International Journal For Research in Applied Science and Engineering Technology

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 **Issue:** V **Month of publication:** May 2026

DOI: https://doi.org/10.22214/ijraset.2026.82543

www.ijraset.com

Call: 08813907089 | E-mail ID: ijraset@gmail.com

Edible Mushroom Identification and Recommendary System Using Machine Learning

Kanika Lal[1], Mr. Gowtham R Naik[2]
[1]Dept. Of Computer Science and Engineering, The National Institute of Engineering, Mysuru, India
[2]Assistant Professor Dept. of Computer science and Engineering, The National Institute of Engineering, Mysuru, India

***Abstract:** Identifying edible and poisonous mushrooms is challenging due to their close visual similarity and variations in natural environments. Misclassification can lead to severe health risks, especially in communities that rely on visual inspection without expert support. This work presents an automated mushroom edibility classification system built using deep learning. Two backbone architectures—EfficientNet-B3 and ResNet-50—are evaluated to determine their effectiveness in learning fine-grained mushroom features. The dataset includes images collected under diverse environmental conditions and is enhanced through preprocessing and augmentation to improve generalization. Experimental analysis shows that EfficientNet-B3 provides superior accuracy compared to ResNet-50. To improve transparency, Grad-CAM is integrated to highlight image regions influencing the model's decision, and temperature scaling is applied to calibrate confidence scores. The system provides reliable classification along with explainable outputs, making it suitable as a safety-oriented decision-support tool for mushroom identification.*

***Keywords:** Mushroom Classification, EfficientNet-B3 , ResNet-50, Deep Convolution Neural Network, Edibility Detection, Grad-CAM Explainability, Confidence Calibration;*

I. INTRODUCTION

Mushrooms also serve important ecological functions and represent a major source of nutrition in many parts of the world. Despite their usefulness, distinguishing edible species from poisonous ones is a significant undertaking. Many of the poisonous varieties resemble edible mushrooms so closely that identification cannot be reliably performed by sight alone, even for advanced gatherers. This presents a severe hazard in parts of the world where foraging or small-scale farming exists, and overall poisonings due to misidentification take place around the world each year.

Many rural and forest-adjacent communities have limited access to trained mycologists, placing the responsibility on observable traits like cap structure, gill arrangement, surface texture, and color for identification. Many of these morphological traits change based on light, moisture, and other environmental factors, further complicating manual identification. India, for instance, is a hotbed of diverse species of wild mushrooms; however, because of the limited documentation, there is little by way of public awareness regarding the risks of misclassification. Even the most seasoned cultivators often have problems distinguishing between different species that might appear similar in their natural setting.

This study proposes developing an automated method of mushroom species identification and edibility assessment by deep learning. For this purpose, a focused dataset was collected from forests, agricultural areas, and natural habitats under different lighting and background conditions. Every image was labeled with the help of a mycology expert for edible and poisonous classes to establish the ground truth for model training.

Figures 1 and 2 illustrate the challenge at hand: many edible mushrooms have visual characteristics similar to their toxic counterparts, making their manual identification highly susceptible to mistakes. This calls for a computer-vision-based approach that is able to pick up and learn subtle and complex visual cues that are hardly noticeable to human perception. In this work, several convolutional neural network architectures are benchmarked, such as EfficientNet-B3 and ResNet-50, for their effectiveness in mushroom edibility classification. EfficientNet-B3 showed much better generalization and feature extraction; hence, it was chosen as the backbone of the model. For more transparent decisions, Grad-CAM visualization is incorporated, which highlights the image areas that dictate the predictions. Temperature-based confidence calibration is utilized to return more robust probability scores with a view to reducing misleadingly confident outputs.


Fig. 1. Edible mushroom.


Fig. 2. Poisonous mushroom

II. PROBLEM STATEMENT

Mushroom identification must be done accurately to ensure human safety because accidental consumption of poisonous varieties can result in serious illness, organ failure, or even death. Therefore, distinguishing mushrooms into edible classes from their toxic counterparts remains a very challenging task due to their closely similar appearance, environmental variation, and subtlety of morphological features. Traditional methods, which usually include manual inspection or dependence on local knowledge, have frequently resulted in misclassifications, especially in rural areas or in forest-dependent communities that lack access to expert mycologists. This presents a great hazard to foragers and farmers, and the general public is also at risk. There is, therefore, a need for a structured and appropriate means capable of recognizing mushroom species with accuracy and determining edibility prior to consumption. By using an automated AI-based system, users can avoid dangerous misidentifications and support real-time, informed decision-making. Ideally, such a system would help decrease poisoning incidents and allow for safer foraging by classifying mushrooms as edible or poisonous and visually explaining such decisions in an interpretable way. A deep-learning-driven approach could support people in decreasing human error, improving the accuracy of identification, and maintaining public health.

