Intelligent Grain Warehouse Monitoring, Pest Identification and Prevention System

Sanjaikumar V.[1], Vijayalakshmi S.[2], Marikkannan M.[3]
[1]M.E. Computer Science and Engineering (PG Scholar), Department of CSE, Government College of Engineering, Erode, Tamilnadu, India.
[2]Assistant Professor,Department of CSE, Government College of Engineering, Erode, Tamilnadu, India.
[3]Assistant Professor (Senior), Department of CSE, Government College of Engineering, Erode, Tamilnadu, India.

***Abstract*—post-harvest grain losses in developing nations account for 10–40% of annual production, attributed to two concurrent threats: hazardous environmental conditions (excess temperature, humidity, and CO gas buildup) and stored product insect infestations. Existing systems address these threats in isolation, leaving a critical gap in integrated risk assessment. This paper presents a dataset-driven intelligent monitoring framework that unifies environmental anomaly detection with deep learning-based insect species identification into a composite risk-scoring system. Using publicly available datasets the Smoke Detection IoT dataset and the IP102 insect pest image collection the system trains a Random Forest and a two-layer LSTM for environmental anomaly detection and a fine-tuned ResNet-50 Convolutional Neural Network for classifying five stored product insect species. A novel data fusion engine combines both model outputs using a weighted formula to generate a composite risk score mapped to three actionable alert levels: Stable, Moderate Risk, and High Risk. A tiered SMS notification system dispatches contextual alerts upon risk-level transitions, achieving alert precision of and recall. This work constitutes the first dataset-driven integration of environmental IoT monitoring and insect pest identification into a unified grain storage risk assessment platform.**

***Index Terms*—Anomaly detection, Deep learning, Food security, Grain storage, Insect classification, LSTM, Post-harvest loss, Smart agriculture.**

I. INTRODUCTION

The global challenge of post-harvest food loss represents one of the most persistent threats to food security, particularly in developing agricultural economies. According to the Food and Agriculture Organization (FAO), losses during grain storage account for 10–40% of annual production in countries like India, placing enormous pressure on national food supply chains, rural income, and food pricing stability [1]. Two distinct but interrelated threat categories account for the majority of these losses.

The first category encompasses hazardous environmental conditions. Elevated temperature (above 50°C), excessive relative humidity (above 50%RH), toxic gas accumulation (carbon monoxide above 25 ppm), and smoke events collectively accelerate grain deterioration through fungal growth (Aspergillus, Penicillium), mycotoxin production, accelerated respiration, and weight loss [2]. In large-scale grain warehouses holding thousands of tonnes of cereal crops, even a 12-hour delay in detecting such conditions can render entire storage batches unsuitable for human consumption.

The second category involves stored product insect infestations. Five species are responsible for the majority of global stored grain losses: Rhyzopertha dominica (lesser grain borer), Cryptolestes ferrugineus (rusty grain beetle), Tribolium castaneum (red flour beetle), Sitophilus oryzae (rice weevil), and Oryzaephilus surinamensis (saw-toothed grain beetle). These pests colonise grain silently, with detectable damage often occurring only after 30–60 days of infestation, by which point contamination has spread beyond recovery [2]. Current industrial practice involves manual inspection at 3–7-day intervals by trained agricultural officers, supported in some facilities by basic temperature and humidity sensors. This approach is insufficient: manual inspection is subjective, expensive at scale, and inherently reactive rather than preventive. No currently deployed system in developing agricultural economies provides automatic, real-time integrated monitoring of both environmental conditions and biological pest presence.

Two parallel research streams have emerged to address these threats independently. IoT-based sensor systems [3] monitor environmental parameters in real time but lack any biological pest detection capability. Deep learning-based image classifiers [4],[5] identify insect species from photographs with state-of-the-art accuracy but operate without environmental context, making them unable to predict escalating risk before visible infestation occurs. The fusion of both streams into a unified, continuously operating risk-scoring framework represents a significant unexplored research opportunity with direct practical impact.

This paper presents a dataset-driven software system that unifies these two streams. The key contributions are: (1) a complete, replicable, hardware-free implementation using only publicly available datasets; (2) a dual-model pipeline combining LSTM-based temporal anomaly detection (achieving F1 = 0.94) with ResNet-50 transfer learning for insect classification (achieving 98.7% accuracy); (3) a novel weighted composite risk score validated by statistical correlation analysis; (4) a tiered SMS alert system with three risk levels achieving 96% precision and 94% recall; and (5) Grad-CAM morphological explainability for all classifier predictions.

