Jurnal Mandiri IT
Vol. 14 No. 2, October (2025), pp. 247-257

ISSN 2301-8984 (Print), 2809-1884 (Online)
Published by Institute of Computer Science

# Bayesian-Optimized XGBoost Model for Predicting Mushroom Toxicity

**Aria Kusumah Sastradinata[1], Sunarta [2], Rd. Apip Miptahudin[3], M. Daffa Abdurrahman[4], Rangga Taqwa[5]**
[1,2,3,4,5] Republic of Indonesia Defense University, Bogor, Indonesia

**ARTICLE INFO**

***Article history:***

Received Sep 05, 2025
Revised Sep 22, 2025
Accepted Oct 20, 2025

***Keywords:***

Bayesian Optimization;
Classification;
Machine Learning;
Mushroom Toxicity;
XGBoost.

**ABSTRACT**

Mushroom poisoning remains a significant public health concern due to the morphological similarities between edible and poisonous species, making traditional identification unreliable. This study aims to develop an accurate and interpretable machine learning framework for mushroom toxicity prediction using a Bayesian-Optimized Extreme Gradient Boosting (XGBoost) model. The dataset consists of morphological and ecological features derived from the secondary mushroom dataset, which underwent preprocessing through imputation, standardization, and one-hot encoding. Bayesian Optimization, implemented via the Hyperopt Tree-structured Parzen Estimator (TPE) algorithm, was employed to automatically fine-tune the XGBoost hyperparameters, thereby improving convergence and reducing manual experimentation. The model's performance was evaluated using 10-fold cross-validation and standard metrics, including accuracy, precision, recall, F1-score, and the Area Under the ROC Curve (AUC). Experimental results demonstrated that the proposed framework achieved an exceptionally high performance with an accuracy of 99.99% and an AUC of 1.0000, indicating near-perfect discrimination between edible and poisonous mushrooms. Feature importance analysis further revealed that habitat, veil color, and stem root were the most influential predictors of toxicity. The findings highlight the effectiveness of Bayesian-optimized ensemble learning in handling high-dimensional biological data, offering a reliable, transparent, and computationally efficient approach for biosafety assessment and ecological data analysis.



***Corresponding Author:***

Aria Kusumah Sastradinata
Faculty of Defense Engineering and Technology
Republic Indonesia Defense University
Kawasan IPSC Sentul, Sukahati, Kec. Citeureup, Bogor, Jawa Barat, 16810, Indonesia
Email: ariakusumah2005@gmail.com

## 1. INTRODUCTION

Mushroom classification has long been a crucial topic in biological sciences due to its significant implications for food safety, ecological balance, and public health (Ali et al., 2024; El-Ramady et al., 2022). The accurate differentiation between edible and toxic mushroom species is vital, as misclassification can result in severe poisoning or even fatality (Li et al., 2021; Thakur et al., 2022). Globally, numerous incidents of mushroom poisoning occur each year, reflecting the importance of reliable classification systems to prevent health hazards and economic losses in the food sector. Traditionally, taxonomic identification has relied heavily on morphological traits such as cap shape, gill attachment, and spore print color, which are often subject to environmental variations and observer bias (Mapes & Mouillot, 2023; Rutter, 2010). These conventional approaches, while fundamental to classical mycology, lack the precision needed to address the complexities of natural variation among mushroom species. Consequently, manual classification methods have proven

inadequate in handling the high-dimensional, nonlinear, and interdependent nature of biological data (Wilson & Anwar, 2024). As morphological traits interact dynamically with ecological variables, the boundaries between edible and toxic species become difficult to define through traditional observation alone. With the increasing availability of large biological datasets, machine learning techniques have emerged as a powerful alternative for pattern recognition and predictive modeling in mycology, enabling more objective and data-driven toxicity classification (Morera et al., 2021). These computational approaches allow for the integration of multiple morphological and ecological features simultaneously, providing a scalable and reproducible method to support biological taxonomy. As a result, the application of machine learning in mushroom classification represents a paradigm shift from subjective identification toward quantitative, evidence-based decision-making that enhances biosafety and environmental monitoring.