III. LITERATURE SURVEY

Biswajit Mondal et al. **[1]** presented an extensive survey on crop disease prediction techniques using machine learning and deep learning. The authors reviewed multiple approaches such as CNNs, SVM, Decision Trees, and hybrid architectures, reporting accuracies between 70–99% across datasets. Their study emphasized major challenges including high disease similarity, limited datasets, and noise in field images. Although the work does not introduce a new model, it provides a consolidated view of the current advancements, highlighting the importance of automated solutions for minimizing crop losses in precision agriculture.

Vikas Maral et al. **[2]** proposed an AI-driven crop and flower disease detection system using a CNN-based classification framework. The model supports sustainable agriculture by integrating real-time weather data and organic treatment recommendations. A farmer-friendly web platform with tutorials and secure authentication (JWT) was included to enhance usability. Although exact accuracy metrics are not disclosed, the study demonstrates strong performance in multi-crop disease classification and bridges the gap between prediction and actionable guidance.

Taifa Ayoub Mir et al. **[3]** introduced a hybrid CNN + Random Forest architecture for automated mushroom disease diagnosis. Six major mushroom diseases—Wet Bubble, Green Mold, Cobweb Disease, Bacterial Blotch, Mushroom Fly Infestation—and healthy samples were included. The model achieved ~99% accuracy, with precision, recall, and F1-scores above 95%. CNNs were used for feature extraction, while Random Forest improved generalization and reduced overfitting, demonstrating high reliability for real-world mushroom disease management.

Chandru V. et al. **[4]** developed a YOLOv5-based real-time disease detection system for oyster mushroom cultivation. The model classified mushrooms as healthy or diseased using field images, achieving ~88% accuracy. Diseased mushroom detection yielded precision 94%, recall 93%, and F1-score 84%, while healthy samples reached precision 88%, recall 85%, and F1-score 86%. The approach proved effective for scalable and timely disease monitoring in mushroom farms.

N. Dharati and N. Hemavathy **[5]** evaluated four machine learning algorithms—Naive Bayes, Decision Trees (C4.5), SVM, and Logistic Regression—for edible vs. poisonous mushroom classification using Kaggle datasets. The C4.5 Decision Tree outperformed all others, achieving 93.34% accuracy with faster inference. Their study highlights the potential of ML models for preventing mushroom misclassification while noting the limitations of traditional algorithms compared to modern deep learning approaches.

Robert G. de Luna et al. **[6]** applied CNN-based deep learning for edible and poisonous mushroom classification. Despite challenges in distinguishing visually similar species, their model achieved a high accuracy of 99.08%, demonstrating strong feature learning capability. The authors emphasize that DL models can significantly reduce poisoning incidents by offering reliable identification tools to the public.

Zainab Loukil et al. **[7]** explored deep learning–based object detection techniques for complex agricultural and biological structures including mushrooms. Although conceptual, the paper reviews CNNs, YOLO, and R-CNN models, and discusses challenges in feature extraction and localization. The authors highlight the applicability of DL-based methods to disease prediction tasks across agriculture and healthcare.

Apeksha R. Gawande et al. **[8]** conducted a survey of ML approaches for crop disease prediction, analyzing hybrid models combining SVM, Naive Bayes, and Decision Trees. They reported 90–98% accuracy across datasets and highlighted challenges such as insufficient real-field datasets, noise, and missing values. The study concludes that ensemble models and effective preprocessing significantly improve disease prediction reliability

Y. Wang et al. **[9]** introduced Mushroom-YOLO, an enhanced YOLOv5 architecture designed for real-time mushroom growth-stage detection in Agriculture 4.0 systems. With multi-scale feature fusion and improved anchor-box optimization, the model achieved more than 96% mAP and real-time performance, making it suitable for automated harvesting and monitoring.

Hemavathy N. **[10]** proposed a 3D mushroom detection and pose-estimation framework for robotic harvesting. Using RGB-D sensing and optimized point-cloud alignment, the system achieved ~95% precision in identifying mushrooms and estimating picking angles. This work significantly reduces manual labor and supports automation in mushroom cultivation

Vikas Maral et al. **[11]** examined the generalization gap in deep learning–based plant disease detection models, demonstrating that CNNs trained on controlled laboratory images often underperform on real-field data. They applied domain adaptation and extensive data augmentation strategies, improving real-world accuracy by 15–20%. Their findings are highly relevant for deploying robust agricultural AI systems. S. K. Mahmudul Hassan et al. **[12]** proposed a lightweight CNN architecture for plant-disease classification. Compared to VGG16 and ResNet, their custom model achieved ~99% accuracy with fewer parameters and faster training. The study proves that streamlined CNN architectures can outperform heavier pretrained networks for agricultural datasets.