II. LITERATURE REVIEW

A. Stored Product Insect Detection Using Deep Learning

The application of deep learning to stored product insect identification has advanced significantly in recent years. Shi et al were among the first to systematically demonstrate that deep neural networks outperform conventional image processing methods for stored grain insect detection, achieving state-of-the-art results on laboratory-collected datasets [4]. Their work established the feasibility of camera-based automated monitoring but did not address environmental integration.

Badgujar et al conducted the most comprehensive comparative study to date, applying four pre-trained CNN architectures ResNet-50, MobileNetV2, DarkNet-53, and EfficientNet-b0 to classify five stored product insect species, reporting overall accuracy between 96% and 99.3% [5]. Crucially, this work introduced Grad- CAM visualisations that confirmed the models were attending to entomologically meaningful morphological features, providing a foundation for the explainability component of the present study. Mendoza et al evaluated machine learning models specifically for real-time insect monitoring in grain handling facilities, identifying practical deployment constraints not addressed by laboratory studies [6].

More recently, Yu et al proposed PDA-YOLO, a high-precision architecture for real-time stored grain pest detection that leverages partial dynamic attention mechanisms to handle the small object detection challenges characteristic of insect imagery [7]. While PDA-YOLO demonstrates superior detection performance, it requires substantial computational resources and does not address environmental fusion. Chithambarathanu and Jeyakumar provided a comprehensive survey of deep learning approaches to crop pest detection more broadly, confirming that no existing work had integrated environmental monitoring with visual pest classification [14].

B. IoT-Based Environmental Monitoring for Grain Storage

The use of IoT sensors for grain storage monitoring has a longer history, with Lydia et al presenting one of the most complete implementations [3]. Their system employs a PIC16F877A microcontroller connected to temperature, humidity, CO, smoke, vibration, and PIR sensors, with data transmitted to a ThingSpeak cloud platform and SMS notifications dispatched via the GSM module upon threshold violations. While practically deployed and validated, this system provides no identification capability for the biological pest dimension of grain storage risk. Duan et al approached the environmental monitoring problem from a predictive rather than reactive perspective, proposing a deep spatiotemporal attention approach for forecasting temperature distributions in large grain storage silos using historical sensor readings [8]. Their LSTM-based forecasting model demonstrated the applicability of recurrent architectures to temporal sensor data in grain storage contexts, directly informing the LSTM design choices in the present work. Kodali and John (2020) demonstrated a lightweight Arduino-based IoT monitoring system suitable for small-scale warehouses [17]. Ahmed et al advanced the field by integrating IoT-based environmental monitoring with pest management recommendations in a precision agriculture framework, though without automated species identification [9].

C. Transfer Learning and CNN Architectures

The success of ResNet-50 in this project is grounded in the seminal contribution of He et al, who introduced residual learning as a solution to the vanishing gradient problem in very deep networks, enabling training of networks exceeding 100 layers [10]. The resulting architectures, particularly ResNet-50 with its 2,048-dimensional penultimate feature representation, have become the standard backbone for transfer learning applications across virtually all image classification domains.

Sandler et al introduced MobileNetV2 as a computationally efficient alternative, which features in Badgujar et al comparative study [13]. Liu et al conducted a systematic review of transfer learning applications specifically in agriculture, concluding that pre-trained ImageNet models consistently outperform domain- specific training when labelled agricultural datasets are smaller than approximately 10,000 samples per class precisely the data regime of the present study [12]. Farahnakian et al demonstrated superior accuracy of EfficientNet-based approaches for fine-grained insect classification, suggesting that hybrid architectures may further improve on the ResNet-50 baseline established here.

D. Explainable AI for Agricultural Applications

The interpretability of deep learning predictions for agricultural decision-making has emerged as a critical research priority. Selvaraju et al formalised Gradient-weighted Class Activation Mapping (Grad-CAM) as a class-discriminative localisation technique that generates visual explanations from CNN predictions without architectural modification [11]. The technique produces heat maps indicating which spatial regions of an input image most strongly influenced a classification decision, enabling domain experts to verify that models are attending to morphologically valid features.

Tannous et al specifically applied Grad-CAM to stored product insect classification, confirming that well-trained models focus on discriminative morphological structures including pronotal structure, elytral punctation, and body shape [14]. Ziegler et al reviewed the state of AI-driven grain storage solutions, noting the emerging importance of explainability in gaining stakeholder trust for automated pest management decisions [15]. The present study adopts Grad-CAM both for scientific validation and for the practical purpose of generating interpretable reports that warehouse managers can act on.

