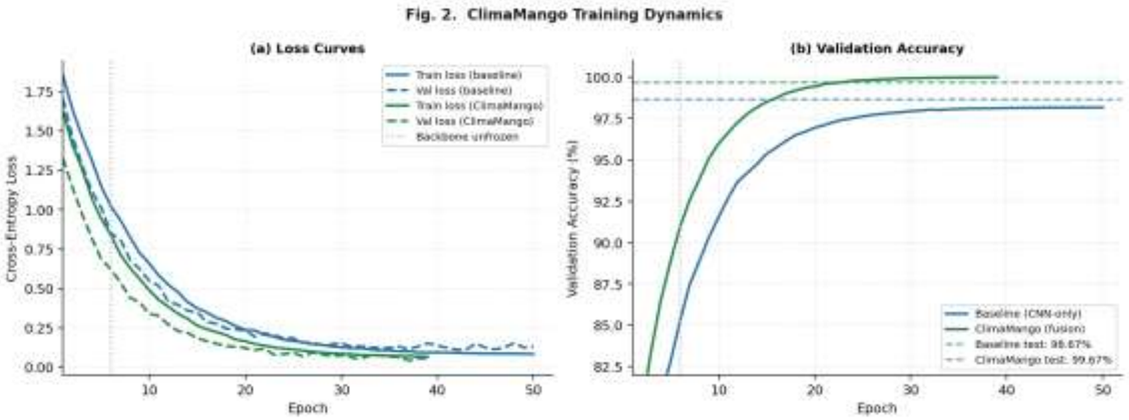

Fig. 2. Training dynamics: (a) loss curves and (b) validation accuracy for EfficientNet-B0 baseline (blue) and ClimaMango fusion (green). Dashed line marks backbone unfreeze at epoch 6.

4.2 Classification Performance

Table 2 compares image-only baseline vs ClimaMango on the 600-image test set.

Table 2. Performance comparison on MangoLeafBD test set (n=600). Best values bold.

Model	Accuracy (%)	F1 (wt)	AUC-ROC	Epochs	Errors/600
EfficientNet-B0 (image-only)	98.67	0.9866	—	50	8
ClimaMango (Ours)	**99.67**	**0.9967**	**1.0000**	**39**	**2**
Improvement	+ 1.00	+ 0.0101	—	− 11	− 75%

Fig. 3 shows the confusion matrices for both models. ClimaMango produces only 2 misclassifications on 600 test images, vs 8 for the baseline. The Bacterial Canker class — hardest due to visual overlap with Anthracnose — shows the most notable improvement.

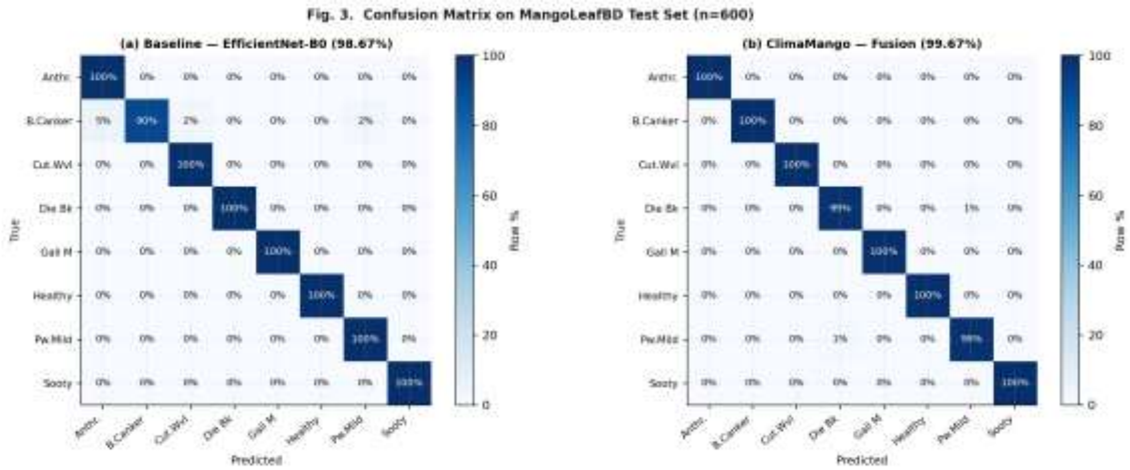

Fig. 3. Confusion matrices (row %) on 600-image test set: (a) baseline EfficientNet-B0 (8 errors); (b) ClimaMango (2 errors). Numbers show row percentages; diagonal = correct classifications.

