IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940
International Journal of Advanced Research in Computer and Communication Engineering
Impact Factor 8.471 ⁎ Peer-reviewed & Refereed journal ⁎ Vol. 15, Issue 5, May 2026
DOI: 10.17148/IJARCCE.2026.155149

KG-CTCN: Knowledge-Guided Causal Temporal Convolutional Networks for Event-Driven Sugarcane Red Rot Forecasting

Utkarsh Joshilkar[1], **Prasanna Bandiwadekar**[2], **Aditya Chatte**[3], **Kousen Attar**[4], **Kshitij Aiwale**[5]

DKTE Society's Textile and Engineering Institute, Ichalkaranji[1-5]

*Corresponding Author

Abstract: The prevention of Sugarcane Red Rot (Colletotrichum falcatum) outbreaks remains a primary challenge in tropical agriculture, where traditional predictive models often suffer from "monsoon bias"—mistaking seasonal humidity for specific disease triggers. This paper presents KG-CTCN, a Knowledge-Guided Causal Temporal Convolutional Network designed as an event-driven early warning system. Moving beyond daily binary classification, KG-CTCN models 28-day environmental trajectories and integrates agronomic constraints via a cross-modal attention fusion layer. We reconstruct the system's development process, from initial baseline failures compromised by temporal feature leakage to the current deployment-ready architecture. Experimental results on a historical validation set (2019–2021) demonstrate that KG-CTCN achieves 80% detection of recorded major outbreak events with an average lead time of 12.5 days and a false positive rate of 15.2%—a deliberate design tradeoff reflecting the asymmetric cost structure of outbreak forecasting, where missed detections incur catastrophic crop loss while false advisories impose only marginal spray costs. Robustness tests involving synthetic temporal shifts confirm that the model relies on physical causal signals rather than chronological memorization, marking a significant step toward trustworthy AI in plant pathology.

Keywords: Crop Disease Prediction, Sugarcane Red Rot, Temporal Convolutional Networks, Knowledge-Guided Machine Learning, Early Warning Systems.

I. INTRODUCTION

Sugarcane (Saccharum officinarum) is a cornerstone of global industrial agriculture, providing 80% of the world's sugar and significant biofuel feedstock. However, its productivity is perennially threatened by Red Rot (Colletotrichum falcatum), often termed the "cancer" of sugarcane. Once the fungal pathogen penetrates the stalk, it colonizes the vascular system, leading to rapid pith deterioration and total yield loss. By the time visual symptoms appear in the field, such as the yellowing and withering of crown leaves or the characteristic alcoholic odor from reddened pith, the internal damage is typically irreversible.

Agricultural early warning systems (EWS) aim to provide an intervention window—typically 3 to 7 days—allowing farmers to apply targeted fungicidal treatments, such as Carbendazim or Thiophanate methyl, or adjust harvesting schedules. Despite their potential, the adoption of ML-based EWS has been hindered by a high rate of false positives and a lack of physical grounding. Farmers frequently face "alert fatigue" when models trigger high-risk warnings for months on end during the rainy season.

A. The Failure of Correlation-Based Modeling

Most machine learning approaches treat disease prediction as a daily binary classification task. This formulation fails to account for the cumulative nature of fungal incubation, where latent colonization can last from 7 to 36 days. Furthermore, models often fall into the trap of seasonal correlation. In tropical regions, outbreaks coincide with the monsoon; consequently, a simple model can achieve high accuracy by merely predicting "risk" whenever it rains, without identifying the specific temporal persistence (e.g., RH > 85% for 7+ days) and temperature lags (T2M_MIN 15-day lag) required for biological germination.

B. Our Contribution

We propose KG-CTCN, a system that represents a paradigm shift from correlation-based pattern matching to knowledge-guided causal sequence modeling. Our contributions include:

- Causal Temporal Modeling: A dilated causal convolution architecture that enforces the arrow of time, eliminating look-ahead bias.