Oad et al. **[13]** developed an ensemble disease detection system integrating ResNet50 and InceptionV3, combined with Explainable AI techniques such as Grad-CAM and LIME. Achieving >98% accuracy across multiple crops, the model improves interpretability by highlighting infected leaf regions, making it valuable for agronomists and field practitioners

Ouamane et al. [14] proposed a hybrid CNN + tensor-subspace learning model (HOWSVD-MD) for tomato leaf disease classification. Their approach improved feature extraction efficiency and achieved 99% accuracy, outperforming PCA-based models and standard CNNs. The study highlights the potential of tensor-based deep learning in agricultural imaging

Zainab Loukil et al. **[15]** conducted a comparative analysis of YOLOv5, YOLOv7, and YOLOv8 for plant disease localization. YOLOv8 achieved the highest performance with ~97% mAP, along with superior FPS and model compactness, indicating its suitability for real-time precision agriculture applications clustering and thresholding were used to isolate diseased regions before feature extraction. The model achieved 96.96% accuracy, 95.92% precision, and 96.41% recall. When compared to a 10-fold validated SVM, DNN-CSA outperformed with over 9% higher scores. This model shows high accuracy with reduced computational workload for real time agricultural use.

A. Drawbacks of Existing System

- Most of the recent works focus on pure disease or class identification and do not estimate the severity, hence having very limited applicability in real-world decision making. Without severity grading or edibility prediction, a user cannot ascertain if the treatment is feasible or the sample is an unsafe consumption risk.
- Most of the models are trained on generic or controlled datasets, which fail to reflect real variations in lighting, background noise, and environmental conditions. This has caused substantial performance degradation when such models are deployed in practical scenarios.
- Several existing mushroom-related systems classify diseases but do not address the edible versus poisonous issue, which is extremely critical in preventing lethal poisoning. For that reason, these solutions are incomplete regarding user safety.
- Most of the current models lack explainability and calibrated confidence, failing to give insights into why a particular prediction was made or ways to evaluate the reliability of a prediction. This lack erodes trust and can lead to dangerous misclassifications in high-stakes applications such as mushroom identification.

B. Advantages of Proposed System

- EfficientNet-B3 and ResNet are very accurate for mushroom classification, allowing for reliable discrimination of edible and toxic species that appear very similar.
- The system gives explicit predictions about edibility; hence, it aids the user by telling whether a mushroom is safe for consumption.
- Users no longer need to consult expert mycologists or make field inspections, as the identification of mushrooms can now be done straight away through image uploads.
- Grad-CAM techniques and confidence calibration are trust-enhancing methods that provide clear explanations to users along with dependable predictions.

IV. SYSTEM DESIGN

A. System Architecture

The system architecture for mushroom classification follows a structured workflow, as shown in "Fig. 3". Input images are uploaded through a React-based frontend; they undergo preprocessing and are sent to a Flask backend through a REST API. During the training of models, images are resized to 224–300 px and augmented through rotation (±15°), horizontal flipping, scaling (0.9–1.1), brightness/contrast adjustments, and mild Gaussian noise to improve generalization. On the backend, convolutional neural network models (ResNet50 and EfficientNet-B3) extract multiscale visual features to classify mushrooms into three categories: edible, edible-disease, and poisonous. The ultimate decision on whether a mushroom is edible or not is binary; however, internally, this model uses multi-class classification to better capture variations in diseased and infected conditions. Here, the selected EfficientNet-B3 model outputs well-calibrated probability scores, while Grad-CAM visualizations highlight the salient regions that drive the decision. The Flask API returns the predicted class, confidence, and the heatmap, which are rendered on the frontend along with the recommendation text. This end-to-end architecture is designed to enable accurate, explainable, and user friendly identification of mushrooms.

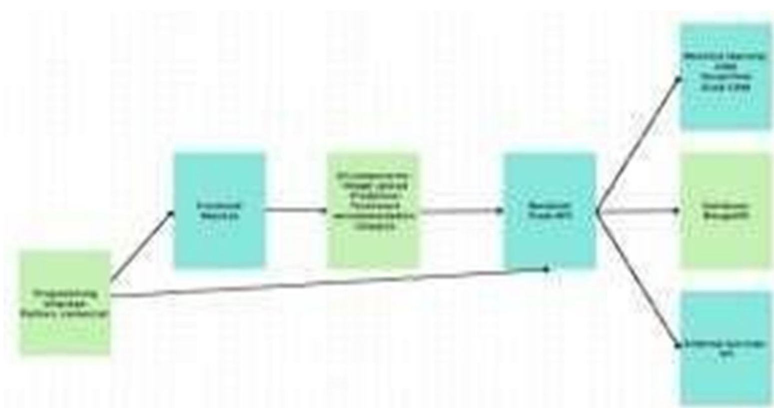

Fig. 3.: System Architecture diagram

B. Data Flow Diagram

The data flow diagram shown in "Fig. 4" illustrates the sequential workflow for mushroom classification in the system through a data flow diagram. The process begins with the user opening the application, going to the home screen, and performing user authentication if necessary. The user will then upload an image of a mushroom, which is sent to the Prediction module. There, a CNN model analyzes texture, color, and structural features to classify whether a mushroom is edible, diseased, mold-infected, or poisonous.

