

A Comparative Analysis of Machine-Learning Models for Plant Acoustic Stress Classification

S. Pradeepa Ph.D Scholar¹,

Sri Ramakrishna Arts and Science College, Coimbatore¹

Dr. Hari Prasad D Associate Professor²,

Sri Ramakrishna Arts and Science College, Coimbatore²

Abstract- In this paper, the authors develop a comparative framework for acoustic stress analysis of plants, which is based on a primary sensing modality, acoustic observations, and a secondary one, environmental measurements. The provided evidence includes a dataset containing 3,991 environmental records labeled with 10 classes and a performance comparison table for six different machine learning models: Support Vector Machine (SVM), Random Forest (RF), CatBoost, XGBoost, Improved XGBoost, and Stacking Classifier. The reported accuracies are 79.97%, 82.73%, 82.10%, 82.35%, 82.23%, and 82.60%, respectively. Due to the fact that the attached tabular data does not include the actual acoustic descriptors but rather environmental variables, these numerical results are explicitly reported as an environmental/contextual computational baseline, and are not considered to be proven acoustic-only performance. The manuscript thus offers a leakage-aware, step-by-step approach for gathering acoustic signals, reducing environmental noise, extracting acoustic features, separating acoustic and contextual information, training the six classifiers, and testing plant level generalization. The framework is designed to enable a subsequent controlled acoustic experiment with plant-level splits to independently label and assess the Healthy, Unstressed and Stressed states. The final manuscript keeps the provided results and also keeps a good separation between evidence measured and the proposed acoustic experiment.

Keywords— Plant acoustics, acoustic emission, plant stress detection, machine learning, random forest, XGBoost, CatBoost, stacking, environmental context.

I. INTRODUCTION

Unfavorable conditions, such as water deficit, heat, soil salinity, nutrition imbalance, pathogens and mechanical injury may cause plant stress before external symptoms are apparent, affecting physiological processes. Visual inspection is then only feasible for conventional use for continuous early monitoring. Recent plant-acoustics studies have revealed that plants under stress can emit air-borne ultrasonic signals and that acoustic-emission studies have long established a link between xylem cavitation and hydraulic dysfunction and detectable acoustic events. Recent reviews and monitoring studies have highlighted the potential for acoustic sensing, as well as some key technical challenges such as environmental noise, species specificity, sensor location, variability in sensor acquisition protocols, and the need for valid and reliable computational interpretation [1]–[4]. An accessible path to plant-state classification from high dimensional signal measurements lies in machine

learning. A scientifically defensible acoustic classifier, however, needs to separate acoustic descriptions from the environmental context. This distinction is particularly significant for the present study, since the spreadsheet provided on the website does include environmental variables such as temperature, humidity, soil moisture, soil temperature, light, CO₂, pH, electrical conductivity, watering status, and days since watering—but no raw acoustic feature columns. Therefore, the numerical model comparison provided with the manuscript is not re-labeled as acoustic evidence and is kept as a context condition for comparison. This work aims to define a reproducible framework that will eventually resolve the question: With the acquisition of controlled acoustic signals and leakage safe validation can plant states be distinguished to a reliable enough extent for plant stress analysis at an early stage? The specific contribution is not in creating a new classifier, but rather in integrating and systematically comparing 6 well-established machine-learning approaches in an acoustic-first experimental

design, where explicit separation of acoustic evidence, environmental context and publication claims were introduced.

II. RELATED WORK

Khait et al. revealed that under stress, tomato and tobacco plants can produce ultrasonic sounds in the air and that these sounds carry information that relates to the condition of the plant [1]. Nardini et al. discussed the biological relevance of plant acoustic emissions and highlighted the importance of cavitation signals to assess

water status [3]. Vande Velde et al. performed synthesis for plant-health monitoring, and emphasized the machine learning, signal processing, standardization of protocols and the development of robust sensor networks as important directions [4]. Saha and Rastogi also suggested that plant acoustic emission can be used as an early stress indicator to remotely and integrated monitor agricultural stress [12]. These studies corroborate the sensing concept but cannot eliminate the necessity for controlled data sets and plant level computation verification.