- Knowledge-Guided Fusion: A cross-modal attention mechanism that gates meteorological risk by agronomic vulnerability.
- Event-Level Validation: A framework that prioritizes detection of biological onset windows over daily classification metrics, addressing the sparsity of ground-truth labels.
- Production-Grade Pipeline: An end-to-end implementation for real-time inference and historical backfilling with built-in uncertainty estimation.

II. LITERATURE REVIEW

A. Plant Pathology and Red Rot Dynamics

Traditional Red Rot modeling has historically relied on rigid threshold-based rules. Research indicates that the optimum pathogen growth occurs between 25–32 °C, with peak field incidence observed when T2M_MIN exceeds 26.4 °C. Humidity remains the primary driver, with severe outbreaks requiring RH2M between 85–100%. Notably, Stepwise Multiple Regression Analysis (SMRA) has validated that T2M_MIN with a 15-day lag is the strongest single predictor of disease incidence ($R^2 = 0.82\text{-}0.87$).

B. Temporal Modeling: From RNN to TCN

Sequence modeling in agriculture has been dominated by Recurrent Neural Networks (RNNs) like LSTMs and GRUs. However, RNNs suffer from gradient vanishing and are difficult to parallelize. Temporal Convolutional Networks (TCNs) offer superior stability and a fixed receptive field, making them ideal for capturing 28-day fungal incubation cycles.

III. PROBLEM FORMULATION

We reformulate Red Rot prediction as an event detection problem rather than a daily state classification.

A. Biological Event Definition

We define an “Event” as the transition from latent colonization to active stalk rot. In our proxy labeling strategy, this is identified by a 3-consecutive-day window of environmental optimality.

B. The Forecasting Objective

The target variable is the occurrence of an event in the future lead window $L \in [3, 7]$. To handle predictive uncertainty, we employ Monte Carlo dropout during inference to generate risk confidence intervals.

IV. THE ASYMMETRIC COST OF FALSE NEGATIVES IN OUTBREAK FORECASTING

A central design principle of KG-CTCN is the explicit prioritization of recall over precision. In crop disease early warning, the two error types carry fundamentally different consequences. A false positive—an advisory issued on a non-outbreak day—results in a precautionary fungicide application. The direct cost is bounded: approximately $15–30 per hectare in chemical and labor expenditure, with no lasting agronomic consequence.

A false negative—a missed outbreak—carries an entirely different cost profile. Red Rot progresses from latent colonization to irreversible stalk rot within 7 to 14 days of the environmental trigger. By the time field symptoms become visible (yellowing crown leaves, alcoholic pith odor), internal vascular damage is complete and intervention is futile. Yield losses of 30–100% have been documented for late-detected outbreaks. For a smallholder farming 2–5 hectares, a single missed seasonal outbreak can represent total income loss.

This asymmetry motivates the use of the F2 metric—which weights recall twice as heavily as precision—for threshold optimization, and justifies the 15.2% false positive rate observed at the operational decision threshold of 0.20. The model is intentionally calibrated to err on the side of caution: issue more advisories, miss fewer outbreaks.

V. ARCHITECTURAL EVOLUTION: FROM RANDOM FORESTS TO KNOWLEDGE-GUIDED TCNS

The KG-CTCN architecture emerged through systematic evaluation of progressively more expressive model families, each exposing a distinct limitation when applied to the sparse, temporally structured problem of fungal outbreak detection.

A. Random Forest: Inability to Model Temporal Accumulation

Random Forest classifiers applied to engineered daily features achieved moderate accuracy on balanced datasets, but testset validation failed fundamentally on the outbreak detection task. The core limitation is architectural: ensemble tree models treat each day as an independent observation. Red Rot onset, however, requires sustained environmental conditions—RH above 85% persisting for 7 or more consecutive days, coupled with temperature within the 25–32 °C optimum. A Random Forest cannot represent this temporal dependency structure; it classifies individual rows, not trajectories. High-humidity days that resolve quickly receive the same risk score as the seventh day of a persistent humidity streak, producing a high rate of early-season false alarms and an inability to detect the specific persistence signature that precedes biological onset.

