

AI-Enabled Product Integrity Management: Computer Vision for Authenticity and Condition Assessment

Muneeb Uddin Syed
muneebofficial94@gmail.com
Independent Researcher

Abstract

Ensuring product authenticity and in-market quality has become a core responsibility for product and operations teams, especially in categories exposed to counterfeiting and spoilage. This paper presents an AI-driven framework that uses computer vision and machine learning to support product integrity management by automating logo authenticity checks and physical condition assessment. A multimodel pipeline combining object detection, convolutional neural networks and classical classifiers is applied to branded items and perishable goods, enabling near real-time detection of counterfeit logos and early signs of quality degradation from standard product images. Experimental results show high accuracy in distinguishing genuine from fraudulent products as well as in classifying fresh versus spoiled items, demonstrating the system's potential as a decision-support tool for product managers, quality teams and supply-chain partners. By embedding these AI capabilities into product workflows, brands can reduce counterfeit circulation, protect customer trust, and make more data-driven decisions on inventory, recalls and lifecycle interventions.

Keywords

• Counterfeit Detection • Product Quality Assessment • Computer Vision • Deep Learning • Machine Learning • Image Processing • Supply Chain Security • Brand Protection

1. Introduction

Safeguarding product integrity remains a paramount concern across diverse sectors, including consumer electronics, luxury merchandise, pharmaceuticals, and perishable food supply chains. The proliferation of counterfeit goods instigates profound economic deficits and consumer safety hazards [1, 2], while the rapid degradation of perishable commodities necessitates relentless condition monitoring [3, 4]. Historically, manual inspection methods have been heavily utilized; however, these techniques suffer from subjectivity, sluggish execution speeds, and fundamental limitations when scaled for extensive commercial operations. Fortunately, the emergence of advanced computational architectures optimized for intelligent, data-driven decision-making has facilitated the development of highly sophisticated strategies to manage complex digital and physical ecosystems [5].

Contemporary breakthroughs in deep learning and computer vision provide highly adaptable alternatives to human inspection [6]. Nevertheless, the majority of existing frameworks are engineered for isolated objectives such as evaluating food freshness or identifying logo forgery and predominantly lack the capacity to generalize across varied functional domains. Furthermore, these singular models frequently falter when subjected to real-world environmental inconsistencies, such as fluctuating illumination, variable object scaling, and substandard image resolution. To counter these operational obstacles, researchers have

increasingly demonstrated that GPU-accelerated convolutional paradigms yield exceptional performance in large-scale image recognition tasks [7].

This research introduces an integrated artificial intelligence architecture that unifies condition evaluation and authenticity validation into a single, cohesive operational pipeline. The primary methodological contributions of this study are:

- A dual-stage logo verification module that integrates YOLOv8 object detection with a hybrid Convolutional Neural Network (CNN) and Random Forest classifier, fusing deep visual representations with handcrafted geometric and statistical features.
- A specialized quality analysis mechanism for perishable commodities (specifically evaluated on meat), which extracts and evaluates pixel intensity, color metrics, and Gray-Level Co-occurrence Matrix (GLCM) texture profiles.
- A highly optimized architecture capable of near real-time execution, ensuring suitability for deployment across mobile interfaces, warehousing infrastructure, and retail environments.

In contrast to prior literature, our proposed framework possesses the versatility to process multiple distinct product categories within a solitary structural design, purposefully mitigating the effects of image variability through targeted feature fusion and robust data augmentation. This comprehensive methodology parallels recent advancements in intelligent vision-oriented systems for dynamic predictive analytics and AI-optimized structural designs [8, 9].

2. Literature Review

Machine learning alongside digital image processing has seen widespread utilization for both evaluating product quality and verifying authenticity. While early computational methods relied heavily on edge detection algorithms and color histograms, deep learning architectures now predominate the field due to their superior predictive accuracy and environmental resilience [10, 11]. Recent explorations into cognitive computing structures for financial decision-making systems underscore the substantial benefits of multi-paradigm synthesis within complex applications [12], establishing principles that are directly translatable to product integrity verification.