Table-1: Comparison of relevant literature and the present study

Study	Signal / modality	Target	Computational focus	Relevance
Khait et al. [1]	Airborne ultrasonic emissions	Stress in tomato/tobacco	Acoustic recording + ML	Evidence that stress-related sounds can be informative
Nardini et al. [3]	Acoustic emissions	Hydraulic/water status	Physiological interpretation	Cavitation-related acoustic monitoring
Vande Velde et al. [4]	Plant acoustics	Plant health	Signal processing + ML	Need for standardized monitoring and robust sensing
Saha & Rastogi [12]	Acoustic emission	Early stress	Integrated monitoring	AEs as potential early indicators
Paul et al. [10]	Multimodal plant sensing	Plant stress	AI/deep learning	General challenges in stress detection
This study	Acoustic + environmental context	Plant state classification	Six ML models + leakage-aware protocol	Explicit separation of acoustic claims from contextual baseline

III. RESEARCH GAP AND OBJECTIVES

There are few standardized plant-acoustic datasets available that have independent stress ground truth. Acoustic noise from wind, insects, irrigation, equipment, handling and laboratory equipment. Possible data mix-up when recordings from the same plant are randomly assigned to training and test sets. Lack of separation of acoustic evidence from contextual variables. Limited reporting of stress class recall, false alert behaviour and early detection lead time. Demand for easily reproducible comparison between classical and boosting classifiers, on the same data protocol. The goals are: (1) to develop a complete workflow for the acoustic acquisition-to-classification chain, (2) to review and

audit the provided contextual dataset, (3) to recreate the provided 6-model numerical baseline without altering its provenance, (4) to specify a leakage-safe acoustic experiment, and (5) to define evaluation and reporting procedures appropriate for a rigorous journal submission.

IV. DATASET AND DATA AUDIT

The spreadsheet attached has a total of 3991 records and 14 columns. It is described by its name, which is based on the number of records, "10,000" records in this case, but the actual file contains 3,991 rows while the file name indicates 10,000 records; the total number of records in the manuscript is therefore the number of rows observed in the file, not the number

in the file name. The columns are Plant_ID, Plant_Species, Growth_Stage, Air_Temperature_C, Relative_Humidity_pct, Soil_Moisture_pct, Soil_Temperature_C, Light_Intensity_lux, CO2_ppm, Soil_pH, EC_dS_m, Watering_Status, Days_Since_Watering, and Class_Label. Plant_ID is an identifier and is not suitable to use as a predictive feature.

Table - 2: Observed class distribution in the supplied spreadsheet.

Class label	Records
Nutrient_Deficiency	444
Unstressed	418
Pathogen_Stress	415
Severe_Drought	407
Mild_Drought	406
Heat_Stress	400
Healthy	380
Optimal_Growth	378
Recovery	375
Salt_Stress	368




Chart - 1 :Class distribution of the supplied environmental/contextual dataset.

The dataset has a relatively balanced distribution across all 10 plant condition classes, with 368–444 records per class.

Table - 3: Missing values observed in the supplied spreadsheet.

Variable	Missing cells
Light_Intensity_lux	91
Air_Temperature_C	86
EC_dS_m	84
Relative_Humidity_pct	82
Soil_Moisture_pct	78

V. PROPOSED METHODOLOGY

The whole methodology is structured as a series of operations, each operation is traceable and is not a

single model training operation. Controlled acquisition, noise characterization and the acoustic branch is a starting process. Environmental variables that could be used to account for or mask acoustic variation are recorded on the contextual branch. It is noted that the two branches are separate during primary evaluation with an acoustic-only claim not being produced from environmental variables

Simple step-by-step classification of plant acoustic signal. First, data acquisition is done by recording the acoustic signals from plants through use of an appropriate data recording device. In the next step, preprocessing takes place where the acquired signals are preprocessed to eliminate unwanted noise, normalized to ensure consistency, and divided into smaller segments of the signals for analysis.