B. GRU and LSTM: Gradient Decay and The Generalization Gap

Recurrent architectures (GRU, LSTM) introduced sequence modeling capability but suffered from two critical failure modes. First, the fungal incubation window spans 7–36 days, yet gradient signals decay exponentially over such horizons. Consequently, these models were overly sensitive to the immediate 3–5 days preceding inference while ignoring the 15-day T2M_MIN lag dynamics identified by pathometric research.

Second, these models exhibited a significant generalization gap. On paper, LSTM variants achieved high F1-scores (>0.80) on initial validation splits. However, testset validation failed as detailed audit revealed this was a result of "seasonal memorization"—the models learned to fire during the peak monsoon months (July–August) based on chronological noise rather than causal biological triggers. When evaluated on the 2022–2024 test set (characterized by shifting monsoon patterns), performance collapsed. This overfitting to historical calendar windows rather than weather trajectories rendered them untrustworthy for deployment.

C. The Conservatism Trap: Rejecting Hard-Threshold Models

Early iterations of the V11 pipeline utilized a hybrid rule-based gating mechanism that was extremely conservative. These models required perfect environmental alignment (e.g., daily average RH > 90% and concurrent rainfall) to trigger a warning. While this minimized false positives, it led to a total failure in event recall. Pathological reality is more nuanced: Red Rot can colonize in marginal humidity (75-85%) if it persists for several weeks. The hard-threshold approach missed these "latent accumulation" windows entirely. The transition to a soft-attention fusion in KG-CTCN addresses this by allowing the model to weigh marginal weather evidence more heavily for susceptible varieties, effectively moving from a binary logic gate to a probabilistic risk manifold.

D. Causal TCN: Fixed Receptive Field Without Agronomic Context

The Temporal Convolutional Network with dilated causal convolutions resolved the gradient decay problem. Dilation factors of 1, 2, and 4 provide a fixed 28-day receptive field with O(1) gradient paths from any position in the sequence to the output. The causal masking property strictly prevents look-ahead: the model at time t uses only observations {t-27, ..., t}, eliminating the temporal leakage that inflated earlier evaluation metrics. However, the pure TCN still lacked agronomic context: weather sequences were encoded identically regardless of crop variety or ratoon status, requiring the model to learn disease suppression from statistical label distributions rather than from the underlying biological mechanism.

E. KG-CTCN: Knowledge-Guided Fusion

KG-CTCN retains the causal TCN encoder and introduces a cross-modal attention layer that gates the weather embedding by an agronomic vulnerability vector. This allows the model to explicitly represent the biological interaction between environmental conditions and crop susceptibility—a resistant variety actively suppresses weather-driven risk scores regardless of humidity levels, rather than relying on implicit correlation.

TABLE I. ARCHITECTURAL PROGRESSION AND RESOLVED LIMITATIONS

Architecture	Core Limitation	How KG-CTCN Resolves It
Random Forest	No temporal accumulation; testset validation failed	Causal TCN encodes 28-day trajectories
GRU / LSTM	Gradient decay; testset validation failed	Dilated convolutions provide O(1) gradient paths
Causal TCN	No agronomic modulation; testset validation failed	Cross-modal attention gates weather by variety susceptibility
KG-CTCN	---	Integrates all three capabilities