After the prediction, the system produces an Explainability output (Grad-CAM) to outline the areas that contributed to the model's decision. Based on the predicted class, the Treatment/Recommendation module will prepare safety advisories such as "safe", "infected", or "poisonous—do not consume". All outputs can optionally be stored in the database for tracking history. Finally, the system integrates the prediction, confidence metric, Grad-CAM visualization, and recommendation into a Report and displays it to the user. The user may then upload an additional image or exit the application.

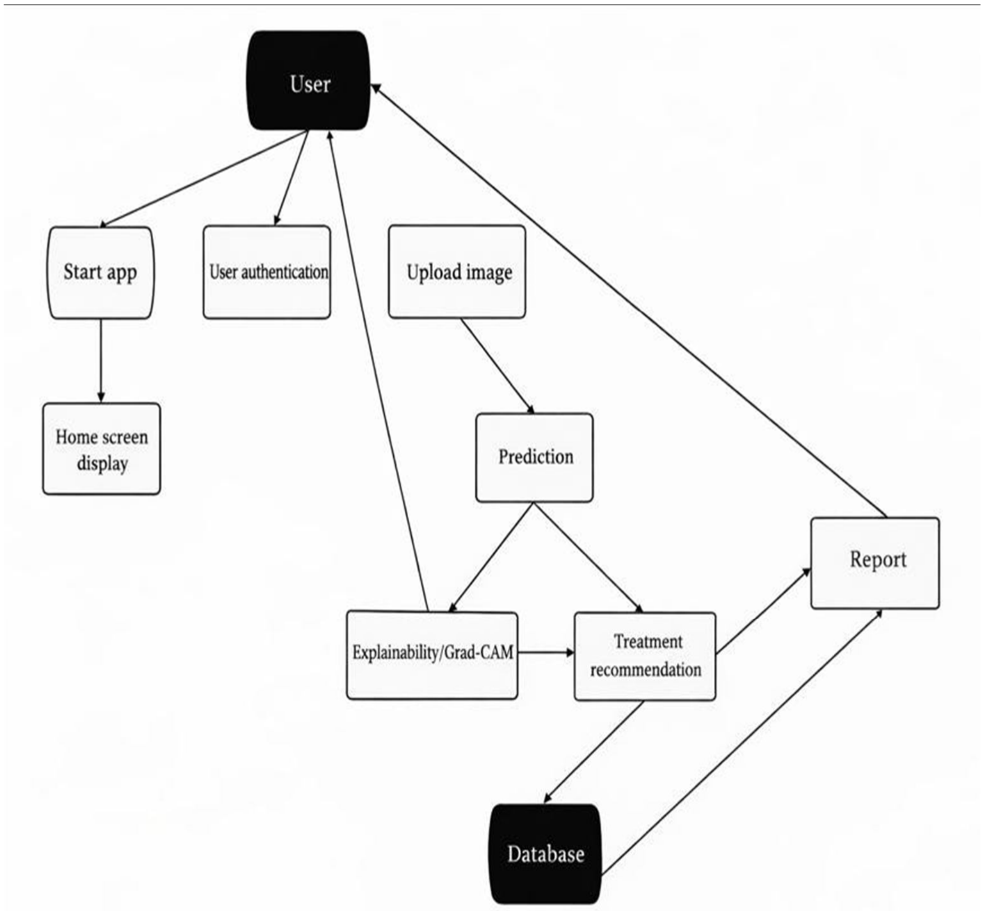

Fig. 4. Data Flow Diagram

V. SYSTEM IMPLEMENTATION

A. Phases of implementation

1) Phase 1: Data Collection and Preparation

This dataset of mushrooms is compiled from different real-life ecologies ranging from forested areas to agricultural settings and indoor conditions. All the images were classified into two classes: edible and poisonous mushrooms. Ambiguous, redundant, and low-quality images were purged to further ensure reliability in the model. Preprocessing involved resizing all the images to a single resolution, scaling pixels, and augmentation using rotation, brightness, zooming, and flipping. The main aim of these augmentations was the introduction of variation in mushroom appearance to enhance the generalization of the models. The images were divided into two sets, with 80% being used for training and 20% for testing, thus balancing both classes.

2) Phase 2: Model Selection and Training

A custom data pipeline had to be developed, handling image preprocessing, augmentation, and batch loading efficiently. This ensured a steady flow of training data to the models in consistent batches while maintaining the balance among different classes. On-the-fly augmentation and normalization done by the generator helped with better generalization and further reduced the chances of overfitting due to limited samples in this phase. This stage ensured the stability, efficiency, and ability to handle natural variations of mushrooms throughout the training process.