Figure 1: Methodology for Plant Acoustic Signal Classification

The feature extraction stage involves the identification of informative features, like frequency, energy, spectral, wavelet, and temporal features of the acoustic signals, and their numerical representation. These features are then given to six

machine-learning models (SVM, Random Forest, CatBoost, XGBoost, Improved XGBoost and Stacking Classifier) for the plant-condition classification. Finally, the accuracy, precision, recall, F1-score and confusion matrix are used to assess the performance of all models for classification of plant acoustic conditions, and comparisons are made to verify the accuracy of the models.

The methodology is systematic machine learning based approach for plant acoustic stress analysis. First, suitable microphone or ultrasonic sensor is used to receive the acoustic signals of plants and the necessary plant and environment information. The raw signals are then preprocessed to suppress background noise caused by wind, insects, and machinery, and standardized and divided into fixed-length frames. Then, a set of meaningful acoustic features such as frequency features, energy features, spectral features, wavelet features and temporal features are extracted to create numerical feature vectors. Using 6 machine-learning models—Support Vector Machine (SVM), Random Forest (RF), CatBoost, XGBoost, Improved XGBoost, and a Stacking Classifier—the features are trained and classified to detect the condition of the plants. Finally, the accuracy, precision, recall, F1-score, and class-wise analysis are used for evaluating the models and their performance for plant stress classification is compared.

1. Problem Definition

Let x be an acoustic feature vector representing a plant recording, and y be a plant-state label. The suggested classes for a controlled first experiment are Healthy, Unstressed, and Stressed. Classifier learns $f(x) \rightarrow y$ and contextual variables remain in a separate branch.

2. Acoustic Acquisition

Depending on the frequency of interest use a suitable microphone or ultrasonic/acoustic-emission sensor. Take note of Plant_ID, Recording_ID, species, growth stage, sensor position, time, treatment condition, and environmental measurements. More than one plant and more than one observation will be needed to allow grouping to be accomplished at the plant level.

3. Noise Characterization and Preprocessing

Recognize wind, irrigation, insect, machinery, handling, sensor/electronic noise. Use the justified denoising, normalization and fixed-duration segmentation. The parameters of the preprocessing must be decided without referencing the test set.

4. Acoustic Feature Extraction

Peak frequency, dominant frequency, RMS energy, spectral centroid, spectral bandwidth, spectral entropy, signal-to-noise ratio, spectral flux, continuous-wavelet-transform energy and event duration are examples of candidate descriptors. Mel-spectrograms can optionally be saved as an additional representation. The final feature list should be obtained from existing acoustic data.

5. Contextual Feature Branch

Contextual predictors can include environmental measurements like air temperature, relative humidity, soil moisture, soil temperature, light intensity, CO₂, soil pH, EC, watering status and days since watering. When the outcome is described as “acoustic classification”, they should not be used in place of acoustic features..

6. Missing-Value Treatment

For continuous values, use training-fold statistics for the missing values, and for categorical values, use training-derived mode or an explicit missing category. It is important to note that the imputation model has to be fit within each training fold.

7. Feature Selection

These include correlation filtering, mutual information, tree-based importance, XGBoost importance and SHAP. Must be done within the training data set in order to avoid information leakage.

8. Model Training

We are providing six classifiers: SVM, Random Forest, CatBoost, XGBoost, Improved XGBoost and Stacking Classifier. The tuning of hyperparameters should be carried out within training folds. The exact changes to regularization, feature selection, weighting of classes, parameters of boosting, etc., or optimization

strategy should be documented for the Improved XGBoost variant.

Validation

Only do stratified splitting if observations are independent. If repeated observations are from the same plant, the first generalisation test is plant level grouped splitting. To be noted that independent field or seasonal testing should be kept as a final external validation.

Evaluation

Report accuracy, precision and recall and F1 score. For stress detection, also include the following information: stress-class recall, stress precision, false-alert rate, confusion matrices, and detection lead time compared to an independently defined ground-truth event.