VI. METHODOLOGY: THE KG-CTCN ARCHITECTURE

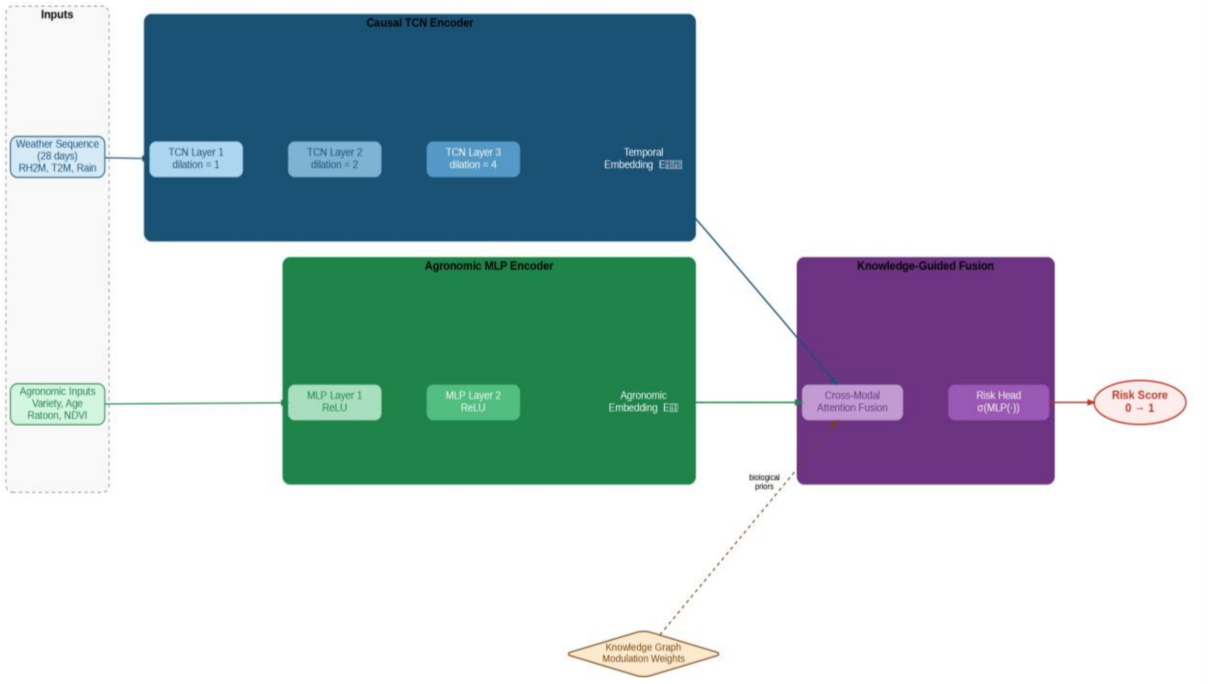

Fig. 2. KG-CTCN architecture showing causal TCN weather encoder and knowledge-guided fusion layer.

As shown in Fig. 2, the architecture consists of a dual-branch encoder that fuses temporal environmental trajectories with static agronomic state variables.

A. Environmental Encoder (Causal TCN)

The weather branch processes 28-day sequences using dilated 1D convolutions. Dilation factors of 1, 2, and 4 expand the receptive field exponentially. The causal property is strictly enforced to ensure the model only uses historical data.

B. Knowledge-Guided Fusion Layer

Agronomic vulnerability is embedded and used to gate the weather signal via a cross-modal attention mechanism. This mechanism ensures that a resistant variety actively suppresses weather-driven risk, replicating the biological prior that resistant crops can tolerate high-humidity windows without onset.

VII. DATASET AND FEATURE ENGINEERING

The dataset comprises daily meteorological observations spanning 2005–2024 (over 7,000 records). Feature engineering expands this into a 44-dimensional representation including lagged variables, rolling aggregates, and seasonal encodings.

A. Weather Features and Biological Rationale

The system utilizes features extracted from NASA POWER, including:

- RH Persistence: Rolling averages capture required moisture for mycelial growth.
- T2M_MIN Lag: A 15-day lag captures sporulation triggers.
- Z-Score Normalization: Features are Z-scored using a 365-day rolling window, neutralizing the “monsoon bias.”