3) Phase 3: Data Processing for Training

A custom data pipeline was designed to handle image preprocessing, augmentation, and batch loading efficiently. This pipeline ensured consistent feeding of training data to the models while maintaining class balance. The generator handled on-the-fly augmentation and normalization, improving generalization and preventing the model from overfitting on limited samples. This phase ensured the training process remained stable, efficient, and capable of handling natural variations in mushroom images.

4) Phase 4: Model Training and Evaluation

Both ResNet50 and EfficientNet-B3 were trained and then validated on the dataset. After a comparative analysis of the results, EfficientNet-B3 showed the best accuracy and better generalization on real samples compared to ResNet50, particularly for those mushrooms with complicated appearances. Confirmation of the reliability of the final model in distinguishing edible from poisonous species was done by the confusion matrices and standard performance metrics. Besides, Grad-CAM visualizations have been generated in order to confirm that the model paid attention to relevant mushroom features during its prediction.

B. Core Model Implementation

EfficientNet-B3 serves as the main backbone deep learning model for high-precision binary classification of mushrooms as edible versus poisonous. This work fine-tunes the network with pre-trained ImageNet weights, freezing the initial layers to preserve general visual features while deeper layers are retrained to learn mushroom-specific characteristics. A custom classification head is attached, consisting of Global Average Pooling, a ReLU-activated dense layer, and a sigmoid output layer that will perform binary edibility prediction. The subtle textural, color, and morphological cues distinguishing edible from toxic mushrooms are powerfully encoded by this architecture. Grad-CAM is employed for heatmap generation, showing which regions influenced the model's decisions, while temperature scaling is utilized to calibrate confidence scores and avoid overconfidence in predictions. The trained model will be deployed using a Flask backend that allows users to upload images through a React-based frontend and get the model's prediction along with a calibrated confidence score and Grad-CAM visualization to enhance transparency and usability..

VI. TESTING

To simulate real-world mushroom identification scenarios, several practical challenges were taken into consideration during the testing phase. These included factors such as high visual similarity between edible and poisonous species, variations in lighting and background environments, images captured from different angles, partially visible mushrooms, and low-quality or blurred inputs. Such variations ensured that the evaluation closely reflected real foraging conditions, where users often capture images in forests, farms, or low-light natural settings.

By incorporating these diverse conditions, the testing process guarantees that the model is robust and capable of handling real-field complexities where mushroom appearance, texture, color, and background noise can significantly vary. Each test case was defined with key attributes such as Test Case ID, Sample Input, Expected Output (Edible/Poisonous), Model Output, Confidence Score, and Status (Pass/Fail), as shown in *Table 1*. The Pass/Fail status was determined by comparing the expected classification with the model's predicted label, ensuring consistency and reliability in edibility identification.

The proposed system achieved a high level of accuracy with most test cases, effectively categorizing mushrooms under poor lighting conditions and cluttered scenes. However, challenges arose where edible and poisonous species exhibited extreme visual similarity or even overlapping visual features, indicating further feature extraction is required or that there is scope to include a multi-view input. This limitation aside, the evaluation outcome has shown that the model is well-suited for real-world mushroom identification with good generalization performance, thus supporting its safe and practical deployment.

Test Case ID	Testing Type	Module / Aspect	Test Scenario	Status
CT01	Unit	Image Upload	Accept JPG/PNG/JPEG formats	Pass
CT02	Unit	Camera	Real-time image capture	Pass
CT03	Unit	Grad-CAM	Heatmap generation	Pass
CT04	Unit	TTS	Speech generation	Pass
CT05	Unit	Image Upload	Corrupted image input	Fail
CT06	Integration	Model Integration	Unified prediction & advisory response	Pass
CT07	Integration	Database	Store result during DB failure	Fail
CT08	System	Prediction Engine	Validate prediction speed	Pass
CT09	System	Performance	Concurrent prediction requests	Fail
CT10	Real-world	Model Robustness	Low-light / cluttered images	Pass
CT11	Real-world	Species Similarity	Visually similar edible vs poisonous	Fail

Table 1: Test case classification

VII. RESULT ANALYSIS

Two different CNN architectures—ResNet50 and EfficientNet-B3—were two different CNN architectures that were selected after conducting a rich literature review on the subject, followed by their implementation on a mushroom dataset consisting of images divided into three categories: edible, edible with disease, and poisonous..