Explainability and Error Analysis

Applying SHAP to select influential features. Differentiate stratification failures by plant species, growth stage, sensor positioning, environmental conditions, and noise level. This will be the model's learning objective from the plant stress or learning the acquisition context.

VI. COMPARATIVE MACHINE-LEARNING MODELS

Table - 4: Comparative Description of Machine-Learning Models

Model	Role in the framework
SVM	Margin-based classifier; useful as a strong non-tree baseline.
Random Forest	Ensemble of randomized decision trees; robust nonlinear baseline.
CatBoost	Gradient boosting with ordered boosting and strong handling of categorical structure.
XGBoost	Regularized gradient-boosted decision-tree framework.
Improved XGBoost	Optimized XGBoost variant; exact modification must be documented.
Stacking Classifier	Combines complementary base learners through a meta-classifier.

VII. EXPERIMENTAL DESIGN

For the ultimate acoustic study to be scientifically interpretable, three experimental conditions should be reported: (A) acoustic only; (B) environmental only; (C) acoustic plus environmental context. This separation quantifies the increment of acoustics and determines if contextual variables are the main drivers of classification. The acoustic experiment should first be kept under control and very precisely defined, for example water stress, and then extended to heat, nutrient, pathogen, salinity and mechanical stress. In a longitudinal design, it is important to capture pre-stress, early-stress, and later-stress recordings. Labels must be based on independent ground truth (treatment schedule, soil moisture, physiological measurements, or validated expert observations). Grouped evaluation: If two groups of recordings are made from the same plant, they should be kept in the same partition.

VIII. RESULTS

As you can see in the accuracy chart provided, there are six model accuracies. Precision, recall and f1 score values are provided in the accompanying manuscript. These numbers are essentially identical to the numbers provided; the accompanying spreadsheet includes environmental variables, not acoustic descriptors.

Table -5: Supplied comparative performance values.

Model	Accuracy (%)	Precision (%)	Recall (%)	F1 (%)
SVM	79.97	78.70	79.98	78.66
Random Forest	82.73	81.71	82.73	81.97
CatBoost	82.10	81.27	82.10	81.48
XGBoost	82.35	81.56	82.35	81.79
Improved XGBoost	82.23	81.34	82.23	81.50
Stacking Classifier	82.60	82.16	82.60	82.25




Chart 2: Supplied comparative accuracy of the six machine-learning models.

The chart is a comparison of accuracy of six machine learning models, and the highest accurate one is Random Forest (82.73%). The accuracy of the models are very close and in the range of 79.97% to 82.73%.

IX. DISCUSSION

The range of accuracy for the baseline provided is 79.97% to 82.73%, which is 2.76 percentage points. The ensemble and boosting methods are very close, suggesting the differences in accuracy in this specific context are fairly small. But it should not be taken as an indication of the superiority of one algorithm over the other in general for plant acoustics. The ranking of the reported variables might not be the same as that of the final acoustic experiment, as acoustic distributions, class overlap, and noise characteristics can vary significantly. The methodological conclusion is that acoustic and environmental information should be treated separately. The framework can be used to uncover that the environmental context is contributing a significant portion of the prediction signal to an environmental-only model that is performing well, but poorly for the acoustic-only model. On the other hand, when validation on a plant-level is performed for grouped sets of data, the acoustic-only model can show evidence of the intended sensing modality, particularly under stress recall, if it is useful. The most important extension to be made is not to add more decimal places to the accuracy values provided, but to use a verified acoustic feature matrix and to report class-wise metrics, grouped validation, independent testing and confidence intervals instead of the contextual benchmark.

Threats to Validity and Limitations

Note that the attached numerical benchmark is not verified acoustic feature columns, but rather environmental/contextual variables. The spreadsheet has 3991 records, but the name of the spreadsheet points to 10,000 records - the number of records in the content of the file as observed is reported. Acoustic emissions can be different for various species, cultivars, growth stage, sensor location, and environmental conditions. Acoustic recordings can be contaminated by wind, insects, irrigation, machinery, handling and electronic noise. If several recordings are repeated from the same plant, then comparing the performance may overestimate due to random row-level splitting. Early stress labels must be handcrafted and cannot be inferred from the acoustic model. Results from the controlled environment may not be directly applicable to field performance.

