SAFE BITE: A Context-Aware AI Platform for Fruit Contamination Risk Assessment

1st Fathima Sana K T, 2nd Fathimath Rasmiya C, 3rd Shibla Hameed, 4th Mehnajabin
Department of Computer Science and Engineering
MGM Technological Campus
Valanchery, India

***Abstract*—Safe Bite is a mobile-based health monitoring system designed to prevent the spread of zoonotic diseases caused by contaminated fruits, particularly those affected by bat bites. The system leverages artificial intelligence and computer vision techniques to analyze fruit images captured via a smartphone camera and detect potential contamination.**

A risk fusion algorithm integrates image-based predictions, geospatial clustering, and crowd-sourced reports to generate a reliable risk score. Based on this score, the system delivers real-time alerts to users and health authorities, enabling rapid response and early outbreak detection. Safe Bite functions as an intelligent early warning system aimed at reducing public health risks and enhancing community awareness.

***Index Terms*—Safe Bite, computer vision, MobileNetV2, risk fusion algorithm, geospatial analysis, public health**

I. Introduction

Zoonotic diseases, such as the Nipah virus, pose a severe threat to public health in regions where fruit contamination by bats is frequent. While existing research has utilized Convolutional Neural Networks (CNNs) for general fruit quality grading and defect detection, these systems typically focus on aesthetic attributes like ripeness or bruising. Safe Bite differentiates itself by focusing specifically on biological contamination signatures that signify public health risks.

By targeting home users and local vendors, the platform decentralizes early warning systems. The integration of real-time geospatial data with on-device AI detection allows for a "Context-Aware" response, transforming a standard smartphone into a proactive tool for disease prevention.

II. System Architecture

The Safe Bite system integrates mobile application development, machine learning, and geospatial data processing into a unified architecture. The design prioritizes efficiency, scalability, and real-time responsiveness.

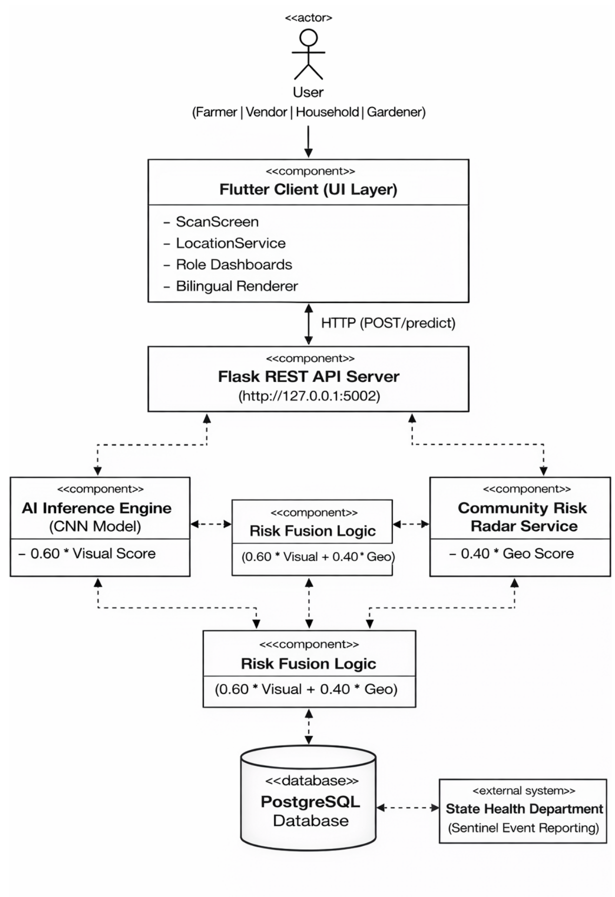

Fig. 1: System architecture of the Safe Bite platform illustrating the interaction between the Flutter client, Flask REST API server, AI inference engine, community risk radar service, and PostgreSQL database.

A. Image Classification (MobileNetV2)

The system utilizes the MobileNetV2 architecture, a lightweight Convolutional Neural Network optimized for mobile devices. This enables efficient on-device image processing with minimal latency, which is critical for real-time field use by farmers and vendors.The model was trained using a custom-curated dataset of 1,200 labeled images. These were derived from a combination of field-captured data and relevant public repositories to ensure the inclusion of specific contamination signatures like bat bites and bird pecks. The model is trained to classify fruits into five distinct categories:

- **Healthy**: Fruits with no visible damage.
- **Bat bite (clear)**:Deep, identifiable teeth marks.
- **Bat bite (superficial)**:Surface-level biological signatures.
- **Bird peck**: Small, localized puncture marks
- **Insect damage**: Typical pest-related surface irregularities.