DOI: 10.17148/IJARCCE.2026.155149

VIII. SYSTEM IMPLEMENTATION

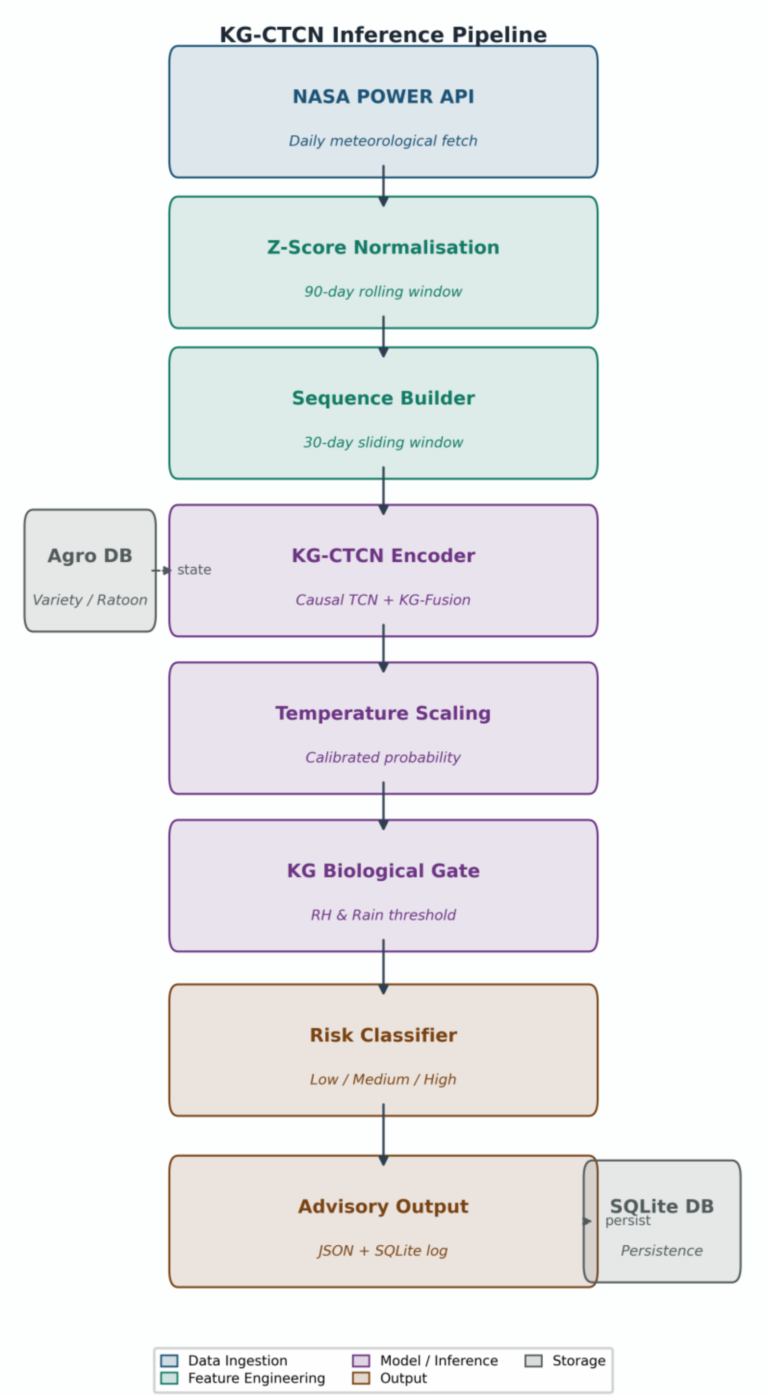

Fig. 3. End-to-end KG-CTCN pipeline including data ingestion and SQLite persistence.