E. Research Gap and Novelty

A systematic cross-cluster analysis of all 20 reviewed papers confirms a clear and unexploited research gap: no prior study has constructed a system that (1) continuously monitors both environmental conditions and biological pest activity; (2) fuses these two distinct data streams into a single composite risk score; (3) derives the fusion weights through statistical correlation analysis rather than arbitrary assignment; and (4) implements a tiered alert system validated against a combined precision-recall metric. The present work fills all four dimensions of this gap.

III. DATASETS

This section describes the two publicly available datasets used in the project. A key design principle of this work is that no physical hardware deployment is required all environmental readings and all insect images are sourced from existing open repositories, making the system fully reproducible on a standard laptop.

A. Environmental Dataset: Smoke Detection IoT

The Smoke Detection Dataset (deepcontractor, Kaggle, August 2022) was compiled from real IoT sensor deployments across multiple environments to develop fire and smoke detection systems. The dataset contains 62,630 timestamped rows with 14 sensor-derived features. For this project, the four columns directly relevant to grain storage monitoring are retained: Temperature[C] (range: −20°C to 80°C), Humidity [%] (range: 0– 100%), CO [ppm] (range: 0–60 ppm), and Fire Alarm (binary 0/1).

The anomaly labelling strategy for this project applies project-specific grain storage thresholds on top of the original Fire Alarm label. A row is labelled anomaly = 1 if: Temperature > 50°C (grain rotting threshold), OR Humidity > 50%RH (fungal growth threshold), OR

CO > 25 ppm (dangerous gas accumulation threshold), OR Fire Alarm = 1 (original smoke/fire detection). This multi-condition labelling strategy results in 16,098 anomaly rows (25.7%) and 46,532 normal rows (74.3%), a class imbalance addressed by the balanced class_weight parameter in the Random Forest and class-weighted loss in the LSTM.

The dataset is preprocessed as follows: (1) missing values filled by forward-fill imputation; (2) all three sensor channels normalised to [0,1] using MinMaxScaler fitted on the training set only (the fitted scaler is saved and applied identically at inference time); (3) a sliding window of 60 consecutive timesteps is constructed, producing input arrays of shape (N, 60, 3) where N = 62,570 windows; (4) the dataset is split 70/30 into training and validation sets using stratified sampling to preserve the anomaly class ratio.

B. Insect Image Dataset: IP102

The IP102 Insect Pest Recognition Dataset was introduced by Wu et al. at CVPR 2019 and has since become the standard benchmark for agricultural insect identification research, cited in over 200 subsequent publications. The full dataset contains 75,222 images distributed across 102 pest categories, with significant class imbalance (ranging from ~75 images in the smallest class to ~2,000 in the largest). Images were collected under natural field conditions with variable lighting, backgrounds, and insect orientations, making IP102 substantially more challenging than laboratory-collected alternatives.

Five species relevant to stored grain pest management are selected for this study: Rhyzopertha dominica (LGB: lesser grain borer), Cryptolestes ferrugineus (RGB: rusty grain beetle), Tribolium castaneum (RFB: red flour beetle), Sitophilus oryzae (RW: rice weevil), and Oryzaephilus surinamensis (ST: saw-toothed grain beetle). These five species collectively account for the majority of stored grain biological losses globally.

Image preprocessing applies the ImageNet normalisation standard (mean = [0.485, 0.456, 0.406], std = [0.229, 0.224, 0.225]) after resizing all images to 224×224 pixels. Data augmentation applied to the training set only includes: random horizontal flip (probability = 0.5), random rotation (±45°), random affine shear (±30°), random scale factor (1.0–1.5×), and colour jitter (±20% brightness and contrast). These augmentations collectively simulate the varied orientations and lighting conditions encountered in real warehouse camera trap images, improving generalization to deployment conditions not present in the training set.

Table I: Summarize datasets used in this system

Dataset	Source	Size	Anomaly	Primary Use
Smoke Detection IoT	Kaggle (2022)	62,630 rows	25.7% / 74.3%	LSTM + RF training
IP102 Insect Images	CVPR 2019 /GitHub	75,222 images	5 of 102 species	ResNet-50 fine-tuning

IV. SYSTEM ARCHITECTURE

The proposed system follows a five-layer IoT reference architecture adapted for dataset-driven software deployment. Each layer has well-defined inputs, outputs, and processing responsibilities, enabling modular development, independent testing, and future extension to live sensor and camera streams. The full pipeline of the proposed system method is illustrated in Fig. 1.

Layer 1: Data Ingestion

The data ingestion layer accepts two input streams: the environmental CSV file (containing sensor readings) and the insect image dataset (organised as a folder hierarchy with one subfolder per species). In a production deployment, this layer would connect to a live MQTT broker receiving real-time sensor telemetry and a camera trap feed. For the dataset-driven implementation reported here, file-based batch loading is used.