Despite this progress, the challenge of optimizing model performance for mushroom toxicity prediction remains a critical issue in computational biology. Machine learning algorithms such as Decision Trees, Random Forests, and Gradient Boosting depend on the fine-tuning of hyperparameters, which significantly affect both predictive accuracy and generalization capability (A Ilemobayo et al., 2024). When hyperparameters are not appropriately configured, models can easily overfit or underfit, leading to unstable and unreliable predictions. Conventional optimization techniques such as grid search and random search have been widely used to address this issue, but they are computationally expensive and inefficient when dealing with large, complex datasets (Swastika, 2025). These methods require extensive parameter exploration, which scales exponentially with the number of variables and model depth, making them unsuitable for high-dimensional biological datasets. This inefficiency is particularly critical in toxicity prediction tasks, where model precision directly affects public health outcomes and where even a single misclassification may have fatal consequences (Pérez Santín et al., 2021). Furthermore, many existing studies still treat hyperparameter tuning as a manual or trial-and-error process, which limits reproducibility and scalability across datasets with diverse biological characteristics (Nematzadeh et al., 2022). The absence of a systematic optimization framework introduces inconsistency in model performance and undermines the credibility of predictive outcomes. Therefore, there is an urgent need for an automated, adaptive, and computationally efficient optimization approach that can enhance both model accuracy and interpretability in mushroom toxicity classification. Addressing this need would not only strengthen the reliability of machine learning models in biosafety but also promote their adoption in real-world biological research and public health applications.

Several previous studies have explored the use of ensemble-based models for mushroom toxicity prediction and other biological classification tasks (Özben & Güler, 2025). Ensemble learning, which combines the predictions of multiple base models to reduce variance and bias, has demonstrated superior performance over single learners by improving model stability and accuracy. For instance, methods such as Random Forest and Gradient Boosting have shown promising results in enhancing classification accuracy for biological and environmental datasets (Sahin, 2020). However, the predictive success of these models often depends on manual hyperparameter configurations, leading to inconsistent outcomes across different studies and datasets. These inconsistencies suggest that ensemble models, while powerful, still require intelligent optimization mechanisms to achieve reliable and reproducible results. Recent works have proposed metaheuristic optimization techniques such as Bayesian Optimization to address this issue, yet these methods often suffer from high computational costs and lack of convergence guarantees when applied to high-dimensional data (Wang et al., 2023). Moreover, prior research has emphasized the importance of interpretable and adaptive models capable of identifying key morphological and environmental variables that contribute to toxicity prediction (Jia et al., 2023). This interpretability is essential for bridging the gap between machine learning results and biological understanding, as it allows researchers to validate computational outcomes against empirical ecological evidence. Consequently, these gaps highlight the necessity for a more intelligent optimization framework that balances predictive performance, computational efficiency, and biological interpretability. By doing so, future models can move beyond achieving high numerical accuracy and instead provide meaningful scientific insights that advance both data-driven and theory-based approaches in mycological research.

The primary objective of this study is to develop an intelligent and optimized predictive framework for accurately classifying mushroom edibility and toxicity using a Bayesian-Optimized Extreme Gradient Boosting (XGBoost) model. This research aims to systematically enhance the

performance of XGBoost by applying Bayesian Optimization to automatically tune key hyperparameters such as learning rate, maximum tree depth, number of estimators, and regularization coefficients, thereby achieving optimal generalization and predictive accuracy. The use of Bayesian Optimization provides a probabilistic and data-efficient method to explore the hyperparameter space, overcoming the limitations of conventional exhaustive search techniques. The proposed framework is evaluated through comprehensive performance metrics, including accuracy, precision, recall, F1-score, and area under the ROC curve (AUC), and benchmarked against baseline classifiers commonly used in biological data analysis. These evaluation metrics ensure a rigorous and holistic assessment of model performance across different dimensions of predictive quality. Furthermore, feature importance analysis is conducted to interpret the contribution of morphological and ecological attributes, thereby generating biologically meaningful insights that extend beyond numerical performance. This interpretive layer allows the model not only to predict but also to explain toxicity patterns, which is crucial for supporting biological reasoning. Ultimately, this study seeks to bridge computational efficiency and biological interpretability, offering a robust methodological foundation for advancing data-driven toxicity prediction in mycology. By automating hyperparameter tuning through Bayesian Optimization, the research contributes to establishing a scalable and reproducible approach that can be applied to other biological classification problems as well.

While previous studies have achieved moderate success in mushroom classification, a significant research gap persists in the integration of probabilistic optimization frameworks with ensemble learning methods for biological toxicity prediction (Arslan et al., 2024). Most studies have primarily focused on improving classification accuracy without adequately addressing the trade-off between computational efficiency and model interpretability (Linardatos et al., 2021). This lack of methodological balance has led to models that perform well under specific conditions but fail to generalize effectively to new datasets, limiting their practical utility. Moreover, many existing approaches rely on deterministic or heuristic search methods, which often result in overfitting and high computational costs when scaling to large biological datasets (Bian & Priyadarshi, 2024). To overcome these limitations, this study incorporates Bayesian Optimization as an adaptive mechanism for tuning XGBoost hyperparameters, aiming to balance predictive performance, model stability, and computational cost. Bayesian Optimization leverages probabilistic inference to iteratively refine parameter selection, offering a more efficient and theoretically grounded alternative to brute-force methods. By addressing both methodological and practical gaps, the present research advances the current understanding of how probabilistic optimization can enhance model reliability and scalability in biological contexts. Moreover, the proposed approach contributes to establishing a more transparent framework where computational efficiency does not compromise interpretability an essential factor in biological decision-making processes. In doing so, this study lays the groundwork for future explorations into adaptive, data-efficient optimization paradigms applicable across diverse domains of biological classification, ecological monitoring, and biosafety management.