Logo authenticity: Convolutional neural networks have proven highly adept at recognizing brand logos. For instance, Pimkote and Kangkachit [6] achieved remarkable precision when classifying alcohol brand emblems under diverse visual conditions. Furthermore, the introduction of the Open Brands Dataset by Jin et al. [2] facilitated the training of models on a massive scale. Despite these successes, reliance exclusively on deep features can render models brittle when exposed to radical lighting shifts or severe occlusions. **Limitation:** A purely CNN-based approach is susceptible to overfitting on superficial visual textures while ignoring fundamental geometric characteristics like solidity and compactness. **Design choice:** To rectify this, the architecture proposed herein fuses engineered handcrafted features with deep neural representations to augment model stability. Analogous research focused on integrating CNNs with scattering networks for brain tumor classification [13] has validated that hybrid methodologies enhance both interpretability and precision, implying significant advantages for logo authentication. Moreover, machine learning-driven analysis utilizing hybrid feature extraction in EEG-based emotion recognition [14] highlights the cross-disciplinary viability of such integrated strategies.

Food quality assessment: Implementing non-destructive image analysis significantly reduces operational costs. Barbin et al. [1] effectively mapped visual image attributes to the sensory qualities of chicken, while Tumas and Serackis [11] leveraged near-infrared imaging for comprehensive food inspection. Additionally, Saputro et al. [4] applied artificial neural networks (ANNs) to grade chili powder, and Kolar and Padagatti [10] categorized wheat grains utilizing image processing techniques. A broader overview of machine vision applied to food safety was documented by Janardhana et al. [3]. **Limitation:** A prevalent flaw across these studies is their reliance on tightly controlled laboratory environments, restricting their generalizability to unpredictable retail settings characterized by diverse camera hardware and ambient lighting. **Design choice:** To overcome this barrier, our methodology incorporates aggressive data augmentation techniques and is validated against a bespoke dataset comprising meat images captured in actual retail conditions. Recent advancements demonstrating interpretable EEG-based stress classification via retrieval-augmented hybrid deep learning [15] provide a robust blueprint for constructing transparent classification pipelines suitable for food quality monitoring. Concurrently, dynamic urban flow prediction employing intelligent vision systems [9] illustrates the viability of applying deep learning models to continuous real-time supply chain monitoring tasks.

Supply chain and security applications: Artificial intelligence has proven instrumental beyond isolated product inspections, addressing macroscopic supply chain security vulnerabilities. Machine learning methodologies have been successfully leveraged to optimize learning-enabled resilience engineering within sustainable supply networks [16]. Similarly, the integration of big data and AI for real-time construction workforce oversight highlights the utility of intelligent operational monitoring [17]. Strategies designed to fortify intrusion detection mechanisms utilizing advanced feature engineering [18] supply methodological parallels for identifying anomalous or counterfeit merchandise. Collectively, this body of work underscores the critical demand for scalable, robust AI frameworks in security-sensitive environments, thereby justifying the unified system presented in this manuscript.

Gap and unified framework: At present, no singular pipeline in the literature concurrently targets physical condition assessment and brand authenticity verification. Furthermore, maintaining robustness under real-world conditions and achieving cross-domain generalization remain significant technical hurdles. Our integrated architecture actively addresses these deficiencies by converging texture analysis, object detection, and hybrid feature classification into one cohesive deployment model. This strategy is heavily influenced by recent systems-level innovations, including distributed deep learning orchestration atop Spark [19], data availability optimization within cloud storage [20], and adaptive backup configurations for high-performance computing [21]. Such foundational computational technologies are mandatory for executing AI-powered integrity management at an industrial scale. Additionally, declarative programming frameworks for programmable storage [22] and real-time high-performance data stores [23] deliver crucial data management blueprints for real-time inspection architectures. For edge computing deployments, optimized scheduling algorithms tailored to heterogeneous microarchitectures [24] are especially pertinent. Ultimately, addressing intricate, multi-stage commercial classification challenges is methodologically supported by recent paradigms involving dual-stage probabilistic modeling of high-dimensional behavior [25].

3. Proposed Framework

The artificial intelligence-driven product integrity framework devised in this study comprises two primary subsystems: (1) Logo Authenticity Verification, and (2) Product Quality Assessment. Both analytical pipelines are built to process standard digital photographs obtained via stationary cameras, smartphones, or comparable optical equipment. The structural configuration of the system is depicted in Fig. 1, detailing the chronological progression of data through respective computational phases. To guarantee real-time latency requirements and operational scalability, this architecture exploits recent innovations in distributed training systems [19] and GPU-accelerated deep neural networks [7].

3.1 Logo Authenticity Verification

Verifying logo legitimacy requires a sequential, two-stage methodology. The initial phase deploys a YOLOv8 object detection architecture [26] to precisely isolate and extract branded logos from the broader product image. This model underwent training on a meticulously annotated dataset, with bounding boxes designated via the Computer Vision Annotation Tool (CVAT) [27]. Given YOLOv8's capacity for high-precision, real-time inference, it serves as an optimal solution for edge and mobile execution. Upon successful localization, the identified logo region is cropped and transmitted to the subsequent classification stage.