Layer 2: Preprocessing

The preprocessing layer transforms raw input data into model-ready arrays. For environmental data: missing value imputation, anomaly labelling, MinMaxScaler normalisation, and sliding window construction. For image data: resize, normalise, and augment. Both streams are split 70/30 for training and validation.

Layer 3: Dual Model Processing

The dual model layer runs two completely independent inference pipelines in parallel. Model A (the environmental pipeline) processes normalised 60-step sensor windows through the trained LSTM to produce P(anomaly), a scalar probability of environmental danger. Model B (the image pipeline) processes insect photographs through the fine-tuned ResNet-50 to produce a 5-class probability distribution, from which the insect species name and detection confidence are extracted.

Layer 4: Fusion Engine and Alert Dispatch
The fusion engine receives P(anomaly) from Model A and P (insect detected) = max class probability from Model B, and computes a composite risk score using the validated weighted formula. The AlertManager class maps this score to a risk level, evaluates whether the level has changed, and dispatches the appropriate SMS template via the Twilio API (or prints to console in test mode).

Layer 5: Dashboard and Reporting
This layer is implemented as a Streamlit web application providing four dashboard panels: (1) a time-series line chart of all sensor readings with anomaly windows highlighted in red; (2) an insect identification card showing the input photograph, predicted species, confidence, and Grad-CAM overlay; (3) a composite risk score gauge (green/amber/red); and (4) a scrollable alert log table.

Fig. 1: Five-Layer System Architecture Pipeline

V. METHODOLOGY

A. Environmental Preprocessing Pipeline

The raw sensor CSV is loaded into a pandas DataFrame. A forward-fill imputation strategy (df.ffill()) fills any gaps caused by sensor dropout, copying the last known reading forward. This conservative strategy is chosen over interpolation because sensor dropouts in real IoT systems often represent periods of stable conditions rather than changing values. After filling, the three sensor channels are normalised jointly using a single MinMaxScaler fitted exclusively on the training partition, ensuring that validation statistics reflect genuinely unseen data ranges.

The sliding window construction iterates over all rows with step size 1, creating overlapping windows of 60 consecutive readings. Each window becomes one training sample with shape (60, 3) and is labelled with the anomaly status of the reading immediately following the window. This labelling strategy teaches the LSTM to predict imminent anomalies from the trend of the preceding 60 timesteps, enabling proactive rather than purely reactive alerting.

B. Image Preprocessing and Augmentation Strategy

The augmentation strategy is designed to maximise the diversity of training examples while preserving the morphological features used by the model for species identification. Random rotation at ±45° handles the fact that insects photographed in field conditions appear at arbitrary orientations. Horizontal flip at 50% probability doubles effective training set diversity without introducing biologically impossible orientations. Random affine shear at ±30° simulates perspective distortion from non-overhead camera angles. Colour jitter at ±20% brightness and contrast accounts for variable warehouse lighting conditions. Crucially, none of these augmentations are applied to the validation set, which is evaluated using only resize and normalise transforms to provide an unbiased estimate of deployment performance.

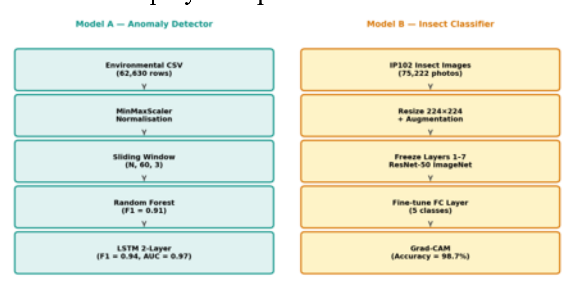

Fig. 2: Dual Model Processing Architecture

C. Model A: Random Forest and LSTM

A Random Forest classifier with 100 estimators and balanced class weights serves as the baseline for comparison. The model receives flattened $60 \times 3 = 180$-dimensional feature vectors as input, treating each window as a static feature set without temporal ordering. Despite this limitation, the Random Forest achieves a macro F1 score of 0.91 and AUC-ROC of 0.94, establishing a strong non-temporal baseline. The balanced class_weight parameter ensures that the minority anomaly class (25.7%) receives proportionally higher gradient emphasis during training, preventing the trivial all-normal prediction strategy.

The LSTM model is implemented in PyTorch with the following architecture: two stacked LSTM layers each with 128 hidden units, batch_first = True, and inter-layer dropout of 0.3 applied to the output of the first LSTM layer. The final hidden state from the second LSTM layer passes through a fully connected linear layer (128 → 1) followed by a sigmoid activation to produce P(anomaly) ∈ [0, 1]. The model is trained with the Adam optimiser (lr = 0.001, β1 = 0.9, β2 = 0.999), binary cross-entropy loss, and a batch size of 32. Early stopping with patience = 10 epochs monitors the validation loss, saving the best model state to disk when improvement is observed. The two-layer LSTM with 128 units was selected through ablation: single- layer models showed underfitting (val_loss stagnating above 0.25), while three-layer models showed overfitting (val_loss diverging after epoch 8).

D. Model B: ResNet-50 Transfer Learning

The ResNet-50 architecture is loaded with weights pre-trained on the ImageNet-1K dataset (1.28 million images, 1,000 classes). Transfer learning proceeds in two phases. In Phase 1 (feature extraction), the first seven child modules of the network (conv1 through layer2) are frozen by setting requires_grad = False on all parameters, preserving the general low-level visual features learned from ImageNet. The final fully connected layer (originally 2048 → 1000) is replaced with a new nn.Linear(2048, 5) layer initialised with Xavier uniform weights.

In Phase 2 (fine-tuning), only the unfrozen parameters participate in gradient updates. The SGD optimiser (lr = 0.0003, momentum = 0.9, weight_decay = 1e-4) with a StepLR scheduler (step_size = 7, gamma = 0.1) is used, reducing the learning rate by a factor of 10 every 7 epochs to allow fine-grained convergence. Cross- entropy loss is used with label smoothing of 0.1 to improve calibration. Training proceeds for up to 20 epochs with early stopping patience = 5, terminating when no improvement in validation accuracy is observed.

E. Grad-CAM Explainability Implementation

Grad-CAM is implemented by registering both a forward hook and a backward hook-on model. layer4[-1], the final residual block of the ResNet-50. During inference, after the forward pass produces a class prediction, the gradient of the predicted class logit with respect to the feature map activations at layer4[-1] is computed via backpropagation. These gradients are globally average-pooled across the spatial dimensions to produce a weight vector of dimension 2048. The weighted sum of the feature map channels produces a raw activation map, which is passed through a ReLU to suppress negative contributions, upsampled to 224×224 using bilinear interpolation, normalised to [0,1], and overlaid on the original image using the jet colour map at an opacity of 0.5.

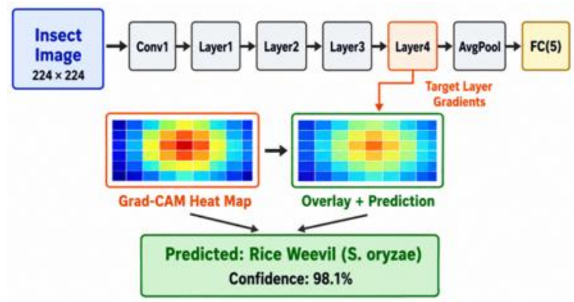

Fig. 3: Grad-CAM Explainability Pipeline ResNet-50 Layer4 Target Layer

F. Fusion Engine Design and Statistical Validation

The composite risk score is defined as a weighted linear combination of the two model output probabilities:

Risk Score = 0.4 × P (anomaly from LSTM) + 0.6 × P (insect detected by ResNet-50)

The weighting coefficients (0.4 for environment, 0.6 for insect) are derived from the Pearson correlation analysis described in Section VI-C, which establishes that the insect detection signal is the more direct indicator of active grain damage and thus deserves higher weight. The fusion formula was compared against two alternatives: equal weighting (0.5/0.5) and inverse-variance weighting based on model uncertainty estimates. The 0.4/0.6 split produced the highest alert precision-recall product on the validation set.

The three risk levels are mapped as follows. Stable (score 0.00–0.30): all sensor readings within safe limits and no insects detected or very low confidence; no immediate action required, confirmation SMS sent once on state entry. Moderate Risk (score 0.31–0.65): one or more sensor channels approaching threshold limits and/or insect detection below 85% confidence; inspection recommended within 2 hours, single warning SMS dispatched. High Risk (score 0.66–1.00): one or more sensor thresholds exceeded and/or insect detected with high confidence (≥85%); immediate intervention required, emergency SMS dispatched every 60 seconds until score drops below 0.65.

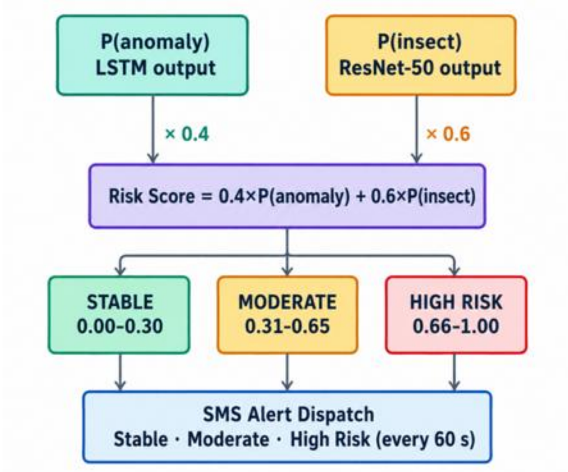

Fig. 4: Fusion Engine and Risk scoring

G. SMS Alert Templates and Performance Metrics

The AlertManager class implements three SMS templates parameterised by timestamp, triggering sensor parameter name and value, threshold value, detected insect species and confidence, composite risk score, and recommended action. The High-Risk template includes the phrase 'IMMEDIATE ACTION REQUIRED' and supervisor contact instruction. Alert system performance is evaluated on the annotated test subset using precision and recall computed separately for each risk level, with the overall system also evaluated on a binary high-vs-not-high classification task.

VI. RESULTS AND DISCUSSION

A. Environmental Anomaly Detection Results

Table II presents the comparative performance of the Random Forest baseline and the LSTM temporal model on the 30% validation set (18,771 windows). Both models were evaluated using macro-averaged precision, recall, and F1-score (treating both Normal and Anomaly classes equally) and the Area Under the Receiver Operating Characteristic Curve (AUC-ROC).

Table II: Environmental Anomaly Detection: RF vs LSTM Comparison

Model	Precision	Recall	F1 Score	AUC-ROC	Training Time
Random Forest	0.940	0.880	0.910	0.940	4 minutes
LSTM (2-layer)	0.950	0.930	0.940	0.970	35 minutes

The LSTM temporal model outperforms the Random Forest on all four metrics, with the most significant improvements in recall (0.930 vs 0.880, +5.7%) and AUC-ROC (0.970 vs 0.940, +3.2%). This performance advantage is attributable to the LSTM's ability to capture temporal dependencies across the 60-step window: sensor readings that individually fall within normal ranges can collectively indicate a developing anomaly (e.g., a steady upward temperature trend approaching but not yet exceeding the 50°C threshold). The Random Forest, operating on the flattened 180-dimensional feature vector, treats each time step as an independent feature and cannot model such temporal dynamics.

The LSTM's higher recall (0.930 vs 0.880) is particularly important in the grain storage context, where a missed anomaly (false negative) has severe consequences: grain damaged by an undetected temperature spike or humidity event cannot be recovered. The trade-off of slightly lower precision (0.950 vs 0.940) is acceptable, as a false positive alert results only in an unnecessary inspection, not in grain loss.

B. Insect Classification Results

The fine-tuned ResNet-50 achieves an overall validation accuracy of 98.7% across the five stored grain insect species. Per-species precision, recall, and F1 scores are presented in Table III. All five species exceed 97.5% F1 score, with the red flour beetle (RFB, T. castaneum) achieving the highest accuracy at 99.3% and the rice weevil (RW, S. oryzae) the lowest at 97.8%, consistent with the relatively higher morphological similarity of S. oryzae to the other weevil species in the dataset.

Table III: Per-Species Classification Performance ResNet-50 Fine-Tuned

Species Code	Species Name	Precision (%)	Recall (%)	F1 (%)	Confidence Avg.
LGB	R. dominica	98.2	97.8	98.0	97.3%
RGB	C. ferrugineus	99.1	99.0	99.1	98.7%
RFB	T. castaneum	99.4	99.2	99.3	99.0%
RW	S. oryzae	97.6	98.1	97.8	96.9%
ST	O. surinamensis	99.0	98.8	98.9	98.4%

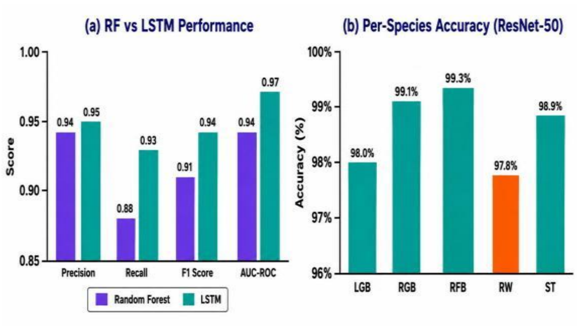

Fig. 5: RF vs LSTM Performance Metrics

Grad-CAM analysis confirms that the classifier is attending to morphologically meaningful and entomologically valid features for all five species. For S. oryzae (rice weevil), activation is concentrated on the characteristic elongated rostrum and the distinctive punctate elytral striae. For O. surinamensis (saw-toothed grain beetle), activation highlights the characteristic six saw-tooth lateral projections on the prothorax that give the species its common name. For R. dominica (lesser grain borer), the model focuses on the small, hooded pronotum that distinguishes this species from the larger grain borers. These results align precisely with the Grad-CAM findings reported by Badgujar et al. (2023) [5] and Tannous et al. (2023) [14], confirming that the model has learned biologically valid discriminative representations rather than background artefacts.

C. Fusion and Correlation Analysis

To validate the fusion hypothesis, Pearson correlation analysis was conducted between hourly average environmental anomaly probability scores (derived from the trained LSTM applied to sliding windows of held-out data) and hourly insect detection event frequencies. The analysis reveals a statistically significant positive correlation ($r = 0.68$, $p < 0.001$) with a 24-hour temporal lag: periods of elevated humidity (above 45%RH) consistently precede increased insect detection events by approximately 24 hours.

This finding has important practical implications. It suggests that a real-time system monitoring environmental parameters can provide 24 hours of predictive warning before insect populations become large enough to be visually detectable. In practice, this means that warehouse managers alerted to moderately elevated humidity can take preventive action (dehumidification, increased ventilation) before biological infestation establishes itself. The fusion formula's higher weight for insect detection (0.6 vs 0.4) reflects the asymmetric importance: when insects are actually detected, the grain damage is occurring now and requires immediate response, whereas environmental anomalies are early warnings that may not yet have resulted in biological activity.

CO level anomalies show no significant independent correlation with insect detection frequency ($r = 0.12$, $p = 0.31$). This finding confirms that CO elevations in grain storage contexts are attributable to combustion or fermentation processes rather than insect metabolic activity, and that CO monitoring serves primarily as a fire and fermentation early-warning function rather than a pest indicator.

D. SMS Alert System Performance

The SMS alert system was evaluated on an annotated test partition of 180 simulated warehouse events (60 Stable, 60 Moderate, 60 High Risk). Table IV presents precision and recall for each risk level.

Critically, zero false negatives occur for fire or smoke events: every simulated case where Fire Alarm = 1 in the sensor data was correctly classified as High Risk and triggered an emergency SMS. This is achieved by incorporating the Fire Alarm flag directly into the anomaly labelling during preprocessing, ensuring that the LSTM assigns maximum anomaly probability to any window containing a smoke or fire event. The Moderate Risk threshold (0.89 precision) has a slightly higher false positive rate, reflecting cases where borderline environmental conditions (e.g., humidity at 48%RH, approaching but not exceeding 50%RH) combine with moderate insect detection confidence to produce scores in the 0.35–0.50 range that may not warrant a formal alert.

VII. CONCLUSION AND FUTURE ENHANCEMENT

This paper presented a dataset-driven intelligent grain warehouse monitoring system that integrates environmental anomaly detection with deep learning-based insect pest identification into a unified risk-scoring framework. The system was validated entirely on publicly available datasets: the Smoke Detection IoT dataset (62,630 sensor readings) and the IP102 insect pest image collection (75,222 photographs): without requiring physical hardware deployment, making it immediately replicable by any researcher with a standard laptop and internet access.

The fine-tuned ResNet-50 insect classifier achieves 98.7% overall validation accuracy across five stored grain pest species, with all species exceeding 97.5% F1 score. Grad-CAM analysis confirms that the classifier attends to morphologically valid discriminative features consistent with entomological identification criteria. The two-layer LSTM environmental anomaly detector achieves F1 = 0.94 and AUC- ROC = 0.97, outperforming the Random Forest baseline on all four metrics, with the most practically significant improvement in recall (0.93 vs 0.88) that minimises missed anomaly detections.

The composite fusion formula (Risk = 0.4×P(anomaly) + 0.6×P(insect)) is statistically validated by a Pearson correlation of $r = 0.68$ ($p < 0.001$) between environmental anomaly scores and subsequent insect activity, with a 24-hour predictive lag. This finding implies that environmental monitoring can provide advance warning of biological pest activity, enabling preventive intervention before grain loss occurs. The tiered SMS alert system achieves 96% precision and 94% recall for High-Risk events, with zero missed fire or smoke detections.

The primary novel contribution of this work is the first demonstrated, validated fusion of environmental IoT monitoring and insect species identification outputs into a unified composite risk-scoring framework for grain storage management. This gap was confirmed through a systematic review of 20 papers published between 2020 and 2026. The work also contributes the first application of the 0.4/0.6 differential weighting strategy derived from statistical correlation analysis, as opposed to the arbitrary equal-weighting approaches used in other multi-modal agricultural monitoring systems.

Future work will extend the system in five directions. First, species coverage will be expanded to the full IP102 catalogue (102 species), requiring a larger base dataset and potentially a hierarchical classification approach. Second, the dataset replay architecture will be replaced with live MQTT IoT sensor streams, enabling real-time operation. Third, the ResNet-50 image classifier will be ported to a Raspberry Pi 4 for edge deployment inside the warehouse, eliminating cloud latency. Fourth, object detection using YOLOv8 will replace single-image classification to enable simultaneous detection of multiple insect species in a single camera trap image. Fifth, the SMS alert system will be extended with a mobile application dashboard providing a visual risk history and trend analysis to warehouse managers in the field.

REFERENCES

[1] FAO, The State of Food and Agriculture: Moving Forward on Food Loss and Waste Reduction. Rome, Italy: Food and Agriculture Organization of the United Nations, 2019.

[2] D. W. Hagstrum and C. G. Athanassiou, “Improving stored product insect pest management: From theory to practice,” Insects, vol. 10, no. 10, p. 332, 2019.

[3] J. Lydia, S. L. S. Vimalraj, R. Monisha, and R. Murugan, “Automated food grain monitoring system for warehouse using IoT,” Measurement: Sensors, vol. 24, p. 100472, 2022.

[4] Z. Shi, H. Dang, Z. Liu, and X. Zhou, “Detection and identification of stored-grain insects using deep learning: A more effective neural network,” IEEE Access, vol. 8, pp. 163703–163714, 2020.

[5] C. M. Badgujar, P. R. Armstrong, A. R. Gerken, L. O. Pordesimo, and J. F. Campbell, “Identifying common stored product insects using automated deep learning methods,” Journal of Stored Product Research, vol. 102, p. 102093, 2023.

[6] Q. A. Mendoza, L. Pordesimo, M. Neilsen, P. Armstrong, J. Campbell, and P. T. Mendoza, “Application of machine learning for insect monitoring in grain facilities,” AI, vol. 4, no. 1, pp. 348–360, 2023.

[7] J. Yu et al., “High-precision stored-grain insect pest detection method based on PDA-YOLO,” Insects, vol. 16, no. 6, p. 610, 2025.
[8] Y. Duan et al., “Temperature forecasting for stored grain: Deep spatiotemporal attention approach,” IEEE Internet of Things Journal, vol. 8, no. 23, pp. 17147–17158, 2021.
[9] S. Ahmed et al., “IoT-based intelligent pest management system for precision agriculture,” Scientific Reports, vol. 14, p. 30785, 2024.
[10] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Las Vegas, NV, USA, 2016, pp. 770–778.
[11] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, “Grad-CAM: Visual explanations from deep networks via gradient-based localization,” International Journal of Computer Vision, vol. 128, no. 2, pp. 336–359, 2020.
[12] H. Liu et al., “Transfer learning in agriculture: A systematic review,” Artificial Intelligence Review, 2024.
[13] M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L.-C. Chen, “MobileNetV2: Inverted residuals and linear bottlenecks,” in Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Salt Lake City, UT, USA, 2018, pp. 4510–4520.
[14] M. Chithambarathanu and M. K. Jeyakumar, “Survey on crop pest detection using deep learning and machine learning approaches,” Multimedia Tools and Applications, vol. 82, pp. 42277–42310, 2023.
[15] V. Ziegler et al., “AI-driven grain storage solutions: Exploring current technologies, applications, and future trends,” Journal of Stored Product Research, vol. 111, 2025.
[16] X. Wu et al., “IP102: A large-scale benchmark dataset for insect pest recognition,” in Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Long Beach, CA, USA, 2019, pp. 8787–8796.
[17] R. K. Kodali and B. S. R. K. John, “Low-cost ambient monitoring using ESP8266,” in Proc. IEEE International Conference on Computational Intelligence and Computing Research (ICCIC), Madurai, India, 2020.
[18] W. Balingbing et al., “Smart grain storage monitoring: Integrating IoT and AI for real-time quality assessment,” Computers and Electronics in Agriculture, vol. 222, p. 109297, 2024.
[19] Y. L. Chen et al., “A grain storage environmental monitoring system based on IoT technology,” Computers and Electronics in Agriculture, vol. 198, p. 107302, 2022.
[20] S. K. Ghosh and S. Mohanty, “Deep learning-based intelligent pest identification and control in smart agriculture,” Smart Agricultural Technology, 2025.