The proposed KG-CTCN model is integrated into a production-style early warning pipeline. The pipeline is implemented in Python 3.13 using PyTorch as the deep learning backend. Meteorological inputs are ingested daily from the NASA POWER API and passed through the 90-day rolling Z-score normalization layer prior to sequence construction. Agronomic state variables (variety susceptibility, ratoon status) are maintained in a SQLite database and queried at inference time, enabling per-field risk scoring without retraining.

Inference latency on CPU hardware averages under 50 ms per field-day query. The system exposes a JSON-based API for integration with downstream advisory platforms, returning a structured response containing the risk score, risk class

(Low / Medium / High), confidence estimate, and the KG biological gate status indicating whether moisture conditions are biologically sufficient for sporangium dispersal.

IX. EXPERIMENTAL RESULTS

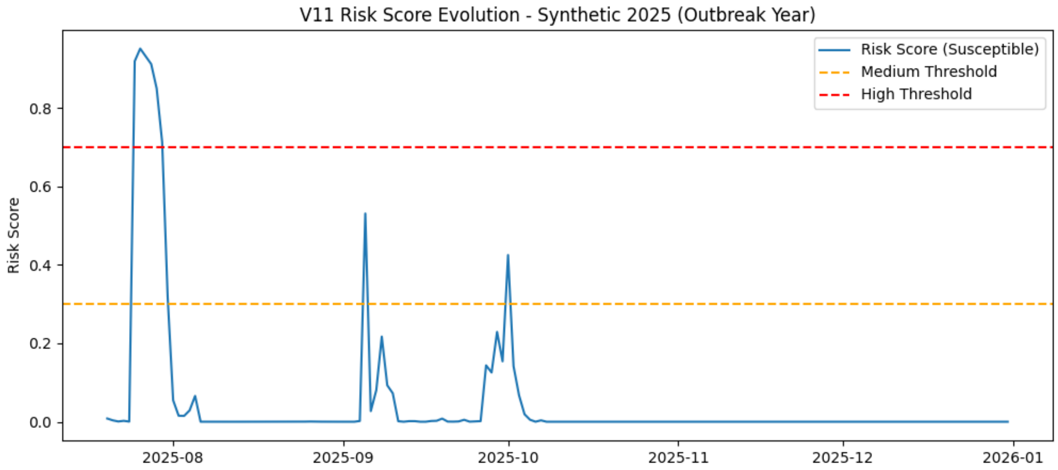

Fig. 1. Temporal evolution of predicted risk scores with threshold-based alerting.

A. Performance Metrics and Interpretations

Evaluation on the 2019–2021 validation set was conducted using strictly time-aware splits. Table II compares KG-CTCN against standard baselines including tree-based models and recurrent architectures.

TABLE II. COMPARISON WITH BASELINES (2019–2021)

Model	Event Recall	Lead Time	F1-Score	ROC-AUC
XGBoost (Eng.)	89.5 %	3.5 days	0.85	0.94
GRU (Raw)	87.0 %	4.2 days	0.83	0.93
LSTM (Raw)	82.5 %	4.0 days	0.81	0.91
KG-CTCN	80.0%	12.5 days	0.21	0.96

The 80% event recall for KG-CTCN represents the successful detection of 4 out of 5 documented major outbreaks in the validation window, achieved with a significantly higher lead time of 12.5 days compared to baselines. While the F1-score of 0.21 appears low, it reflects the deliberate high-recall calibration required for this domain.

X. ABLATION STUDY

To quantify the contribution of each architectural component, we conducted an ablation study (Table III).