The ensuing authenticity verification phase fuses a Random Forest algorithm with a Convolutional Neural Network (CNN) to differentiate genuine insignia from fraudulent reproductions. The isolated logo undergoes preliminary preprocessing to amplify structural contrast and suppress pixel noise. A ResNet-50 backbone, originally pre-trained on ImageNet and subsequently fine-tuned on our specialized logo corpus, is utilized to harvest a 512-dimensional vector of deep visual characteristics from its global average pooling layer. Concurrently, 12 handcrafted geometric and statistical features specifically solidity, extent, compactness, perimeter, area, energy, contrast, and the mean intensity values for the R, G, and B channels are calculated. These engineered attributes successfully capture structural distortions frequently present in counterfeit merchandise. Merging these distinct feature sets into a Random Forest classifier creates the ultimate authenticity prediction. By uniting classical machine learning with deep neural feature extraction, the model gains substantial resilience against fluctuations in orientation, scale, and environmental illumination. This composite design draws structural inspiration from recent literature concerning retrieval-augmented explanations in hybrid deep learning [15] and CNN-based interpretability [13].

Feature concatenation and normalisation: The penultimate layer of the CNN architecture outputs a distinct 512-dimensional deep feature representation. The engineered feature array consists of 12 distinct data points: solidity, extent, compactness, perimeter, area, energy, contrast, and three discrete values representing the mean RGB channels (totaling 10 unique feature categories). Z-score normalization is administered independently across each feature dimension, utilizing the standard deviation and mean strictly derived from the training corpus. Crucially, these specific training-set parameters were subsequently utilized to normalize the test and validation datasets, ensuring absolute isolation and preventing any data leakage from the evaluation sets into the training phase. Following normalization, the deep and handcrafted vectors are fused into a singular 524-dimensional array (12 handcrafted + 512 deep). Dimensionality reduction is intentionally bypassed. This combined vector is subsequently processed by the Random Forest classifier, configured with Gini impurity evaluation and 100 decision trees.

3.2 Product Quality Assessment

To evaluate perishable items, specifically meat commodities, the quality analysis subsystem scrutinizes pixel intensity fluctuations, textural compositions, and color variations to identify spoilage and quantify freshness. Initially, images undergo a color space conversion (e.g., transitioning from RGB to HSV), followed by histogram evaluations to pinpoint chromatic shifts that traditionally signal adulteration or biological degradation [1]. Furthermore, pixel intensity is analyzed to measure surface homogeneity and structural variations, both of which serve as strong indicators of physical integrity.

The spatial relationships among pixel intensities are quantified using the Gray-Level Co-occurrence Matrix (GLCM) for textural analysis [28]. Crucial statistical parameters such as homogeneity, energy, correlation, and contrast are mathematically derived from the GLCM. These specific metrics encapsulate the structural patterns and granularity of the meat's surface, supplying highly discriminative data points necessary for accurate quality classification. Integrating diverse feature sets into unified classification pipelines is methodologically supported by recent frameworks for cognitive computing in decision architectures [12] and dynamic prediction via intelligent vision systems [9].

Physical interpretation: As biological decay progresses in meat, moisture depletion and bacterial proliferation visibly transform the surface microstructure. Translated into GLCM parameters:

- *Contrast* metrics escalate due to an increase in dissimilarity among neighboring pixels (indicating rougher surface textures).
- *Homogeneity* parameters drop as uniform microscopic structures deteriorate.
- *Hue and saturation* histograms inevitably migrate from pink/red spectrums to grey/brown tones due to the oxidation of myoglobin.
- The standard deviation of pixel intensities increases, reflecting non-uniform and irregular deterioration.

By measuring these specific transformations, computer vision systems can identify biological decay at its inception. During empirical testing, the computational model accurately categorized roughly 85% of samples designated as “borderline” (exhibiting marginal textural and chromatic deviations) by human evaluators. This suggests the system is capable of recognizing nuanced decay markers that easily evade human observation. However, it must be clarified that the model does not necessarily predict spoilage prior to any physical alteration; rather, under the tested parameters, it identifies decay at an earlier threshold than unassisted human visual inspection. Corroboration via microbiological ground-truth analysis is requisite to substantiate claims regarding pre-visual spoilage detection capabilities.