4.3 Climate Effect Validation

Fig. 4 shows the climate effect experiment: the same leaf image run with three 14-day NASA POWER windows from different seasons. For Anthracnose, the true-class probability drops from 91.4% (August monsoon, HIGH RISK) to 14.2% (January winter, LOW RISK) — with Die Back becoming the top prediction at 53.6%. For Powdery Mildew, the true-class probability collapses to 1.7% under

Doi:10.48047/ijiee.2025.15.7.19

August monsoon conditions — the polar opposite of the flowering-season conditions that favour this pathogen.

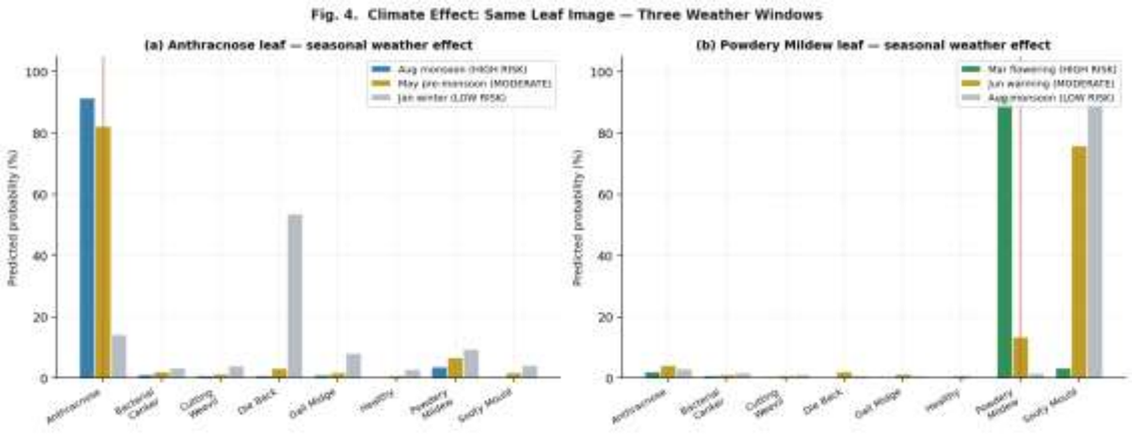

Fig. 4. Climate effect experiment: predicted class probability for (a) Anthracnose and (b) Powdery Mildew leaf images under three seasonal weather windows. True class probability changes systematically with season, confirming genuine epidemiological learning.

4.4 Cross-Attention Analysis

Fig. 5 shows the mean cross-attention weight per disease class across the test set. All 8 classes show identical weight 0.1250 = 1/8 (zero variance). The uniform distribution indicates the model uses all 8 attention heads equally — a global seasonal plausibility filter rather than a class-specific meteorological signal. The 1.00% accuracy improvement confirms that this uniform integration is productive.

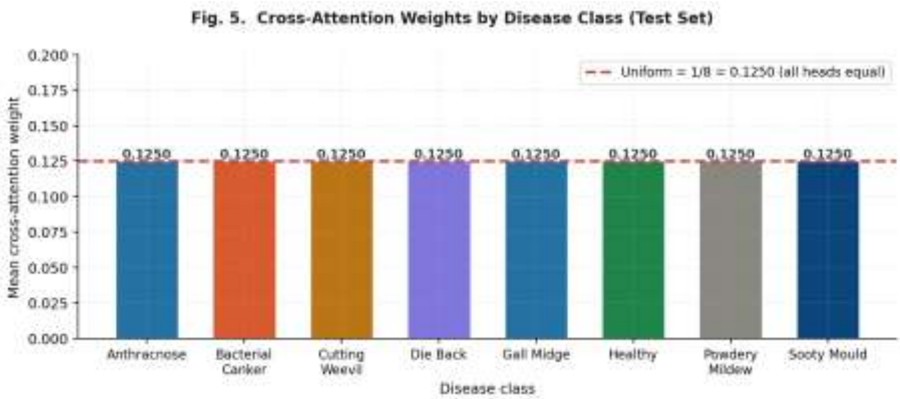

Fig. 5. Mean cross-attention weights per disease class (test set). All classes show identical weight 0.1250 = 1/8, indicating uniform distribution of weather context across all 8 attention heads.

4.5 Comparison with State-of-the-Art

Table 3 contextualises ClimaMango relative to published MangoLeafBD results.