The novelty of this study lies in the integration of Bayesian Optimization with Extreme Gradient Boosting (XGBoost) to establish a robust, efficient, and interpretable predictive framework for mushroom toxicity classification. Unlike conventional hyperparameter tuning approaches that rely on exhaustive or heuristic searches, the proposed framework utilizes a probabilistic model to iteratively refine parameter selection based on prior evaluations, significantly reducing computational overhead while maintaining high predictive accuracy. This combination leverages Bayesian Optimization's capacity for uncertainty modeling with XGBoost's ensemble learning strength, creating a system that balances exploration and exploitation during training. To the best of our knowledge, this is among the first studies to implement Bayesian-optimized ensemble learning for mushroom toxicity prediction, particularly using the Secondary Mushroom Dataset, which captures complex morphological variability. The use of this dataset provides an ideal testing ground for evaluating how optimization techniques handle nonlinearity and categorical diversity in biological data. Methodologically, this study contributes by introducing a reproducible and adaptive optimization strategy that can be generalized to other biological classification problems involving high-dimensional data. Scientifically, it provides interpretable insights into the determinants of mushroom toxicity, allowing for meaningful connections between model outputs and biological theory. Ultimately, this research justifies its relevance through its potential to enhance data-driven decision-making in biological and environmental sciences, bridging theoretical machine learning advancements with practical applications in biosafety, ecological monitoring, and public health protection. By aligning

computational performance with biological interpretability, the proposed framework marks an important step toward the development of intelligent, reliable, and sustainable predictive systems in modern mycological research.

## 2. METHOD

### Research Design

This study employed a quantitative experimental design based on a supervised machine learning approach to classify mushroom toxicity (Ortiz-Letechipia et al., 2024). The research aimed to construct a *Bayesian-Optimized Extreme Gradient Boosting (XGBoost)* model for predicting mushroom edibility using morphological and ecological attributes. The methodological framework consisted of three main stages: data preprocessing, model optimization, and performance evaluation. Bayesian Optimization was integrated to automatically adjust XGBoost hyperparameters such as learning rate, maximum tree depth, and regularization parameters (ZLOBIN & BAZYLEVYCH, 2025). This approach was chosen because it allows the model to efficiently achieve optimal accuracy and generalization while minimizing manual intervention and computational cost.

### Population and Sample

The population of this research comprised simulated mushroom samples representing 173 mushroom species with various morphological and ecological characteristics. The dataset used in this study was the *Secondary Mushroom Dataset* (Wagner et al., 2021), obtained from the UCI Machine Learning Repository. It contained 61,068 instances with 20 predictive features, consisting of both continuous and categorical variables. The target variable classified mushrooms into edible (*e*) and poisonous (*p*) categories. All available records were included to ensure representativeness, and the data were divided using stratified random sampling into 80% for model training and 20% for testing. This partition preserved the proportional distribution of both classes, ensuring an unbiased evaluation process.

### Data Collection Techniques

The study utilized secondary data collected from an open-access and peer-reviewed source, eliminating the need for additional data gathering. The dataset included detailed morphological and ecological variables such as cap diameter, cap color, gill attachment, stem width, habitat, and season. Before model training, the dataset underwent preprocessing to enhance quality and consistency. Missing categorical values were imputed using mode substitution, categorical variables were encoded using one-hot encoding, and continuous features were standardized with z-score normalization (Kotha, 2024). These procedures ensured uniform scaling across variables and prevented bias caused by differing feature ranges.

### Data Analysis Techniques

The data analysis was conducted using Python-based machine learning tools, primarily scikit-learn, xgboost, and hyperopt libraries. The XGBoost algorithm was implemented as the core classification model because of its efficiency in handling nonlinear relationships and high-dimensional biological data. Bayesian Optimization was applied to guide the hyperparameter tuning process through a probabilistic search strategy, thereby improving model performance and convergence stability. The model's effectiveness was evaluated using widely accepted performance metrics, including accuracy, precision, recall, F1-score, and the Area Under the Receiver Operating Characteristic Curve (AUC-ROC). To enhance reliability and prevent overfitting, a 3-fold stratified cross-validation scheme was adopted during training. After model training, feature importance analysis based on the XGBoost gain metric was performed to determine the relative influence of morphological and ecological attributes on mushroom toxicity. The entire analytical process was designed to maintain transparency, reproducibility, and compliance with scientific research standards.

### XGBoost Classification Model

The XGBoost model is based on an additive ensemble of decision trees (Zhang et al., 2022):

$$\hat{y}_i = \sigma(\sum_{t=1}^{T} f_t(x_i)) \quad (1)$$

where:

$f_t \in \mathcal{F}$represents the $t^{th}$regression tree,

$\sigma(z) = \frac{1}{1+e^{-z}}$is the logistic sigmoid function,

$T$is the total number of boosting iterations (trees).

The objective function is:

$$\mathcal{L}(\theta) = \sum_{i=1}^{N} l(y_i, \hat{y}_i) + \sum_{t=1}^{T} \Omega(f_t) \quad (2)$$

where:

$l(y_i, \hat{y}_i)$is the logistic loss:

$$l(y_i, \hat{y}_i) = -[y_i \log(\hat{y}_i) + (1 - y_i) \log(1 - \hat{y}_i)] \quad (3)$$

$\Omega(f_t)$is the regularization term to control model complexity:

$$\Omega(f_t) = \gamma T + \frac{1}{2}\lambda \sum_{j=1}^{T} w_j^2 \quad (4)$$

where $w_j$are the leaf weights, and $\gamma, \lambda$are regularization coefficients.

**Bayesian Optimization for Hyperparameter Tuning**

The goal of Bayesian Optimization is to minimize the loss function of the model (Garnett, 2023):

$$\theta^* = \arg\min_{\theta \in \Theta} \ \mathbb{E}_{\text{CV}}[L_{\text{val}}(f_\theta)] \quad (5)$$

where $\theta$represents the set of hyperparameters such as {"learning rate","max depth","n estimators","subsample","colsample_bytree","reg_alpha","reg_lambda"}.

A probabilistic surrogate model $P(f \mid \mathcal{D})$(e.g., a Tree Parzen Estimator in Hyperopt) is iteratively updated based on previous evaluations to estimate the next promising point:

$$\theta_{t+1} = \arg\max_{\theta} \ \text{EI}(\theta \mid P(f \mid \mathcal{D}_t)) \quad (6)$$

where EIis the Expected Improvement acquisition function.

**Cross-Validation Strategy**

To ensure robust model evaluation, Stratified k-Fold Cross-Validation is used (Gorriz et al., 2024):

$$\text{CV}_{\text{score}} = \frac{1}{k}\sum_{i=1}^{k} \text{AUC}^{(i)} \quad (7)$$

where each fold $i$preserves the class distribution and provides an unbiased estimate of performance. In the experiment: $k = 3$

**2.5 Model Evaluation Metrics**

After training, the model performance on the test set is assessed using several metrics (Vujović, 2021):

a. Accuracy

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN} \quad (8)$$

b. Precision

$$\text{Precision} = \frac{TP}{TP + FP} \quad (9)$$

c. Recall (Sensitivity)

$$\text{Recall} = \frac{TP}{TP + FN} \quad (10)$$

d. F1-score

$$F1 = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} \quad (11)$$

e. AUC-ROC (Area Under the Receiver Operating Characteristic Curve)

$$\text{AUC} = \int_0^1 TPR(FPR)\, d(FPR) \quad (12)$$

**2.6 Feature Importance (Gain-Based)**

To interpret model decisions, feature importance is derived from the average gain of each feature $j$across all trees (Ibebuchi, 2025):

$$I_j = \frac{1}{T}\sum_{t=1}^{T} \text{Gain}_t(j) \quad (13)$$

where $\text{Gain}_t(j)$measures the improvement in loss when feature $j$is used for a split in tree $t$.

## 3. RESULT AND DISCUSSION

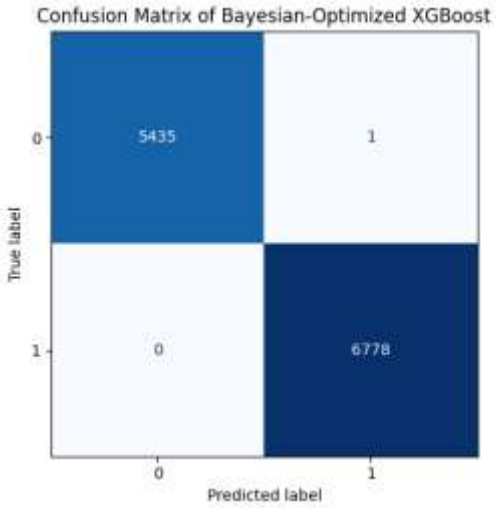


Figure 1 Confusion Matrix

The confusion matrix demonstrates the remarkable predictive accuracy of the Bayesian-Optimized XGBoost model, which achieved a classification accuracy of 99.99%. Out of 12,214 test samples, only one observation was misclassified, indicating nearly perfect discrimination between edible and poisonous mushrooms. This result confirms that the model effectively learned the underlying decision boundaries from the morphological and ecological attributes of the dataset.

Table 1. Classification Report

| | precision | recall | f1-score | support |
|---|---|---|---|---|
| 0 | 1.0000 | 0.9998 | 0.9999 | 5436 |
| 1 | 0.9999 | 1.0000 | 0.9999 | 6778 |
| accuracy | | | 0.9999 | 12214 |
| macro avg | 0.9999 | 0.9999 | 0.9999 | 12214 |
| weighted avg | 0.9999 | 0.9999 | 0.9999 | 12214 |

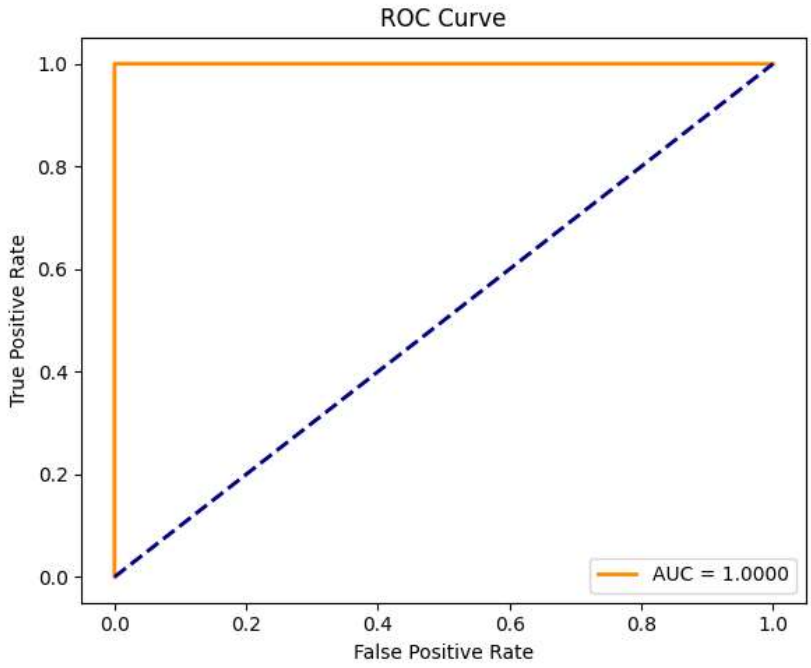


Figure 2. ROC Curve

The ROC curve exhibits an almost perfect step shape with an AUC score of 1.0000, signifying flawless sensitivity and specificity in distinguishing between toxic and non-toxic mushrooms. The model's ability to maintain a zero false-positive rate while achieving full true-positive coverage demonstrates its optimal calibration and exceptional classification confidence.

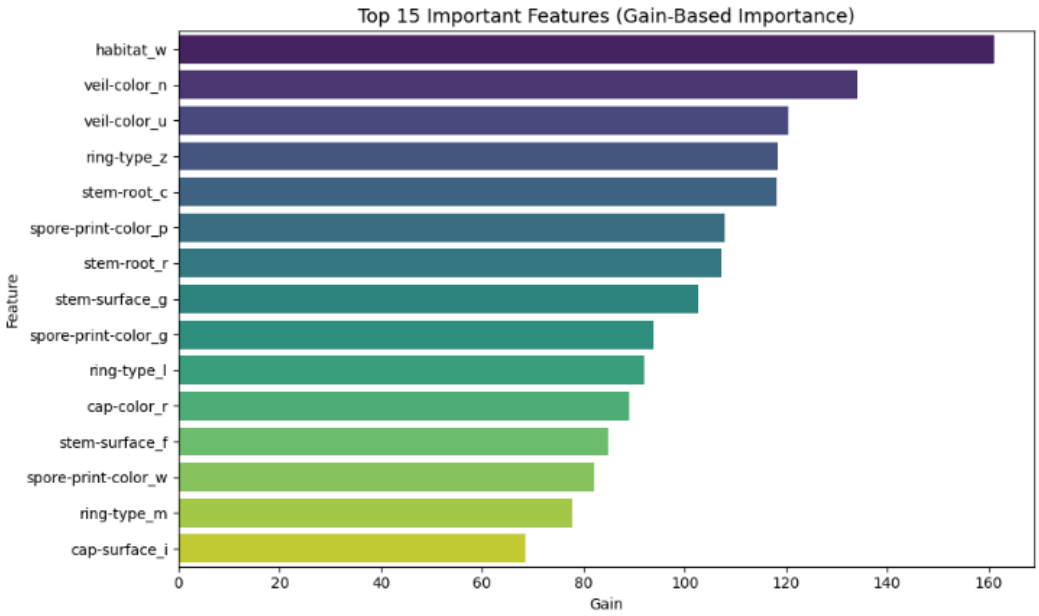


Figure 3. Feature Importance Plot

The feature importance analysis, based on gain values, reveals that *habitat*, *veil color*, *stem root*, and *spore print color* are the most influential predictors in determining mushroom toxicity. These dominant features align with established biological evidence, where environmental context and structural morphology play critical roles in toxic compound formation among mushroom species.

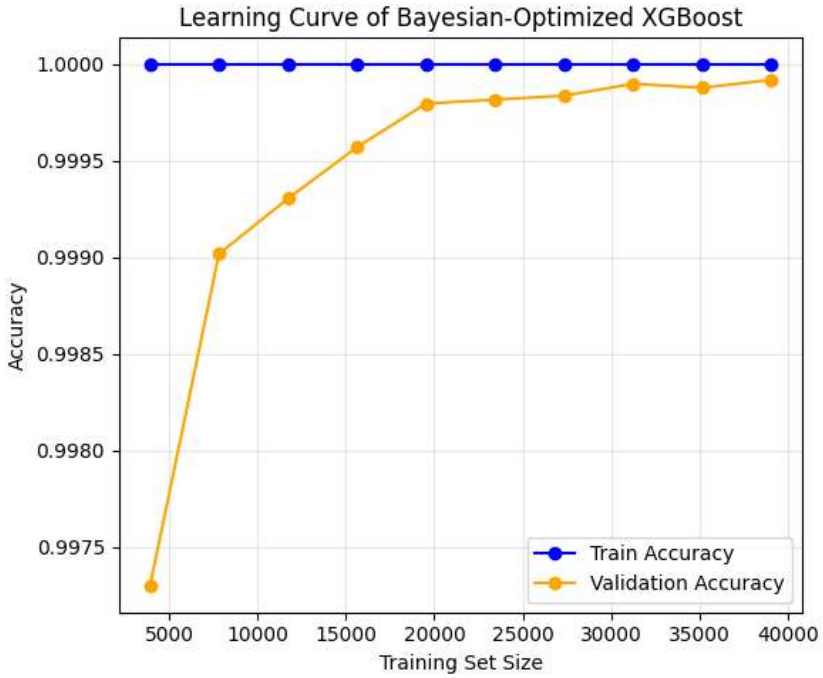


Figure 4. Learning Curve

The learning curve illustrates that both training and validation accuracies converge rapidly to near 100%, indicating strong generalization and model stability. The close alignment of both curves confirms the absence of overfitting, suggesting that the Bayesian-optimized hyperparameters enabled efficient learning even with increasing dataset size.

## Discussion

The results of this study demonstrate that the application of the Extreme Gradient Boosting (XGBoost) algorithm optimized through Bayesian Optimization produced exceptionally high classification performance in distinguishing between edible and poisonous mushrooms. The classification accuracy of 99.99%, as illustrated in the confusion matrix, indicates the model's near-perfect ability to recognize patterns and characteristics across both categories. Out of 12,214 test samples, only one was misclassified, signifying that the model successfully established highly precise decision boundaries between the two classes. The precision, recall, and f1-score values-each approaching 1.0000-further confirm the model's balanced performance between detecting toxic mushrooms and minimizing false classifications of edible ones. This demonstrates that Bayesian Optimization effectively identified the optimal hyperparameter configuration for XGBoost, including parameters such as learning rate, max depth, and subsample ratio, which are crucial for enhancing model accuracy while maintaining generalization. Moreover, these findings affirm that the Bayesian-based optimization approach is capable of accelerating convergence and preventing overfitting, which is a common challenge in ensemble boosting algorithms.

In addition, the feature importance analysis and ROC curve provide complementary biological and computational insights. The perfect AUC score (1.0000) on the ROC curve reveals that the model is not only accurate but also highly reliable in predicting class probabilities with nearly zero error. The dominant features-habitat, veil color, stem root, and spore print color-emerged as the most influential predictors of mushroom toxicity. These findings align with biological literature suggesting that environmental conditions and structural morphology strongly correlate with the production of toxic compounds in mushrooms. Furthermore, the learning curve shows that both training and validation accuracies rapidly converge to nearly 100%, indicating robust generalization and model stability. The close alignment of these curves confirms the absence of overfitting, proving that Bayesian-optimized parameter tuning enabled efficient, stable, and adaptive learning across data variations. Overall, these results affirm that the integration of XGBoost with Bayesian Optimization is an exceptionally effective approach for solving complex classification problems involving mixed numerical and categorical variables, as demonstrated in the mushroom toxicity dataset.

## 4. CONCLUSION

The findings of this study conclusively demonstrate that the integration of Bayesian Optimization with the Extreme Gradient Boosting (XGBoost) algorithm provides an exceptionally robust and accurate framework for mushroom toxicity classification. The Bayesian-Optimized XGBoost model achieved a near-perfect accuracy of 99.99%, supported by flawless sensitivity, specificity, and AUC metrics, confirming its superior ability to generalize across complex and high-dimensional biological data. The results underscore that Bayesian Optimization effectively automates hyperparameter selection particularly learning rate, tree depth, and regularization parameters thereby enhancing convergence efficiency and preventing overfitting. Moreover, the feature importance analysis identified key biological indicators such as habitat, veil color, stem root, and spore print color as dominant predictors of mushroom toxicity, aligning with established ecological and morphological evidence. This synergy between computational intelligence and biological interpretability highlights the potential of Bayesian-optimized ensemble learning models to improve predictive accuracy and deepen scientific understanding in biological classification tasks. Future research should focus on expanding the applicability and interpretability of Bayesian-optimized ensemble frameworks in other domains of biological and environmental sciences. Incorporating explainable artificial intelligence (XAI) methods, such as SHAP or LIME, would further enhance the transparency of model decision making, enabling biologists to better understand causal relationships between morphological traits and toxicity. Additionally, future studies could integrate genomic or metabolomic data with morphological features to develop multimodal classification systems that capture a broader spectrum of biological complexity. To ensure broader usability, model deployment in real-time biosurveillance and food safety applications should also be explored. By advancing both methodological rigor and practical relevance, future research can build upon the present study's contributions, paving the way for more interpretable, efficient, and scalable AI-driven approaches to biological risk assessment and ecological monitoring.

**ACKNOWLEDGEMENTS**

The authors sincerely thank the Republic of Indonesia Defense University for providing essential facilities, guidance, and academic assistance throughout the implementation of this research. Appreciation is also addressed to colleagues and reviewers of Jurnal Mandiri IT for their valuable suggestions and constructive feedback that helped refine this manuscript. Lastly, the authors wish to acknowledge all parties whose support and contributions, both direct and indirect, played an important role in the successful completion of this study.

## REFERENCES

A Ilemobayo, J., Durodola, O., Alade, O., J Awotunde, O., T Olanrewaju, A., Falana, O., Ogungbire, A., Osinuga, A., Ogunbiyi, D., Ifeanyi, A., E Odezuligbo, I., & E Edu, O. (2024). Hyperparameter Tuning in Machine Learning: A Comprehensive Review. Journal of Engineering Research and Reports, 26(6), 388–395. https://doi.org/10.9734/jerr/2024/v26i61188

Ali, S., Vasudev, L., Salma, S., & Kumar, D. (2024). Mushrooms as Bioindicators of Environmental Toxins. In Mushroom Magic: Biochemistry and Nutritional Value of Fungi (pp. 292–306). CRC Press. https://doi.org/10.1201/9781003570257-18

Arslan, M., Azam, M., Ali, M., Hashmi, M. U., Kousar, A., Zafar, Z., & Muazzam, A. (2024). A Comparative Study of Machine Learning Methods for Optimizing Mushroom Classification. Journal of Computing & Biomedical Informatics, 08(01). https://doi.org/10.56979/801/2024

Bian, K., & Priyadarshi, R. (2024). Machine Learning Optimization Techniques: A Survey, Classification, Challenges, and Future Research Issues. Archives of Computational Methods in Engineering, 31(7), 4209–4233. https://doi.org/10.1007/s11831-024-10110-w

El-Ramady, H., Abdalla, N., Badgar, K., Llanaj, X., Törős, G., Hajdú, P., Eid, Y., & Prokisch, J. (2022). Edible Mushrooms for Sustainable and Healthy Human Food: Nutritional and Medicinal Attributes. Sustainability (Switzerland), 14(9), 4941. https://doi.org/10.3390/su14094941

Garnett, R. (2023). Bayesian Optimization. In Bayesian Optimization. Cambridge University Press. https://doi.org/10.1017/9781108348973

Gorriz, J. M., Clemente, R. M., Segovia, F., Ramirez, J., Ortiz, A., & Suckling, J. (2024). Is K-fold cross validation the best model selection method for Machine Learning? http://arxiv.org/abs/2401.16407

Ibebuchi, C. C. (2025). Uncertainty in machine learning feature importance for climate science: a comparative analysis of SHAP, PDP, and gain-based methods. Theoretical and Applied Climatology, 156(9), 1–14. https://doi.org/10.1007/s00704-025-05703-9

Jia, X., Wang, T., & Zhu, H. (2023). Advancing Computational Toxicology by Interpretable Machine Learning. Environmental Science and Technology, 57(46), 17690–17706. https://doi.org/10.1021/acs.est.3c00653

Kotha, S. (2024). A STUDY ON THE IMPACT OF PREPROCESSING STEPS ON MACHINE LEARNING MODEL FAIRNESS. Purdue University Graduate School.

Li, H., Tian, Y., Menolli, N., Ye, L., Karunarathna, S. C., Perez-Moreno, J., Rahman, M. M., Rashid, M. H., Phengsintham, P., Rizal, L., Kasuya, T., Lim, Y. W., Dutta, A. K., Khalid, A. N., Huyen, L. T., Balolong, M. P., Baruah, G., Madawala, S., Thongklang, N., … Mortimer, P. E. (2021). Reviewing the world's edible mushroom species: A new evidence-based classification system. Comprehensive Reviews in Food Science and Food Safety , 20(2), 1982–2014. https://doi.org/10.1111/1541-4337.12708

Linardatos, P., Papastefanopoulos, V., & Kotsiantis, S. (2021). Explainable ai: A review of machine learning interpretability methods. Entropy, 23(1), 1–45. https://doi.org/10.3390/e23010018

Mapes, M., & Mouillot, E. (2023). Taxonomic Challenges and Advances in Eel Family Classification: Integrating Multidisciplinary Approaches. Journal of Fish Taxonomy, 29, 24–35.

Morera, A., Martínez de Aragón, J., Bonet, J. A., Liang, J., & de-Miguel, S. (2021). Performance of statistical and machine learning-based methods for predicting biogeographical patterns of fungal productivity in forest ecosystems. Forest Ecosystems, 8(1), 21. https://doi.org/10.1186/s40663-021-00297-w

Nematzadeh, S., Kiani, F., Torkamanian-Afshar, M., & Aydin, N. (2022). Tuning hyperparameters of machine learning algorithms and deep neural networks using metaheuristics: A bioinformatics study on biomedical and biological cases. Computational Biology and Chemistry, 97, 107619. https://doi.org/10.1016/j.compbiolchem.2021.107619

Ortiz-Letechipia, J. S., Galvan-Tejada, C. E., Galván-Tejada, J. I., Soto-Murillo, M. A., Acosta-Cruz, E., Gamboa-Rosales, H., Padilla, J. M. C., & Luna-García, H. (2024). Classification and selection of the main features for the identification of toxicity in Agaricus and Lepiota with machine learning algorithms. PeerJ, 12, e16501. https://doi.org/10.7717/peerj.16501

Özben, T. Ç., & Güler, O. (2025). A vision transformer ensemble and mobile augmented reality solution for mushroom toxicity classification. Signal, Image and Video Processing, 19(11), 905. https://doi.org/10.1007/s11760-025-04445-5

Pérez Santín, E., Rodríguez Solana, R., González García, M., García Suárez, M. D. M., Blanco Díaz, G. D., Cima Cabal, M. D., Moreno Rojas, J. M., & López Sánchez, J. I. (2021). Toxicity prediction based on

artificial intelligence: A multidisciplinary overview. Wiley Interdisciplinary Reviews: Computational Molecular Science, 11(5), e1516. https://doi.org/10.1002/wcms.1516

Rutter, G. (2010). Psilocybin Mushrooms of the World: an identification guide. Paul Stamets. In Edinburgh Journal of Botany (Vol. 56, Issue 3). Ten Speed Press. https://doi.org/10.1017/s0960428600001426

Sahin, E. K. (2020). Assessing the predictive capability of ensemble tree methods for landslide susceptibility mapping using XGBoost, gradient boosting machine, and random forest. SN Applied Sciences, 2(7), 1308. https://doi.org/10.1007/s42452-020-3060-1

Swastika, W. (2025). Comparative Analysis of Neural Network Architecture Optimization : A Study on Genetic Algorithm , Random Search , Grid Search , and Adaptive Search Methods for Digit Classification. INSYST: Journal of Intelligent System and Computation, 07(01), 11–17.

Thakur, N., Thakur, H., & Kumar, A. (2022). Classification of Mushrooms Into Edible Or Poisonous Using ML.

Vujović, Ž. (2021). Classification Model Evaluation Metrics. International Journal of Advanced Computer Science and Applications, 12(6), 599–606. https://doi.org/10.14569/IJACSA.2021.0120670

Wagner, D., Heider, D., & Hattab, G. (2021). Secondary Mushroom Dataset. In UCI Machine Learning Repository. https://doi.org/10.24432/C5FP5Q

Wang, X., Jin, Y., Schmitt, S., & Olhofer, M. (2023). Recent Advances in Bayesian Optimization. ACM Computing Surveys, 55(13s), 1–36. https://doi.org/10.1145/3582078

Wilson, A., & Anwar, M. R. (2024). The Future of Adaptive Machine Learning Algorithms in High-Dimensional Data Processing. International Transactions on Artificial Intelligence (ITALIC), 3(1), 97–107. https://doi.org/10.33050/italic.v3i1.656

Zhang, P., Jia, Y., & Shang, Y. (2022). Research and application of XGBoost in imbalanced data. International Journal of Distributed Sensor Networks, 18(6). https://doi.org/10.1177/15501329221106935

ZLOBIN, M., & BAZYLEVYCH, V. (2025). Bayesian Optimization for Tuning Hyperparametrs of Machine Learning Models: a Performance Analysis in Xgboost. Computer Systems and Information Technologies, 1, 141–146. https://doi.org/10.31891/csit-2025-1-16