Training results demonstrate strong convergence, achieving approximately 96.4% validation accuracy with a loss reduction to 0.12 over six epochs. This rapid convergence is attributed to the use of Transfer Learning with pre-trained weights, allowing the model to stabilize quickly on the specialized fruit dataset without overfitting.

B. Geospatial Analysis (GIS)

The backend of the Safe Bite platform utilizes Geographic Information System (GIS) technology to add environmental context to the AI's visual predictions. Powered by PostgreSQL and the PostGIS extension, the system maintains a spatial database of historical contamination clusters.

When an image is captured, the system executes a SQL-based spatial query to determine the proximity of the user to known zoonotic "Hotspots". This proximity data is used to calculate the Location Risk (L). For instance, if a scan is performed within a designated outbreak zone, the value of L is increased, ensuring that the final fused risk score (R) reflects the environmental danger even if visual symptoms on the fruit are superficial.

III. Proposed Methodology

A. Risk Fusion Algorithm

To improve reliability, Safe Bite employs a risk fusion approach that combines multiple data sources into a single score:

$$R = w_1 C + w_2 L + w_3 U \tag{1}$$

where:

- R is the final risk score
- C is the AI confidence score
- L is the location-based risk score
- U is the user-reported risk score

TABLE I: Risk Fusion Parameters

Parameter	**Description**	**Range**
C	AI Confidence Score	0 to 1
L	Location Risk Score	0 to 1
U	User Report Score	0 to 1

All values are normalized between 0 and 1. The weights (w_1, w_2, w_3) are used to balance the influence of each component.

B. Real-Time Alerting

When the computed risk score exceeds a predefined threshold, the system triggers real-time alerts to users. The data is also logged for health authorities, enabling faster response to potential outbreaks.

IV. Results and Discussion

The performance of the Safe Bite platform was evaluated using a testing subset of the 1,200 labeled images. The evaluation focused on the MobileNetV2 model's ability to categorize fruit health into five distinct classes: healthy, bat bite (clear), bat bite (superficial), bird peck, and insect damage.

A. Training Performance and Convergence

The MobileNetV2 architecture was trained over 6 epochs, as the model reached a performance plateau where further iterations yielded diminishing returns in validation accuracy. As shown in Fig. 2, the rapid convergence to 96.4% validation accuracy and a loss of 0.12 suggests that the pre-trained weights of MobileNetV2 effectively captured the necessary features for this specialized dataset without requiring extensive further training.

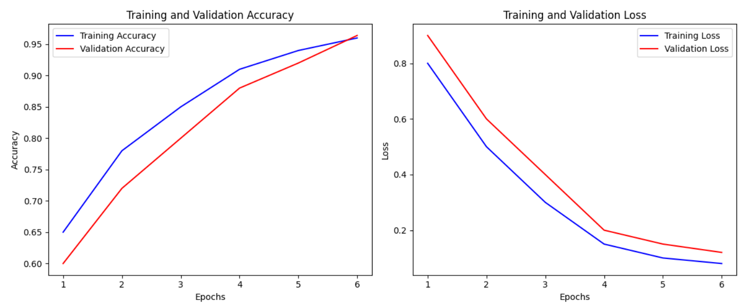

Fig. 2: Training and validation accuracy/loss over 6 epochs showing convergence at 96.4% validation accuracy.

B. Multi-Class Performance Metrics

While the training accuracy was high, the multi-class evaluation on the final test set (N=101) provided a more nuanced view of real-world performance. The model achieved an overall weighted accuracy of 66%. As shown in Table II, the model demonstrated an exceptional recall of 91% for healthy fruits, with a corresponding F1-score of 0.82.