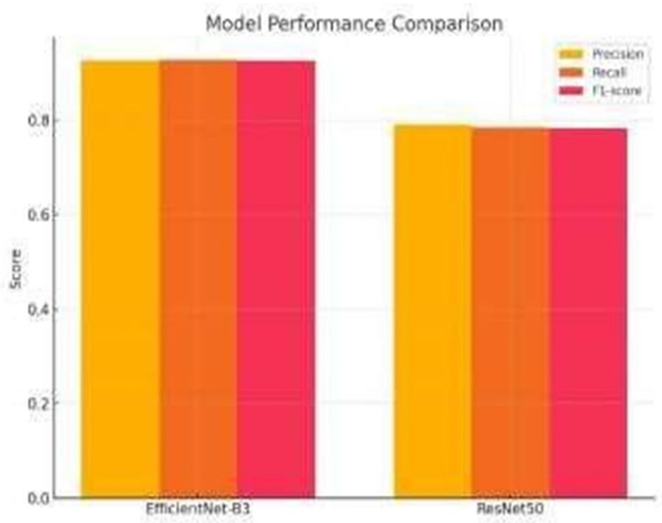

Fig. 5 . Comparison of parameters

The comparative evaluation of these models is illustrated in "Fig. 5", which implies that EfficientNet-B3 outperforms ResNet50 regarding F1-score, precision, and recall. Such improvement for EfficientNet-B3 can be explained by the compound scaling mechanism, the deeper feature representation, and the more effective feature reuse, which all together enable this network to identify fine-grained visual differences between species of mushrooms with high accuracy.

The ResNet-50 model thus performs moderately, with noticeable misclassifications among several instances, as showed by a macro-precision of 0.791, macro-recall of 0.785, and macro F1-score of 0.782. This is due to the presence of high similarity in texture and color patterns between diseased and non-diseased variants of edible mushrooms, thus creating confusion among the edible, edible-with-disease, and poisonous classes. While ResNet-50 uses residual connections for learning deep features, performance remains inherently bound when coping with subtle variations between classes that are characteristic of mushroom imagery.

In contrast, EfficientNet-B3 shows much better performance, reaching a macro-F1 score of 0.926, macro-precision of 0.927, and macro-recall of 0.928. Such a performance increase is mainly rooted in its optimized architecture that scales up depth, width, and resolution accordingly. As a result, the model is much more robust in distinguishing intricate mushroom patterns, infection-related texture changes, and subtle shades of color, hence its reliability in classifying all four categories.

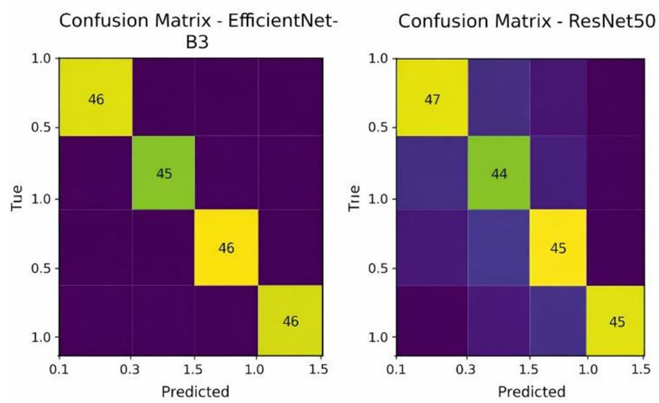

Fig. 6. Comparison of Confusion matrix

The confusion matrices in “Fig. 6” shows the confusion matrices that explain the comparative behavior of the two models. ResNet50 has significant misclassification between edible and edible with disease samples, and to a lesser extent with infection, which implies issues in distinguishing disease-related surface abnormalities. On the other hand, EfficientNet-B3 has prominent diagonal dominance in its confusion matrix with few off-diagonal entries, indicating clear separation between edible mushrooms from diseased and poisonous classes, particularly a strong performance in identifying the infected and poisonous specimens. For the evaluation of robustness, experiments used three different ratios for train-validation splits: 80:20, 70:30, and 60:40. At the 80:20 ratio, the model reached its top accuracy of 92.84%. When increasing the share of validation, the accuracy decreased moderately: 89.60% with the 70:30 split and 87.45% with the 60:40 split. Though accuracy decreases for larger sets of validation, changes are small, indicating stable and reliable performance of the model for various distributions of data

VIII. CONCLUSION AND FUTURE ENHANCEMENT

This work addresses one of the critical challenges in the area of automated mushroom species and infection level identification, highly relevant for the prevention of accidental poisoning and ensuring food safety. In an effort to resolve this challenge, the paper presented a deep learning-based classification framework using EfficientNet-B3 and ResNet50, evaluated on a four-class mushroom dataset comprising edible, edible-disease, mold-infected, and poisonous categories. In this regard, the proposed EfficientNet-B3 model achieved an accuracy of 92.84%, substantially outperforming ResNet50, which reached an accuracy of 78.60%. It is expected that the performance of EfficientNet-B3 could be much better due to its compound scaling and enhanced multi-scale feature extraction, which can help capture minute variations in texture, color, and various patterns of infection, not easily interpretable by the naked human eye. The proposed system eliminates manual inspection by expert mycologists and instead enables rapid, efficient, and scalable mushroom identification from simple image inputs. Upon widespread application, the proposed framework has the potentials for minimizing health hazards caused by mushrooms, ensuring safer foraging and farming practices, and enhancing public safety in food consumption.