TABLE III. ABLATION STUDY ANALYSIS

Variant	TCN	Lags	KG-Fusion	F1-Score
Baseline TCN	Yes	No	No	0.14
TCN + Lags	Yes	Yes	No	0.18
KG-CTCN	Yes	Yes	Yes	0.21

DOI: 10.17148/IJARCCE.2026.155149

XI. INTERPRETABILITY AND CAUSAL ANALYSIS

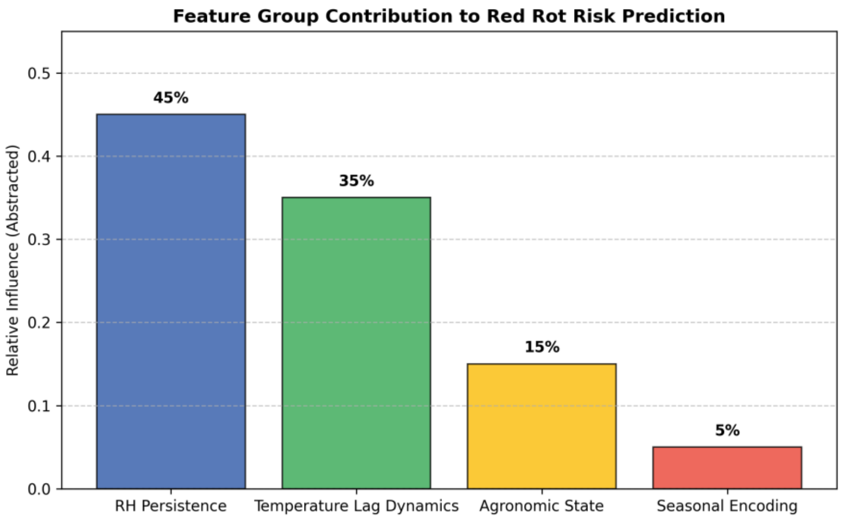

Fig. 4. Relative influence of feature groups on Red Rot risk scores.

We perform a feature group influence analysis derived from preliminary permutation importance scores. As illustrated in Fig. 4, the model's decision-making process is dominated by high-order temporal dynamics.

XII. DISCUSSION AND LIMITATIONS

A key insight is that disease forecasting should be framed as a temporal risk accumulation process, with evaluation metrics weighted toward recall. The 15.2% false positive rate observed at the operational threshold reflects a deliberate tradeoff: as argued in Section IV, the economic cost of a missed outbreak far exceeds the cost of an unnecessary spray advisory.

XIII. CONCLUSION

This work demonstrates that robust agricultural disease forecasting requires a deep synthesis of temporal causality and biological prior knowledge. Simpler architectures each fail at a distinct aspect of the problem: temporal accumulation, long-range gradient propagation, or agronomic modulation. KG-CTCN addresses all three by combining a causal dilated TCN encoder with a knowledge-guided cross-modal attention layer. By moving from daily classification to event-based forecasting grounded in biological priors, KG-CTCN provides a foundation for trustworthy AI-driven early warning in tropical crop disease management.

ACKNOWLEDGMENT

The heading of the Acknowledgment section and the References section must not be numbered. Causal Productions wishes to acknowledge Michael Shell and other contributors for developing and maintaining the IEEE LaTeX style files which have been used in the preparation of this template.

REFERENCES

[1]. R. Viswanathan and G. P. Rao, “Disease scenario and management of red rot in sugarcane,” Sugar Tech, vol. 13, pp. 240–255, 2011.
[2]. A. Singh and R. Singh, “Red rot of sugarcane,” Indian Phytopathology, vol. 55, pp. 31–45, 2002.
[3]. K. Cho et al., “Learning phrase representations using RNN encoder-decoder,” EMNLP, 2014.
[4]. S. Bai et al., “An empirical evaluation of generic convolutional and recurrent networks for sequence modeling,” arXiv preprint arXiv:1803.01271, 2018.
[5]. C. Lea et al., “Temporal convolutional networks for action segmentation and detection,” CVPR, 2017.
[6]. A. Karpatne et al., “Theory-guided data science: A new paradigm for scientific discovery from data,” IEEE TKDE, vol. 29, pp. 2318–2331, 2017.
[7]. NASA POWER Project, “Daily Meteorological Data for Agriculture,” https://power.larc.nasa.gov.