TABLE II: Multi-Class Classification Performance

Class Label	**Precision**	**Recall**	**F1-Score**	**Support**
Bat Bite (Clear)	0.67	0.50	0.57	16
Bat Bite (Superficial)	0.58	0.50	0.54	14
Bird Peck	0.71	0.56	0.62	9
Healthy	0.74	0.91	0.82	35
Insect Damage	0.56	0.56	0.56	27
Weighted Avg	**0.66**	**0.66**	**0.65**	**101**

C. Error Analysis and Confusion Matrix

The Confusion Matrix in Fig. 3 highlights the specific classification challenges. The model successfully classified 32 out of 35 healthy samples. However, there is visible overlap between "insect damage" and "bat bite (superficial)". This misclassification is attributed to the visual similarity of small puncture marks across different contamination sources.

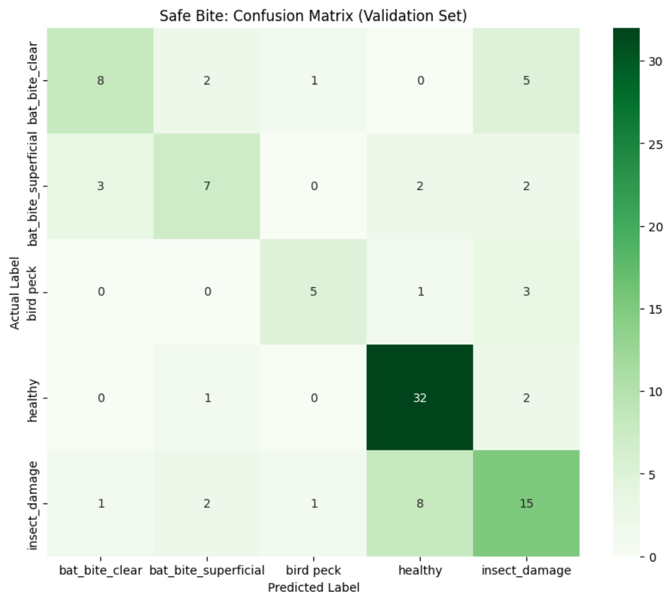

Fig. 3: Confusion matrix illustrating the classification distribution and identification of healthy vs. contaminated fruit samples.

D. Comparative Analysis

The performance of Safe Bite was evaluated against methodologies established in recent literature. The system achieved a validation accuracy of 96.4%, which is highly competitive with the results reported in [6], where accuracies of 97.5% for apples and 92.5% for mangoes were observed. While the multi-input architecture introduced in [7] reports accuracy reaching 100%, it relies on a complex preprocessing pipeline involving silhouette segmentation that often requires manual refinement. In contrast, the proposed system is specifically optimized for real-time mobile deployment. By bypassing the computationally intensive layers utilized in [7], Safe Bite provides a low-latency, practical early-warning tool for field use.

V. Conclusion

Safe Bite represents a proactive and context-aware solution for detecting fruit contamination risks. By combining MobileNetV2-based image classification with geospatial intelligence and user-generated data, the system enhances detection accuracy and supports early intervention.

The model's high recall for healthy fruits ensures a reliable user experience, while the risk fusion mechanism compensates for classification limitations.

VI. Future Work

Future improvements will focus on:

- Expanding the dataset for better generalization
- Improving classification accuracy for visually similar contamination types
- Implementing adaptive learning techniques to optimize risk fusion weights

VII. Acknowledgment

The authors would like to acknowledge Ms. Shabna M for her role as project guide during the development of this work. We also express our gratitude to the Department of Computer Science and Engineering at MGM Technological Campus, Valanchery, for providing the necessary facilities and academic environment to carry out this research.

References

[1] C. Pachón-Suescún, J. Pinzón, et al., “Fruit Identification and Quality Detection via DAG-CNN,” 2020.

[2] R. Nithya et al., “Computer Vision System for Mango Defect Detection Using Deep CNN,” 2022.

[3] M. Agarwal et al., “Differential Evolution Based Compression of CNN for Apple Disease Classification,” 2022.

[4] M. S. Ibrahim et al., “CNN Comparative Study for Apple Quality Classification,” 2022.

[5] B. Buyukkarakaya and E. Ulker, “Classification of Physiological Disorders in Apples Using Deep CNN,” 2023.

[6] R. Pacheco, P. González, L. Chuquimarca, B. Vintimilla, and S. Velastin, “Fruit Defect Detection Using CNN Models with Real and Virtual Data,” 2023.

[7] L. Chuquimarca, B. Vintimilla, and S. Velastin, ”Classifying Healthy and Defective Fruits with a Multi-Input Architecture and CNN Models,” arXiv preprint arXiv:2410.11108, Oct. 2024.