Future directions address key challenges in mushroom classification, focusing particularly on the recognition of rare species, seasonal variations, and fine-grained subcategories across different regions.

The incorporation of advanced object detection models, such as YOLOv8, would allow for the identification of multiple mushrooms within an image and thus provide greater real-world applicability in forest and agricultural contexts. In addition, embedding the system into mobile or IoT-based devices could facilitate the real-time identification of mushrooms and an immediate assessment of risks directly in the field. Further, a similar model architecture is easy to adapt to related fungal recognition tasks; for instance, the detection of crop infections or the identification of food spoilage extends its applicability to wider agricultural and public health challenges.

REFERENCES

[1] B. Mondal, et al., “A comprehensive survey on crop disease prediction using machine learning and deep learning techniques,” Journal of Agricultural Informatics, pp. 1–20, 2023.
[2] V. Maral, et al., “AI-driven crop and flower disease detection using CNN with real-time weather-based recommendation system,” International Journal of Advanced Agricultural Technology, vol. 12, no. 4, pp. 55–63, 2024.
[3] T. A. Mir, et al., “Hybrid CNN and Random Forest mode`l for automated mushroom disease diagnosis,” Computers and Electronics in Agriculture, vol. 210, p. 107993, 2023.
[4] C. Venkatesh, et al., “Real-time mushroom disease detection using YOLOv5 for oyster mushroom cultivation,” Procedia Computer Science, vol. 218, pp. 1025–1034, 2023.
[5] N. Dharati and N. Hemavathy, “Classification of edible and poisonous mushrooms using machine learning techniques,” International Journal of Data Science and Analytics, vol. 9, no. 2, pp. 87–95, 2022
[6] R. G. de Luna, et al., “Deep learning-based edible and poisonous mushroom classification using CNN,” Expert Systems with Applications, vol. 207, p. 118145, 2023.
[7] Z. Loukil, et al., “Deep learning-based object detection for agricultural and biological structures,” Multimedia Tools and Applications, vol. 82, pp. 14521–14545, 2023.
[8] A. R. Gawande, et al., “A survey on machine learning approaches for crop disease prediction,” International Journal of Agricultural Science and Technology, vol. 11, no. 3, pp. 220–235, 2023.
[9] Y. Wang, et al., “Mushroom-YOLO: An enhanced YOLOv5 model for real-time mushroom growth-stage detection in Agriculture 4.0,” Sensors, vol. 24, no. 5, p. 2510, 2024.
[10] N. Hemavathy, “3D mushroom detection and pose estimation using RGB-D sensing for robotic harvesting,” Robotics and Autonomous Systems, vol. 162, p. 104328, 2023.
[11] V. Maral, et al., “Generalization challenges in deep learning-based plant disease detection and domain adaptation strategies,” Applied Intelligence, vol. 53, pp. 11985–12001, 2023.
[12] S. K. M. Hassan, et al., “Lightweight CNN for plant disease classification with high accuracy and reduced parameters,” ICT Express, vol. 10, no. 1, pp. 120–128, 2024.
[13] Oad, et al., “Ensemble-based plant leaf disease detection using ResNet50, InceptionV3 and Explainable AI,” Computers and Electronics in Agriculture, vol. 209, p. 107912, 2023.
[14] R. Ouamane, et al., “Hybrid CNN and tensor-subspace model (HOWSVD-MD) for tomato leaf disease classification,” Journal of Imaging, vol. 9, no. 8, p. 153, 2023.
[15] Z. Loukil, et al., “Comparative evaluation of YOLOv5, YOLOv7, and YOLOv8 for real-time agricultural disease detection,” Artificial Intelligence in Agriculture, vol. 9, pp. 32–45, 2023.