Future Research Directions

Gather an exclusive set of plant-acoustic data for Healthy, Unstressed and Stress states. Increase the size of the data set by adding to the list of plant species and cultivars, stress types, growth stages, and locations. Compare acoustic-only, environmental-only and multimodal models with the same grouped splits. Improve source-isolation and denoising techniques used in the field. Identify the degree of change in the sound that is detectable before independent visible/physiological changes are present. Identify biologically meaningful acoustic signatures with the aid of SHAP and class-wise confusion analysis. Make independent field validation over seasons and locations prior to deployment claims.

REFERENCES

1. B. Saha and A. Rastogi.(2026). Plant acoustic emission as early stress signals: Towards remote integrated monitoring for sustainable agriculture. *European Journal of Agronomy*, vol. 177, 128100, doi: 10.1016/j.eja.2026.128100.
2. B. Dey and R. Ahmed.(2025). A comprehensive review of AI-driven plant stress monitoring and embedded sensor technology: Agriculture 5.0,

- Journal of Industrial Information Integration*, vol. 47, 100931, doi: 10.1016/j.jii.2025.100931.
3. B. Dey and R. Ahmed.(2025). A comprehensive review of AI-driven plant stress monitoring and embedded sensor technology: Agriculture 5.0, *Journal of Industrial Information Integration*, vol. 47, 100931, doi: 10.1016/j.jii.2025.100931
 4. J. Vande Velde, B. da Silva, A. Touhafi, and L. Segers.(2025). Plant Health Monitoring using Acoustics. *Forum Acusticum/EuroNoise*, pp. 4427–4434, doi: 10.61782/fa.2025.0682.
 5. [5] N. Paul, G. C. Sunil, D. Horvath, and X. Sun. (2025). Deep learning for plant stress detection: A comprehensive review of technologies, challenges, and future directions. *Computers and Electronics in Agriculture*, vol. 229, 109734, doi: 10.1016/j.compag.2024.109734.
 6. Khait et al.,(2023). Sounds emitted by plants under stress are airborne and informative. *CellPress*, vol. 186, no. 7, pp. 1328–1336.e10, doi: 10.1016/j.cell.2023.03.009.
 7. M. Hussain et al.,(2023). Plants can talk: A new era in plant acoustics. *Trends in Plant Science*, vol. 28, no. 9, pp. 987–990, 2023, doi:10.1016/j.tplants.2023.06.014.
 8. A.A. Fernandez-Jaramillo, C. Duarte-Galvan, L. Garcia-Mier, S.N. Jimenez-Garcia, L.M. Contreras-Medina. (2018).Effects of acoustic waves on plants: an agricultural, ecological, molecular and biochemical perspective, *Sci. Hortic.* 235 340–348. PubMed PMID: WOS:000430762600040.
 9. J.Y. Kim, J.S. Lee, T.R. Kwon, S.I. Lee, J.A. Kim, G.M. Lee, et al.,. (2015). Sound waves delay tomato fruit ripening by negatively regulating ethylene biosynthesis and signaling genes, *Postharvest Biol. Technol.* 110 (2015) 43–50. PubMed PMID: WOS:000362919300007.
 10. L. Prokhorenkova et al.,.(2018). CatBoost: Unbiased boosting with categorical features. *Advances in Neural Information Processing Systems* 31.
 11. B. Choi, R. Ghosh, M.A. Gururani, G. Shanmugam, J. Jeon, J. Kim, et al.,. (2017). Positive regulatory role of sound vibration treatment in *Arabidopsis thaliana* against *Botrytis cinerea* infection. *Sci.*
- Rep. 7 (2017). PubMed PMID: WOS:000402312200001.