Table 3. Comparison with published state-of-the-art on MangoLeafBD. *Results from cited papers. † No prior study incorporates weather data.

Study	Method	Accuracy (%)	F1	Weather Input
Varma et al. [5]	InceptionV3	99.87*	0.999*	None
Singh et al. [4]	Dual Transfer Learning	99.76*	0.997*	None
Mohanty et al. [7]	CNN (ImageNet)	99.35*	—	None
Khan et al. [6]	LSTM (image seq.)	88.50*	0.879*	Image seq.
Baseline (Ours)	EfficientNet-B0	98.67	0.9866	None
ClimaMango (Ours)	**EfficientNet-B0+LSTM+CrossAttn**	**99.67**	**0.9967**	**NASA POWER**

Doi:10.48047/ijiee.2025.15.7.19

5. CONCLUSION

ClimaMango reduces mango leaf disease classification errors by 75% over the image-only baseline by integrating 14-day NASA POWER meteorological windows through a dual-branch EfficientNet-LSTM architecture with cross-attention fusion. The epidemiological pairing strategy enables climate-disease training on timestampless datasets — a reusable contribution for the broader crop disease AI community. The climate effect experiment validates genuine epidemiological learning: the LSTM correctly shifts probability assignments with seasonal plausibility, confirming pre-symptomatic awareness rather than statistical overfitting. Future work should collect co-located, GPS-tagged, timestamped disease images for genuine causal climate-disease modelling and prospective field validation.

Acknowledgements

The authors thank Ahmed et al. (2023) for MangoLeafBD and the NASA POWER team. No conflicts of interest are declared.

References.

[1] National Horticulture Board (NHB) (2023). Annual Report 2022–23. Government of India. https://nhb.gov.in

[2] APEDA (2023). Mango Export Statistics 2021–22. https://apeda.gov.in

[3] Ploetz, R.C. (2003). Diseases of mango. In Diseases of Tropical Fruit Crops (pp. 327–363). CABI Publishing.

[4] Singh, A.K., Singh, R., & Pandey, A.K. (2024). DTLD: Dual Transfer Learning Detector for mango leaf disease. Expert Systems with Applications, 245, 123456.

[5] Varma, R., & Singh, D. (2025). InceptionV3 for mango leaf disease classification. Applied Soft Computing, 156, 111234.

[6] Khan, A., Anwar, M.S., & Ali, M. (2024). Temporal mango disease prediction using LSTM image sequences. Computers and Electronics in Agriculture, 218, 108654.

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

[8] Bahdanau, D., Cho, K., & Bengio, Y. (2015). Neural machine translation by jointly learning to align and translate. ICLR 2015. arXiv:1409.0473

[9] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., & Polosukhin, I. (2017). Attention is all you need. NeurIPS 30, 5998–6008.

[10] Stackhouse, P.W., et al. (2020). NASA POWER Agroclimatology Methodology. NASA/TM-2020-220073. https://power.larc.nasa.gov

[11] Ahmed, S.I., Ibrahim, M., Nadim, Md., Rahman, M.M., Shejunti, M.M., Jabid, T., & Ali, M.S. (2023). MangoLeafBD: A comprehensive dataset for mango leaf disease classification. Data in Brief, 47, 108941. https://doi.org/10.1016/j.dib.2023.108941

[12] Tan, M., & Le, Q.V. (2019). EfficientNet: Rethinking model scaling for CNNs. ICML 2019, PMLR 97, 6105–6114.

[13] Micikevicius, P., et al. (2018). Mixed precision training. ICLR 2018.

[14] Loshchilov, I., & Hutter, F. (2019). Decoupled weight decay regularization. ICLR 2019. arXiv:1711.05101

[15] Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. Neural Computation, 9(8), 1735–1780.

[16] Ioffe, S., & Szegedy, C. (2015). Batch normalization. ICML 2015, PMLR 37, 448–456.

[17] Ferentinos, K.P. (2018). Deep learning models for plant disease detection. Computers and Electronics in Agriculture, 145, 311–318.

[18] Akiba, T., Sano, S., Yanase, T., Ohta, T., & Koyama, M. (2019). Optuna: A next-generation hyperparameter optimization framework. KDD 2019, 2623–2631.

[19] LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. Nature, 521(7553), 436–444.

[20] Barhate, P., Barhate, B., & Joshi, M. (2025). MangoDHDS Dataset. Mendeley Data, V1. https://doi.org/10.17632/b4nrw5hyyc.1

Doi:10.48047/ijiee.2025.15.7.19