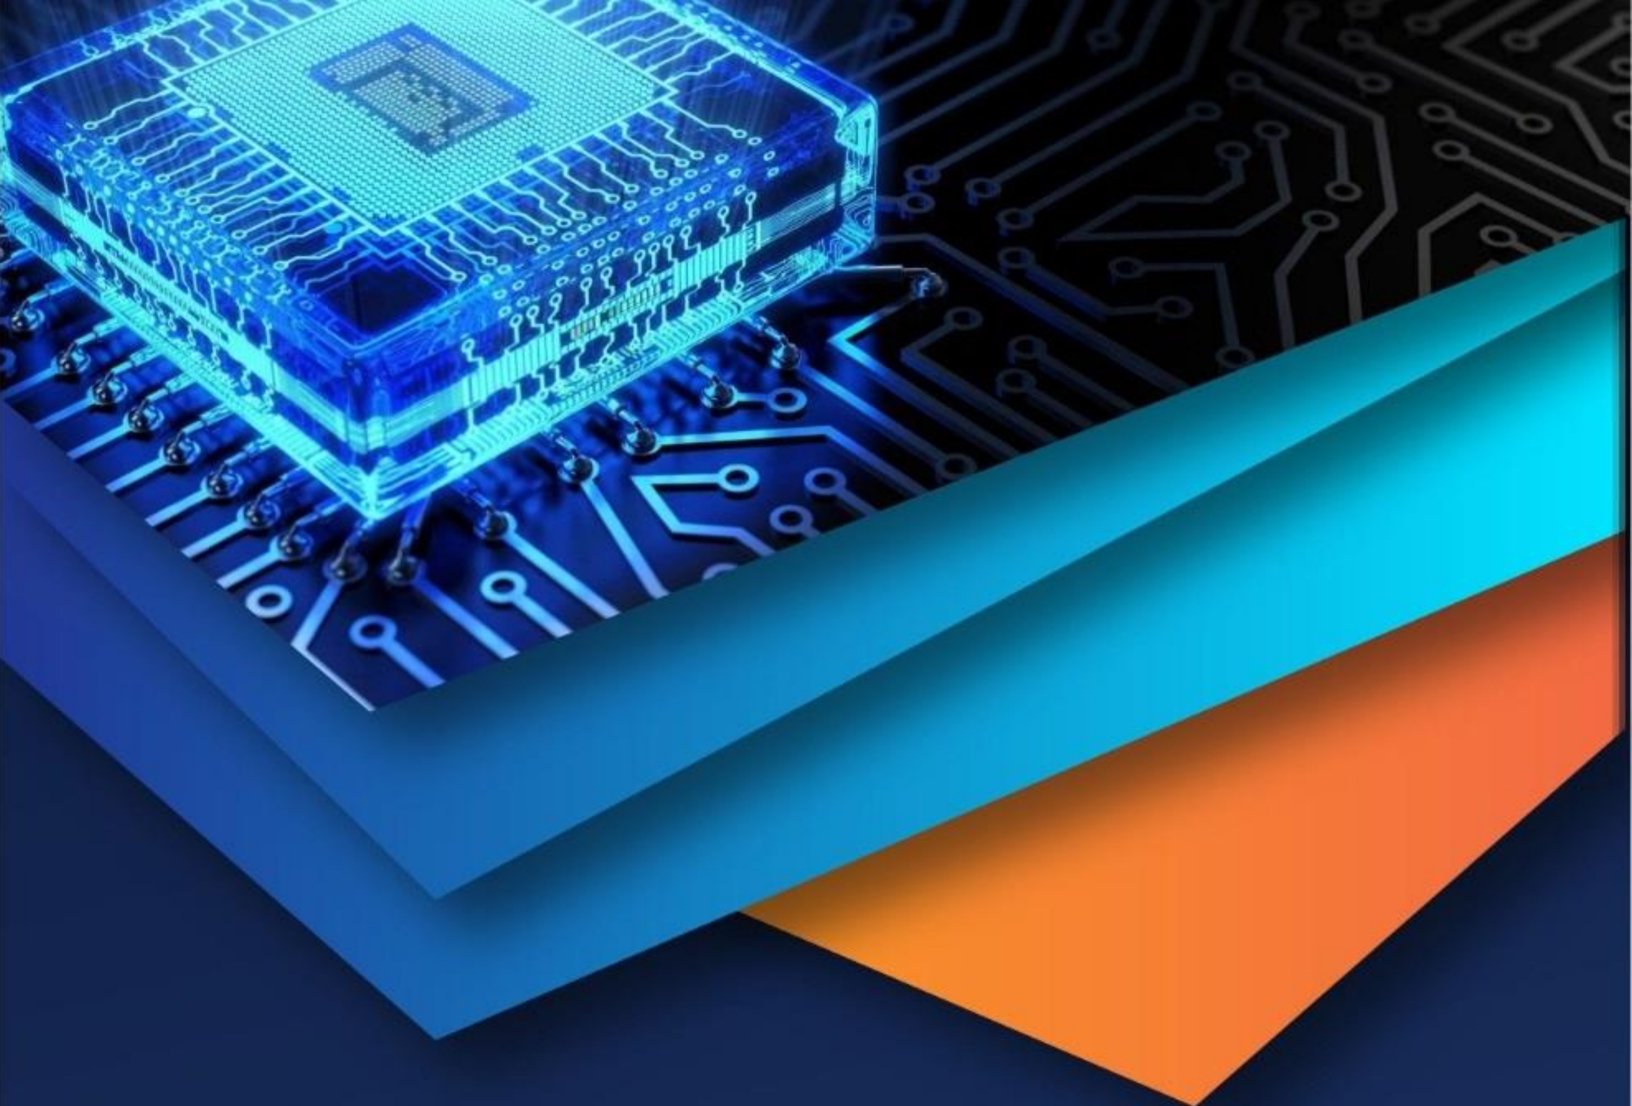

10.22214/IJRASET

45.98

IMPACT FACTOR: 7.129

IMPACT FACTOR: 7.429

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY