SvedbergOpen
DISSEMINATIONOFKNOWLEDGE

International Journal of Artificial Intelligence and Machine Learning

Publisher's Home Page: https://www.svedbergopen.com/

Research Paper Open Access

Machine-Learning Prediction of Practitioner-Level Rooftop Gardening Success in Tamil Nadu

Dr.S.Santhana Iyyappa Sundararaj

B.Arch, M.Arch (Landscape), Ph.d, Professor, Kalasalingam School of Architecture, Kalasalingam Academy of Research and Education, Tamil Nadu, India, Email id: iyyappa@klu.ac.in, Orcid ID: 0009-0002-3876-1043

Abstract
Rooftop gardening is increasingly promoted as a multi-functional urban intervention addressing food security, thermal regulation, and community resilience in rapidly urbanising Indian cities. Existing scholarship has concentrated on feasibility assessments, thermal benefit quantification, and policy barrier analysis, leaving the predictability of individual practitioner success empirically under-examined. This study addresses that gap by developing binary classification models of high gardening success using structured questionnaire data from 300 rooftop gardening practitioners across Chennai, Madurai, and Coimbatore, Tamil Nadu. The outcome variable distinguished high-success from non-high-success gardeners, and the continuous overall success score was excluded from predictors to prevent target leakage. Four supervised classifiers: Logistic Regression, Random Forest, Extra Trees, and Gradient Boosting were compared using a stratified 75:25 training-test partition and five-fold stratified cross-validation, with performance assessed through accuracy, precision, recall, F1-score, and ROC-AUC. Gradient Boosting achieved the highest held-out ROC-AUC (0.760) and recall (0.816), whereas Extra Trees achieved the highest mean cross-validated ROC-AUC (0.692 ± 0.062). Random Forest and Extra Trees shared the highest test accuracy (0.707), with Random Forest recording the highest F1-score (0.732). This divergence between held-out and cross-validated rankings indicates that no single algorithm was uniformly superior and that model selection should be guided by the intended application and validation context. SHAP and impurity-based importance analyses identified gardening experience and cost-related constraints as the most influential predictors, with visual plant-health score (a concurrent correlate) and kitchen-waste reuse also contributing. The findings establish a predictive baseline for practitioner-level rooftop gardening success and shift the analytical question from whether rooftop gardening is beneficial toward which practitioner profiles are associated with success, providing a foundation for targeted advisory interventions. Independent external validation and prospective data collection are required before generalisation.

Keywords: rooftop gardening; urban agriculture; machine learning; classification; SHAP; Tamil Nadu

1. Introduction

Rooftop gardening has attracted growing policy and research attention as a multi-functional urban intervention addressing food security, thermal regulation, and community resilience. In India, where rapid urbanisation intersects with climate stress, the practice is increasingly framed as a scalable response to multiple urban challenges. Existing scholarship, however, has concentrated on feasibility assessments, thermal benefit quantification, and policy barrier analysis. What remains empirically under-examined is the predictability of individual practitioner success. This gap constrains the translation of rooftop gardening from scattered grassroots practice into evidence-based intervention strategy.

The present study addresses this gap by shifting analytical focus toward the predictability of success. Using observational data encompassing gardening experience, time allocation, constraint rankings, cultivation management, and harvest outputs, it compares four classifiers in predicting a binary high-success outcome. By excluding the overall success score to prevent target leakage and conducting preprocessing within pipelines, the study seeks to identify combinations of factors exhibiting stable predictive power. The predictive orientation does not claim causality; feature importance rankings reflect associations, not intervention effects. This cautious position is methodologically necessary at a stage when rooftop gardening research lacks large-sample longitudinal data.

2. Related Work

Research on rooftop agriculture has expanded rapidly across agronomic, engineering, and urban planning disciplines, yet the practitioner-level determinants distinguishing successful from unsuccessful gardeners remain insufficiently synthesised. Existing scholarship broadly clusters into four interconnected domains, each of which stops short of modelling individual-level success in real-world domestic settings.

2.1. Growing media and substrate engineering

The distinct microclimatic and structural constraints of rooftop environments, specifically shallow substrate depth, rapid drainage, elevated ambient temperatures, high solar radiation, and wind exposure, make growing media formulation a primary determinant of plant survival [1], [2]. Experimental studies demonstrate that

amending conventional soil or coir substrates with vermicompost and biochar significantly enhances water-holding capacity, cation exchange capacity, microbial activity, and leaf relative water content, yielding up to 26.28% higher fruit output than unamended controls [1], [2]. Similarly, incorporating recycled inorganic aggregates such as crushed brick and foamed glass has been shown to match or outperform commercial rooftop substrates while lowering structural dead loads and environmental footprints [3]. While these controlled trials establish optimal biophysical formulations, they do not explain why household practitioners utilizing comparable commercial potting mixes achieve markedly divergent outcomes.

2.2. Crop selection and agronomic productivity

Within container-based rooftop systems, fruiting vegetables, particularly compact and dwarf cultivars of cherry tomato (Solanum lycopersicum var. cerasiforme) have received substantial attention for their high harvest index and adaptability to spatial constraints [1], [4]. Empirical trials indicate that container-grown rooftop cherry tomatoes can achieve yields exceeding 2,000 g per plant annually, outperforming conventional soil-grown benchmarks [4]. At a broader technological scale, systematic reviews of building-integrated rooftop greenhouses confirm high annual productivity (10–50 kg/m^2), water-use efficiency, and thermal symbiosis with host buildings, while identifying high upfront capital expenditure, structural load limits, and stringent municipal building codes as major barriers [5]. Consequently, productivity literature establishes the upper bounds of agronomic and engineering potential without accounting for the behavioural and management variability of non-commercial practitioners.

2.3. Spatial potential and economic feasibility

At the macro scale, geospatial frameworks integrating high-resolution satellite imagery and building footprint datasets have quantified aggregate urban rooftop capacity and projected economic returns [6]. Recent spatial-economic assessments in Indian tier-II cities, such as Coimbatore, demonstrate substantial city-wide revenue and food-provisioning potential alongside rainwater harvesting and photovoltaic co-benefits [6]. However, because geospatial models operate at the building or municipal polygon level, they treat rooftops as uniform productive surfaces and cannot predict individual household adoption, retention, or yield success.

2.4. Governance, adoption barriers, and the practitioner-level gap

Socio-economic and policy studies in urban India highlight that institutional outreach, hands-on training, and input subsidies can substantially raise adoption rates, whereas lack of technical know-how, structural leakage concerns, and poor follow-up support drive abandonment [7], [8]. Together, existing trajectories map isolated dimensions of the system: agronomic studies define optimal inputs under controlled conditions [1]–[4]; geospatial analyses quantify aggregate urban capacity [6]; and governance studies catalogue structural barriers to entry [5], [7], [8]. What remains missing is a predictive, practitioner-level framework modelling how specific combinations of gardener demographics, management routines, resource constraints, and cultivation choices interact to distinguish high-success from low-success practitioners in real-world domestic environments. This study addresses that gap by formulating rooftop gardening success as a binary classification problem driven by practitioner characteristics, operational practices, and resource constraints.

3. Materials and Methods

3.1 Study Area and Data Collection

The study was conducted across Chennai, Madurai, and Coimbatore, Tamil Nadu, India. Primary data were collected from individuals engaged in gardening activities using a structured questionnaire. A non-probability convenience sampling technique was employed to select respondents based on accessibility and willingness to participate. A total of 310 records were documented. After removal of the outliers, 300 records were included: Chennai (n = 106, 35.33%), Madurai (n = 97, 32.33%), and Coimbatore (n = 97, 32.33%). The collected information encompassed demographic characteristics, gardening experience, infrastructure, growing media composition, cultivation practices, resource utilisation, plant-health conditions, and gardening outcomes. Visual plant-health scores were recorded at the time of questionnaire administration, concurrently with the self-reported gardening outcome; they therefore represent a concurrent correlate rather than a strictly prospective predictor.

3.2 Questionnaire Development and Expert Validation

A structured questionnaire was developed and validated by subject experts from the Tamil Nadu Agricultural University. Expert evaluation ensured that items adequately represented important dimensions of urban gardening and were technically appropriate for rooftop garden information collection. The finalised instrument covered demographic characteristics, gardening infrastructure, growing media, water and nutrient management, mulching and crop management, pest and disease management, constraints, and plant-health outcomes.

3.3 Potting-Mix Standardisation

Respondents reported the relative proportions of soil, cocopeat, compost, and other materials. To ensure comparability, proportions were standardised to a cumulative total of 100% for each mixture:

$$\sum_{i=1}^{n} P_i = 100\%$$

where Pi represents the percentage contribution of each component. This standardisation ensured that component proportions represented complete composition rather than independent quantities.

3.4 Outcome Classification and Dataset Preparation

The dependent variable was the binary high_success_binary: Class 1 (high success) for overall gardening success scores of 4 or 5; Class 0 (non-high success) for scores of 1, 2, or 3. The resulting dataset contained 153 high-success observations (51%) and 147 non-high-success observations (49%). The continuous overall_success_score_1_to_5 was excluded from predictors to minimise target leakage. The final dataset comprised 300 observations and 46 variables. With 153 events in the positive class and 46 candidate predictors, the events-per-variable ratio exceeds 3:1, which is modest but consistent with the exploratory, comparative nature of the modelling exercise; the sample size is acknowledged as a constraint on generalisability in Section 6. The visual plant-health score was retained despite its concurrent timing because it captures observable crop condition rather than a self-assessed success judgment; nevertheless, it is treated as a concurrent correlate rather than a prospective predictor throughout the interpretation. (Sections 4.6, 4.7, and 5.3)

3.5 Machine Learning Platform and Computational Analysis

The analysis was performed using Python (version 3.12.7), with Pandas (version 2.2.2) [9] and NumPy (version 2.0.2) [10] for data manipulation, Scikit-learn (version 1.5.2) [11] for machine learning model implementation, and SHAP (version 0.46.0) for model interpretation [12]. The dataset was partitioned into training (75%) and testing (25%) subsets using stratified sampling. Five-fold stratified cross-validation assessed consistency across data subsets, following standard practice for model selection in classification tasks [13]. Evaluation metrics included accuracy, precision, recall, F1-score, and ROC-AUC. Confusion matrices and ROC curves were generated for comparative assessment.

For model interpretation, the SHAP library was employed to calculate feature-level Shapley values [12]. SHAP's additive property ensures that predictions decompose into feature contributions, providing a theoretically grounded basis for interpreting ensemble models. Random Forest impurity-based feature importance was additionally examined as a complementary measure.

3.6. Analytical Workflow

The workflow followed: primary data collection (n = 300); expert validation and dataset preparation; preprocessing and binary classification; 75:25 stratified partition; model training (Logistic Regression, Random Forest, Extra Trees, Gradient Boosting); five-fold cross-validation and performance evaluation; SHAP and feature-importance interpretation; and final comparative outputs. The overall analytical workflow followed the sequence illustrated in Figure 1.

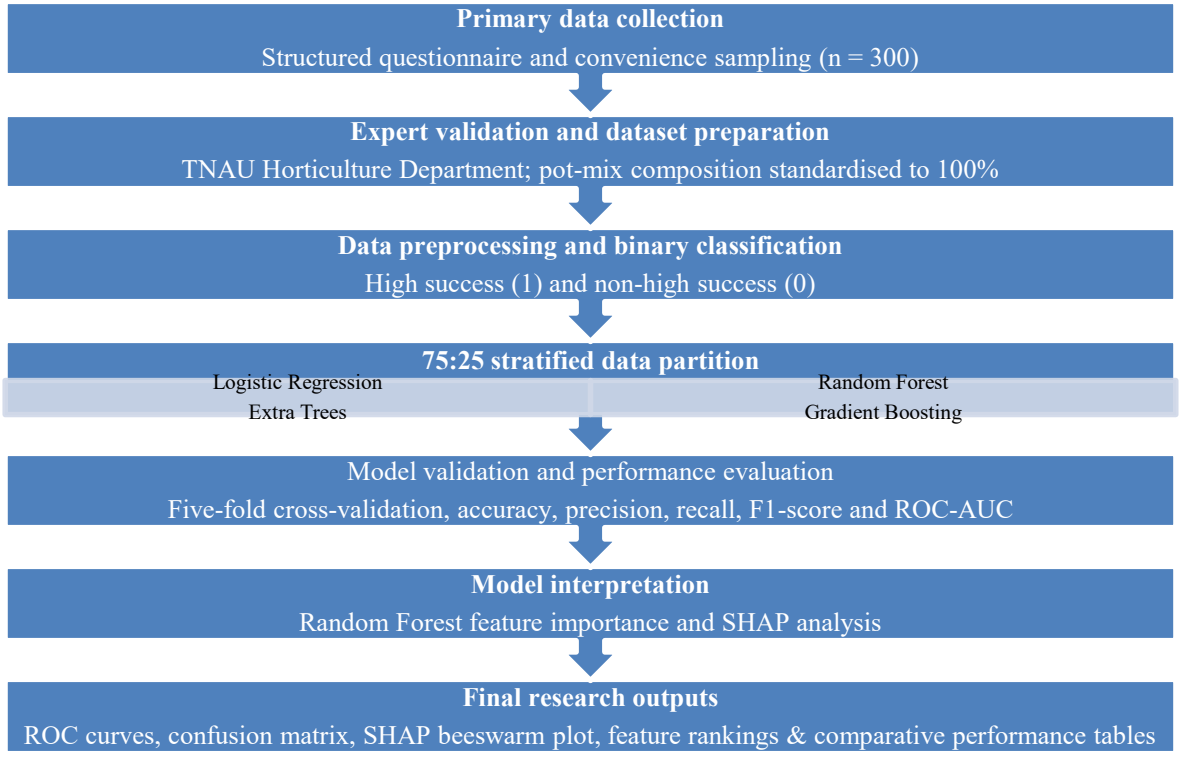


Figure 1. Summary of analytical workflow

3.7 Essential Formulations

Logistic Regression probability:

$$P(y = 1 \mid x) = \frac{1}{1 + e^{-(\beta_0 + \beta_1 x)}}$$

Performance metrics from confusion-matrix counts:

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN}, \text{ Precision} = \frac{TP}{TP + FP}, \text{ Recall} = \frac{TP}{TP + FN}$$

$$\text{F1} = \frac{2 \cdot \text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}}$$

SHAP additive property:

$$f(x) = \phi_0 + \sum_{j=1}^{p} \phi_j$$

where ϕ0is the base value and ϕjis the Shapley value of the j-th feature [12].

4. Results

4.1 Dataset Characteristics

Chennai contributed 106 observations (35.33%), while Madurai and Coimbatore contributed 97 each (32.33%). Regarding gardening experience: beginners (<1 year) comprised 109 respondents (36.33%), intermediate (1–3 years) 112 (37.33%), and expert (>3 years) 79 (26.33%), as illustrated in Table 1. The dependent variable distribution was 153 high-success (51%) and 147 non-high-success (49%), suitable for binary classification without substantial imbalance. All potting-mix components summed to exactly 100% across all 300 observations.

Table 1. Distribution of respondents by city, gardening experience, and success category (n = 300)

Parameter	Category	Frequency	Percentage
City	Chennai	106	35.33
	Madurai	97	32.33
	Coimbatore	97	32.33
Experience	Beginner (<1 yr)	109	36.33
	Intermediate (1–3 yrs)	112	37.33
	Expert (>3 yrs)	79	26.33
Outcome	High success	153	51.00
	Non-high success	147	49.00

4.2 Comparison of Classifier Performance

On the held-out test set, Gradient Boosting achieved the highest ROC-AUC (0.760), indicating the strongest ranking discrimination among the four models for this particular partition. Its accuracy, precision, recall, and F1-score were 0.693, 0.660, 0.816, and 0.729, respectively, as specified in Table 2. Extra Trees and Random Forest yielded the highest observed accuracy (both 0.707), while Random Forest recorded an F1-score of 0.732. Logistic Regression produced a test ROC-AUC of 0.688. These differences indicate that the preferred algorithm depends on the performance criterion: Gradient Boosting ranked highest by ROC-AUC, whereas Extra Trees and Random Forest achieved higher overall accuracy on this test sample.

Table 2. Test-set performance of the four classification algorithms

Algorithm	Accuracy	Precision	Recall	F1-score	ROC-AUC
Logistic Regression	0.640	0.641	0.658	0.649	0.688
Random Forest	0.707	0.682	0.789	0.732	0.701
Extra Trees	0.707	0.690	0.763	0.725	0.725
Gradient Boosting	0.693	0.660	0.816	0.729	0.760

4.3 ROC Curve Analysis

The ROC analysis compares sensitivity against the false-positive rate across all possible classification thresholds. The Gradient Boosting curve produced the largest test AUC (0.760), followed by Extra Trees (0.725), Random Forest (0.701), and Logistic Regression (0.688).

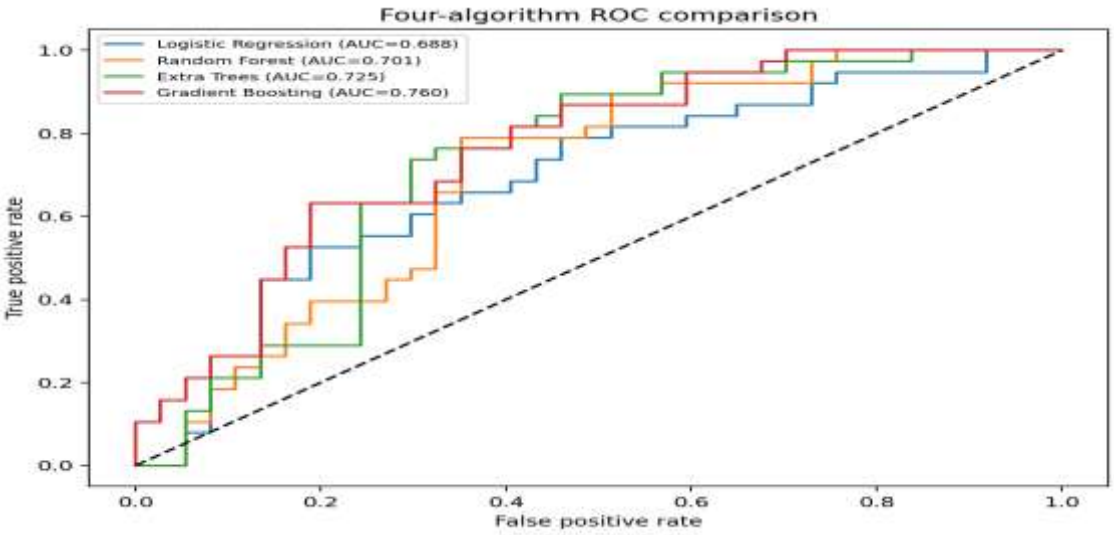


Figure 2. Four-algorithm ROC comparison

These results suggest moderate discrimination rather than near-perfect separation. The AUC ordering is based on a single held-out partition and should not be interpreted as evidence that the same ordering will persist in a new sample. Figure 2. Receiver operating characteristic (ROC) curves for Logistic Regression, Random Forest, Extra Trees, and Gradient Boosting evaluated on the held-out test set. The diagonal line represents no discrimination (AUC = 0.50).

4.4 Cross-Validation Performance

Five-fold stratified cross-validation on the training set revealed a different ranking (Table 3). Extra Trees achieved the highest mean cross-validated ROC-AUC and the highest mean values across the reported performance metrics. However, its standard deviation was not the lowest among the evaluated algorithms. Thus, Extra Trees demonstrated the strongest average cross-validation performance, rather than the greatest fold-to-fold stability. Gradient Boosting recorded the largest ROC-AUC standard deviation (0.070), indicating greater sensitivity to data partition. The discrepancy between held-out and cross-validated rankings underscores the importance of multiple validation strategies and cautions against declaring a universally superior algorithm.

Table 3. Five-fold stratified cross-validation performance (mean ± SD)

Algorithm	Accuracy	Precision	Recall	F1-score	ROC-AUC
Logistic Regression	0.644 ± 0.027	0.643 ± 0.032	0.646 ± 0.094	0.648 ± 0.048	0.677 ± 0.056
Random Forest	0.596 ± 0.058	0.610 ± 0.072	0.617 ± 0.071	0.609 ± 0.044	0.672 ± 0.053
Extra Trees	0.649 ± 0.073	0.649 ± 0.067	0.687 ± 0.104	0.665 ± 0.075	0.692 ± 0.062
Gradient Boosting	0.622 ± 0.047	0.626 ± 0.039	0.643 ± 0.084	0.633 ± 0.059	0.662 ± 0.070

The comparative results did not identify a uniformly superior classifier. Gradient Boosting achieved the highest ROC-AUC on the held-out test set, Random Forest achieved the highest test F1-score, and Extra Trees achieved the highest mean cross-validation performance. Accordingly, model preference depends on whether the primary objective is discrimination, identification of high-success cases, balanced classification performance, or average performance across validation folds.

4.5 Random Forest Confusion Matrix

At the 0.50 probability threshold, the Random Forest classifier correctly classified 23 non-high-success cases (true negatives) and 30 high-success cases (true positives). It incorrectly classified 14 non-high-success cases as high-success (false positives) and failed to identify 8 high-success cases (false negatives). The corresponding test accuracy was 0.707, with a precision of 0.682, a recall of 0.789, and an F1-score of 0.732. The confusion matrix indicates that the model detected a substantial proportion of high-success cases, although some positive cases were missed and some negative cases were incorrectly flagged. Figure 3. Confusion matrix for the Random Forest classifier on the held-out test set at a probability threshold of 0.50. Rows represent observed classes, and columns represent predicted classes.

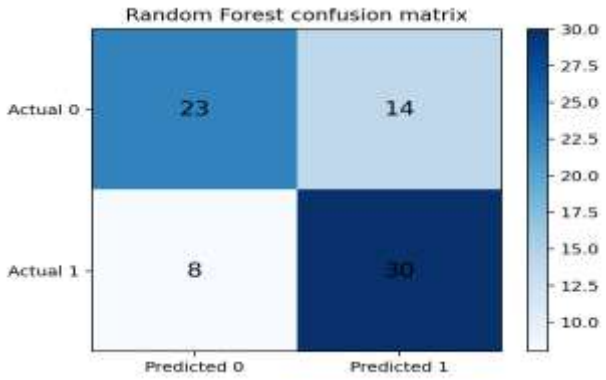


Figure 3. Confusion matrix for the Random Forest classifier

4.6 SHAP Feature Interpretation

The highest mean absolute SHAP contribution was associated with beginner gardening experience (0.0372), followed by encoded visual plant-health category (0.0169) and expert gardening experience (0.0161). The visual plant-health score was retained despite its concurrent timing because it captures observable crop condition rather than a self-assessed success judgment; nevertheless, it is treated as a concurrent correlate rather than a prospective predictor throughout the interpretation.

Other influential encoded features included kitchen-waste reuse, mulching application, geographical location, mulch material, and pest-control methods. The prominence of experience categories suggests that gardening background contributed substantially to classification decisions; however, beginner and expert categories represent different encoded levels of the same categorical variable and should be interpreted separately.

Figure 4. SHAP beeswarm summary plot for the Random Forest classifier. Each point represents an individual observation; the horizontal position indicates the SHAP value (impact on the model output), and the color represents the feature value (red = high/yes, blue = low/no). Features are ranked in descending order of global importance. Positive SHAP values push the prediction toward the high-success class, while negative values push it toward the non-high-success class.

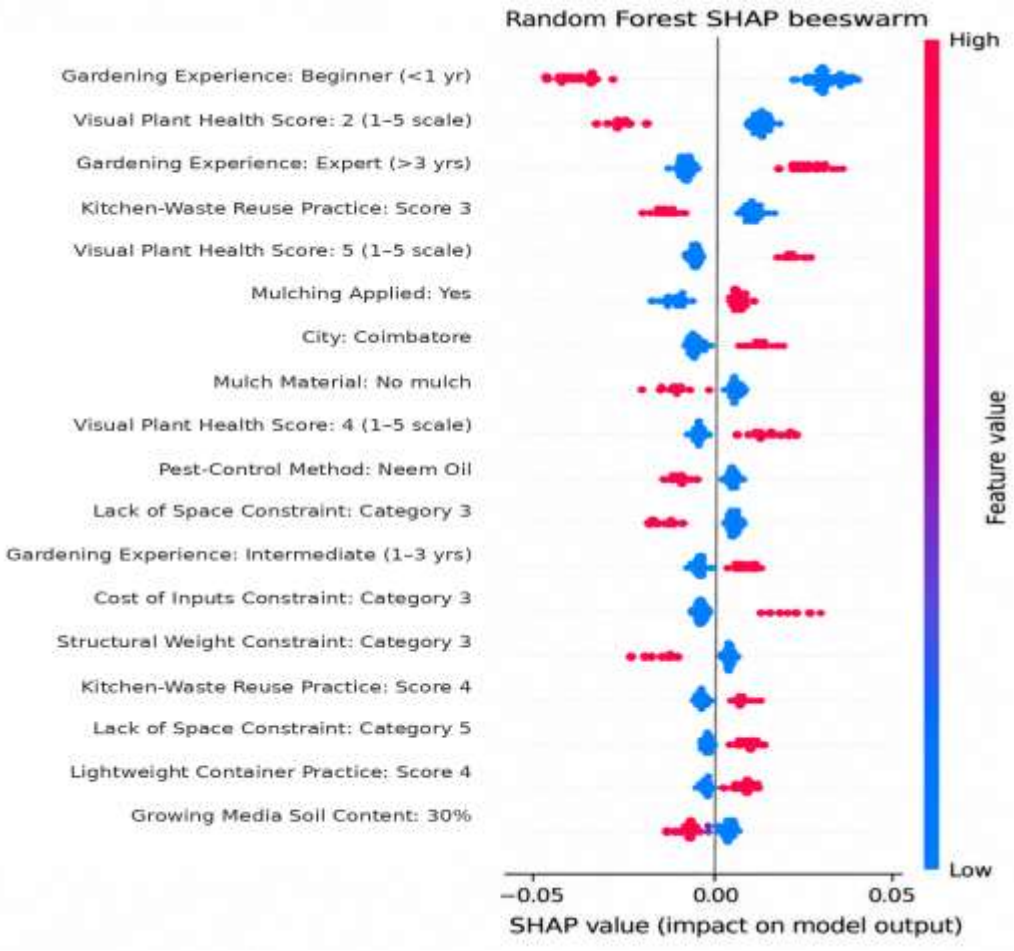


Figure 4. SHAP beeswarm plot of feature importance and effects for the Random Forest model

4.7 Random Forest Feature Importance

Aggregated impurity-based importance identified gardening experience as the leading original predictor (0.0705), followed by visual plant-health score (concurrent correlate) (0.0482) and cost-related constraint rank (0.0327). Kitchen-waste reuse and estimated monthly yield were also prominent. The findings suggest that gardening success is associated with a combination of experiential, plant-related, and resource-management characteristics rather than a single isolated variable.

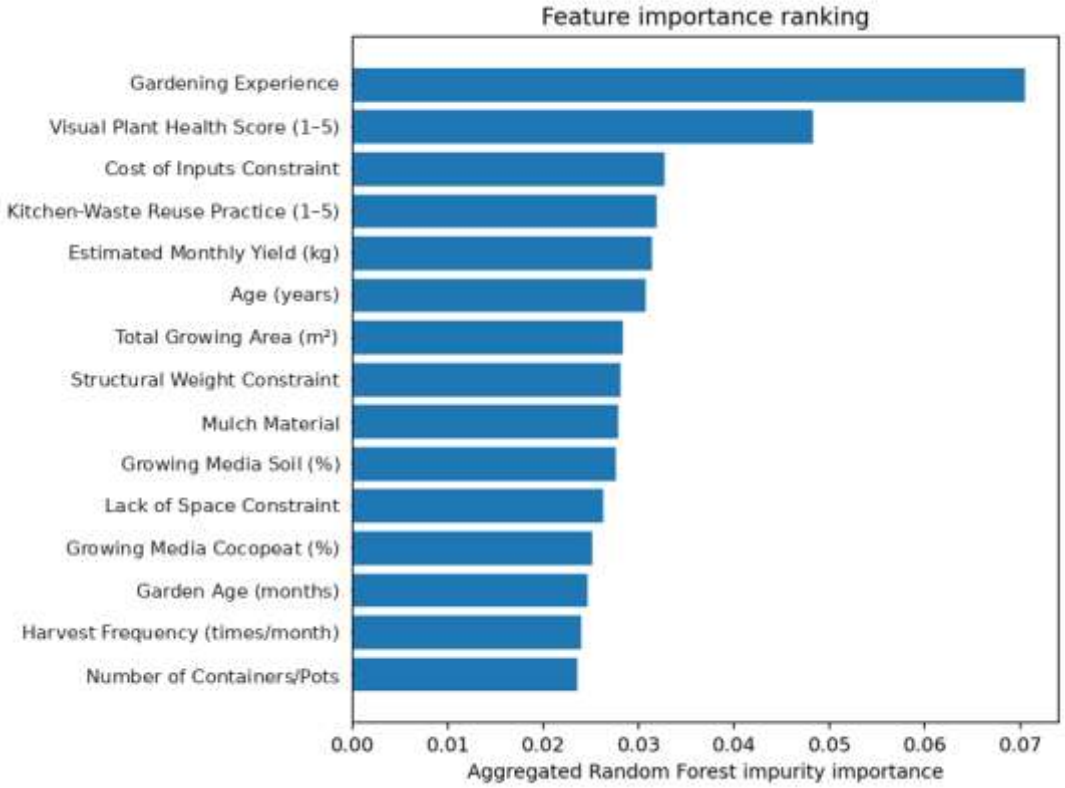


Figure 5. Random Forest-Based Ranking of Predictors Influencing Gardening Success.

Figure 5. Random Forest feature-importance ranking aggregated to the original predictor variables. Higher values indicate greater contribution to the fitted model's impurity reduction.

5. Discussion

5.1 Predictive Performance and Algorithm Selection

Supervised machine learning can identify predictive patterns associated with high gardening success using questionnaire-derived information. Gradient Boosting achieved the highest held-out ROC-AUC (0.760), whereas Extra Trees achieved the highest cross-validated ROC-AUC (0.692 ± 0.062). This divergence in model rankings reflects sampling variability inherent to modest tabular datasets and specific train-test partition characteristics. Methodological evaluations of resampling techniques establish that model selection anchored to a single held-out split can be unstable, and that stratified cross-validation provides a more reliable basis for comparing classifiers [13], [14].

Furthermore, the stronger cross-validated stability of Extra Trees aligns with its underlying algorithmic design: by randomizing both feature subsets and split cut-points during node construction, Extra Trees substantially reduces prediction variance compared to deterministic or sequential boosting ensembles on noisy, small-sample datasets. Similarly, recent systematic reviews of machine learning in green roof and urban agricultural modelling confirm that tree-based ensembles such as Random Forest and Gradient Boosting consistently excel at capturing non-linear interactions across heterogeneous environmental and management inputs [15]. Extra Trees therefore merits prioritization as a baseline candidate due to its fold-to-fold consistency, while final algorithm selection should weigh validation stability against operational misclassification costs.

5.2 Gardening Experience as Leading Predictor

Gardening experience emerged as the most prominent predictor in both SHAP and impurity-based analyses. At the encoded level, the beginner category recorded the highest mean absolute SHAP contribution (0.0372), followed by the expert category (0.0161). This empirical prominence directly corroborates survey-based urban agriculture research showing that technical unfamiliarity with specialized rooftop practices, such as drainage management, lightweight substrate preparation, and localized pest control, constitutes the primary barrier for novice practitioners [16].

Practitioners with greater experience possess stronger experiential familiarity with cultivation routines and adaptive responses to rooftop microclimatic stresses, whereas beginners are disproportionately vulnerable to early crop failure [17], [18]. Importantly, however, the cross-sectional model does not establish whether experience causally drives higher yields; the observed importance reflects predictive association within the sample. Longitudinal cohort studies are needed to track how skill acquisition and targeted training interventions alter practitioner success trajectories over time.

5.3 Plant Health and Management Practices

Visual plant-health score (concurrent correlate) ranked second in impurity importance (0.0482) and contributed notably in SHAP analysis (encoded category: 0.0169), alongside circular cultivation practices such as kitchen-waste reuse and mulching. These patterns align with experimental rooftop agronomy literature demonstrating that organic waste valorization, compost integration, and moisture-retaining mulches significantly improve substrate water retention, nutrient cycling, and crop productivity under elevated rooftop evaporative demand [1], [19]–[20].

Systematic reviews of rooftop vegetable production similarly highlight circular economy practices and organic waste recycling as core pillars of resilient urban horticulture. An important methodological caveat concerns temporal ordering: because visual plant-health scores were recorded concurrently with the gardening outcome (Section 3.1), they reflect the outcome state rather than functioning as strictly prospective predictors. Future early-warning models should isolate baseline management practices from mid-cycle or post-hoc physiological indicators to prevent temporal leakage.

5.4 Resource Constraints

Cost-related constraint rank was the third most influential original predictor (0.0327). Urban rooftop cultivation entails recurring capital and operational expenses, including containers, lightweight growing media, organic fertilizers, and irrigation fixtures that frequently deter sustained household participation [5], [17], [18], [21]. Recent empirical assessments of rooftop gardening determinants confirm that perceived procurement costs for logistics and inputs rank immediately behind technical complexity as major obstacles to adoption and expansion.

Moreover, because rooftops offer valuable thermal co-benefits such as indoor cooling and urban heat island mitigation, alongside city-scale food-energy-water synergies, reducing financial entry barriers carries broader urban sustainability value. The salience of cost constraints in our model underscores the need for applied research into low-cost, locally sourced substrate alternatives, such as household compost and recycled aggregates [3], [16], while recognizing that the present observational model quantifies predictive association rather than the causal magnitude of financial subsidies [16].

5.5 Classification Behaviour and Practical Implications

Random Forest correctly identified 30 of 38 actual high-success observations (recall 78.9%), with 14 of 44 positive predictions representing false positives (precision 68.2%). In an urban gardening extension or municipal support system, such as community training or do-it-yourself kit distribution programs, [6], [7], [22], [23], prioritizing higher recall is advantageous when the goal is identifying and supporting the widest pool of potentially successful households.

Conversely, if high-success classifications trigger expensive capital subsidies or structural retrofitting grants, reducing false positives to improve precision becomes critical to prevent resource misallocation. Accordingly, operational deployments should calibrate the probability classification threshold to match institutional cost-benefit trade-offs rather than defaulting to a standard 0.50 cutoff.

5.6 Methodological Reflection on SHAP Interpretation

SHAP's local accuracy and additive feature-attribution framework provide a theoretically grounded decomposition of complex ensemble predictions. Nevertheless, standard tree-based and interventional SHAP implementations rely on assumptions of additivity and feature independence that are frequently violated in socio-agronomic datasets, where practitioner experience, financial capacity, and management habits naturally co-vary. When predictors are correlated, marginal feature perturbations can evaluate unrealistic combinations and distribute attribution weights unevenly across collinear variables [24]. Consequently, the SHAP values reported here are interpreted strictly as measures of predictive association within the fitted model rather than as causal effect estimates, serving as an exploratory screening tool to prioritize candidate variables for controlled agronomic or behavioral trials.

6. Limitations and Implications for Further Research

The limited sample size may render the test-set estimates sensitive to the random partition, and although five-fold cross-validation was performed, it was confined to the training partition; external validation on new observations therefore remains necessary. Furthermore, dichotomising the five-point overall success score at a threshold of 4 simplifies practitioner outcomes and warrants sensitivity testing across alternative cutoffs in future work. In addition, the timing of predictor measurement should be audited to exclude post-outcome

information from any intended prospective prediction model. A nested cross-validation or repeated stratified cross-validation design may afford more stable model selection and hyperparameter assessment, while external validation should ideally draw on a separately collected sample from different gardening locations, households, or growing seasons. Finally, the classification threshold should be selected according to the practical costs of false positives and false negatives rather than assumed to be 0.50.

7. Conclusion

This study demonstrates that supervised machine learning can extract meaningful predictive patterns of rooftop gardening success from practitioner-level questionnaire data. Using a dataset of 300 urban gardeners across Chennai, Madurai, and Coimbatore, four classification algorithms were compared to predict a binary high-success outcome. Gradient Boosting achieved the highest discrimination on the held-out test partition (ROC-AUC = 0.760), whereas Extra Trees demonstrated the most consistent and highest mean performance during five-fold stratified cross-validation (ROC-AUC = 0.692 ± 0.062). This divergence between test-set and cross-validated rankings underscores the critical importance of employing multiple validation strategies on modest-sized tabular datasets, cautioning against the premature selection of a single algorithm based on a solitary data split.

Feature attribution analyses, utilizing both SHAP and impurity-based importance, identified gardening experience and cost-related constraints as the most influential predictors of success, with visual plant-health score, a concurrent correlate, also contributing substantially to model output. The prominence of gardening experience and resource constraints highlights that successful rooftop cultivation is not solely a function of agronomic inputs, but is Associated with practitioner skill and economic limitations.

By shifting the analytical focus from aggregate spatial potential and controlled experimental efficacy to individual-level predictability, this study provides a foundational baseline for evidence-based urban agriculture interventions. The findings suggest that municipal and extension programs could move beyond generalized advocacy to targeted support, such as specialized training for novices and subsidies for affordable inputs. However, these predictive associations must not be misconstrued as causal effects. Future research should prioritize prospective, longitudinal data collection to establish temporal ordering, conduct external validation across diverse geographic and climatic contexts, and explore the deployment of these models within practical, cost-sensitive advisory systems for urban gardeners.

References

1. CM. M. Islam et al., "Synergistic effects of vermicompost and biochar amendments on cherry tomato production in an urban rooftop garden: A 2-year study," Urban Agric. Reg. Food Syst., vol. 11, no. 1, Art. no. e70045, 2026, doi: 10.1002/uar2.70045.
2. M. M. Islam et al., "Eco-efficient synergy of vermicompost and biochar in coir-based substrate enhances cherry tomato performance under urban rooftop garden conditions," Not. Bot. Horti Agrobot. Cluj-Napoca, vol. 53, no. 4, Art. no. 14878, Dec. 2025, doi: 10.15835/nbha53414878.
3. F. Carrier, M. Dorais, A.-C. Laliberté, and G. Grégoire, "Greening the sky: Incorporating recycled inorganic materials in rooftop growing media," Urban Agric. Reg. Food Syst., vol. 11, no. 1, Art. no. e70046, 2026, doi: 10.1002/uar2.70046.
4. M. L. Richardson and C. G. Arlotta, "Producing cherry tomatoes in urban agriculture," Horticulturae, vol. 8, no. 4, Art. no. 274, Mar. 2022, doi: 10.3390/horticulturae8040274.
5. A. Drottberger, Y. Zhang, J. W. H. Yong, and M.-C. Dubois, "Urban farming with rooftop greenhouses: A systematic literature review," Renew. Sustain. Energy Rev., vol. 188, Art. no. 113884, Dec. 2023, doi: 10.1016/j.rser.2023.113884.
6. V. S. Manivasagam, G. V. S. Shreyaa, and S. D. Ganesh, "Food-energy-water sustainability in urban spaces: Insights from open geospatial datasets in Coimbatore, India," Comput. Urban Sci., vol. 5, no. 1, Art. no. 54, Oct. 2025, doi: 10.1007/s43762-025-00213-w.
7. V. Singh, V. N. Singh, and S. K. Pandey, "Rooftop gardening for enhanced food security and nutrition: A case study from urban India," Int. J. Environ. Clim. Change, vol. 15, no. 12, pp. 868–880, Dec. 2025, doi: 10.9734/ijecc/2025/v15i125203.
8. A. Dubey, A. R. Radhakrishna, R. Narendra, N. Madhu, and R. Rajesh, "An engineering approach to roof top gardening," Int. J. Environ. Clim. Change, vol. 10, no. 5, pp. 14–23, May 2020, doi: 10.9734/ijecc/2020/v10i530197.
9. W. McKinney, "Data structures for statistical computing in Python," in Proc. 9th Python Sci. Conf. (SciPy), S. van der Walt and J. Millman, Eds., 2010, pp. 56–61, doi: 10.25080/Majora-92bf1922-00a.
10. C. R. Harris et al., "Array programming with NumPy," Nature, vol. 585, no. 7825, pp. 357–362, Sep. 2020, doi: 10.1038/s41586-020-2649-2.
11. F. Pedregosa et al., "Scikit-learn: Machine learning in Python," J. Mach. Learn. Res., vol. 12, pp. 2825–2830, Oct. 2011.
12. S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," in Adv. Neural Inf. Process. Syst., vol. 30, I. Guyon et al., Eds. Red Hook, NY, USA: Curran Associates, Inc., 2017, pp. 4765–4774.

13. R. Kohavi, "A study of cross-validation and bootstrap for accuracy estimation and model selection," in Proc. 14th Int. Joint Conf. Artif. Intell. (IJCAI), Montreal, QC, Canada, 1995, vol. 14, no. 2, pp. 1137–1145.
14. S. Varma and R. Simon, "Bias in error estimation when using cross-validation for model selection," BMC Bioinformatics, vol. 7, Art. no. 91, Feb. 2006, doi: 10.1186/1471-2105-7-91.
15. S. Seidu, D. W. M. Chan, and M. Santamouris, "Application of machine learning in multifunctional green roof assessments: A systematic review," Environ. Res. Commun., vol. 8, no. 9, Art. no. 092002, Sep. 2026, doi: 10.1088/2515-7620/aea67d.
16. M. N. A. S. Mithun, S. Rahman, M. A. Rahman, M. H. Sabit, J. L. Donaldson, and S. Sheheli, "Perceived barriers and determinants of rooftop gardening in urban agriculture," Urban Agric. Reg. Food Syst., vol. 11, no. 1, Art. no. e70054, May 2026, doi: 10.1002/uar2.70054.
17. Y. Harada and T. H. Whitlow, "Urban rooftop agriculture: Challenges to science and practice," Front. Sustain. Food Syst., vol. 4, Art. no. 76, Jun. 2020, doi: 10.3389/fsufs.2020.00076.
18. G. Picca, A. Goni-Urtiaga, L. Lozano de Sosa, M. V. Colombo Rodriguez, I. Fernandez Navarro, C. Plaza, and M. Panettieri, "Sustainability in urban agriculture: The role of biochar in enhancing the productive capacity of compost-based growing media for rooftop farming," Urban For. Urban Green., vol. 107, Art. no. 128774, May 2025, doi: 10.1016/j.ufug.2025.128774.
19. B. J.-P. Grard, C. Chenu, N. Manouchehri, S. Houot, N. Frascaria-Lacoste, and C. Aubry, "Rooftop farming on urban waste provides many ecosystem services," Agron. Sustain. Dev., vol. 38, no. 1, Art. no. 2, Feb. 2018, doi: 10.1007/s13593-017-0474-2.
20. M. Oueld Lhaj, R. Moussadek, L. Mouhir, H. Sanad, A. Zouahri, R. Spaccini, and R. Paradiso, "Compost as a sustainable tool for greenhouse horticultural productivity: Benefits, limitations, and state of the art," Front. Sustain. Food Syst., vol. 10, Art. no. 1833979, Jun. 2026, doi: 10.3389/fsufs.2026.1833979.
21. G. Visvanathan, K. Patil, Y. Suryawanshi, V. Meshram, and S. Jadhav, "Mitigating urban heat island and enhancing indoor thermal comfort using terrace garden," Sci. Rep., vol. 14, no. 1, Art. no. 9697, 2024, doi: 10.1038/s41598-024-60546-0.
22. L. Ieviņa, P. Lakovskis, S. Lakovska, S. Zeipiņa, and L. Lepse, "Way up for sustainable food: A review on challenges and opportunities in rooftop vegetable production," J. Hum. Earth Future, vol. 7, no. 3, pp. 966–984, Sep. 2026, doi: 10.28991/HEF-2026-07-03-018.
23. Department of Horticulture and Plantation Crops, Government of Tamil Nadu, Operational Manual: Do-It-Yourself Kit Program for Promotion of Rooftop Vegetable Gardening in Chennai and Coimbatore. Coimbatore, India: Tamil Nadu Agricultural University (TNAU) Agritech Portal, 2014. [Online]. Available: https://agritech.tnau.ac.in/horticulture/horti_schemes_doit_yourself.html
24. K. Aas, M. Jullum, and A. Løland, "Explaining individual predictions when features are dependent: More accurate approximations to Shapley values," Artif. Intell., vol. 298, Art. no. 103502, Sep. 2021, doi: 10.1016/j.artint.2021.103502.