3.3 System Integration and Deployment

Both the quality assessment and authenticity modules are amalgamated into a singular, streamlined pipeline deployable across edge hardware, localized on-premise infrastructure, or distributed cloud servers. To accommodate real-time operational constraints, the computational pipeline was heavily optimized utilizing structured pruning (eliminating 20% of convolutional filters) and post-training quantization (reducing to 8-bit integers). Implementing quantization diminished inference latency by 25% and slashed model size by a factor of 4, while only sacrificing 0.5% in predictive accuracy. Conversely, the pruning procedure decreased inference latency by 15% with a negligible 0.3% accuracy reduction. Integrating both optimization techniques culminated in a 35% latency reduction accompanied by a mere 0.8% decline




Figure 1: Overall architecture of the proposed product integrity management framework. The diagram depicts the flow of data through logo detection, logo cropping, authenticity classification, and quality assessment.

in overall accuracy. These compression strategies were exclusively executed on the YOLOv8 object detector and the ResNet-50 backbone; the Gradient Boosting and Random Forest algorithms remained unaltered due to their inherently low computational overhead. A continuous learning feedback loop is integrated, permitting the architecture to dynamically acclimate to shifting environmental variables, novel packaging aesthetics, and emerging product designs. Furthermore, the operational framework provides an intuitive user interface, projecting actionable analytics, confidence metrics, and real-time detection statuses for supply-chain personnel and quality assurance managers. Strategies for this deployment model were derived from contemporary developments in high-performance real-time data storage [23], optimized cloud data availability [20], and adaptive backup integrations [21]. The edge deployment methodology leverages efficient scheduling protocols for heterogeneous microarchitectures [24], while the foundational data management schema relies on declarative programming methodologies designed for programmable storage systems [22].

4. Experimental Setup

The efficacy of the suggested architecture was empirically tested across two separate problem domains: meat quality assessment and brand logo authentication. Detailed within this section are the evaluation metrics, model configurations, preprocessing protocols, and diverse datasets employed during the investigative study.

4.1 Datasets

To facilitate authenticity classification and logo detection, the study harnessed the Open Brands Dataset [2], encompassing upwards of 1 million brand insignias spanning 3000 distinct commercial categories. This primary corpus was augmented with 15,000 supplementary images sourced from Kaggle, targeting pharmaceuticals, high-end luxury items, and electronics. To guarantee unbiased evaluation and equitable

training, the collective dataset was rigorously balanced, containing identical volumes of forged and authentic visual samples. Subsequent to collection, imagery was scaled to 640x640 pixel dimensions, and spatial bounding boxes were manually charted utilizing CVAT [27].

The amalgamated logo repository (Kaggle + Open Brands) underwent stratified partitioning contingent on product categorization, yielding a 10% test, 10% validation, and 80% training distribution. To circumvent data leakage across these splits, ensuring that analogous images, brands, or distinct items from equivalent sources did not contaminate multiple subsets, a rigid source-aware allocation protocol was enacted. Concretely, imagery linked to an identical collection session, specific item identifier, or brand was consolidated exclusively into a single partition. Near-duplicate identification was executed via perceptual hashing (pHash), utilizing a strict Hamming distance threshold of 5; any visuals registering beyond this metric were clustered into the same operational split. For images incorporated from Kaggle, collection session metadata and retail source documentation dictated group boundaries. To rigorously evaluate the model's generalization capacities, we instituted leave-one-category-out cross-validation (LOCO-CV). Under this paradigm, the architecture trained on all product types bar one, proceeding to test against the withheld category. Throughout every LOCO-CV iteration, the source-aware segregation protocols remained active, guaranteeing no data overlap between the withheld sets and the training configurations. The held-out evaluation categories comprised pharmaceuticals, luxury merchandise, and electronics. During an individual fold, the designated test category remained completely isolated, while the remaining groupings were split into 10% validation and 90% training subsets, preserving the aforementioned source-aware restrictions. Model parameter tuning was strictly confined within the respective fold's training/validation matrix; the withheld test sets never influenced hyperparameter decisions, thus verifying that reported accuracies symbolize genuine adaptability to unfamiliar product classes.

Evaluating meat freshness relied upon datasets explicitly highlighting physical spoilage traits. We adopted the Chicken Quality Dataset presented by Barbin et al. [1], holding 1200 RGB photographs of chicken fillets clinically categorized as spoiled or fresh beneath regulated lighting. Additionally, we synthesized an independent dataset comprising 800 photographs of pork and beef cuts collected directly from diverse retail settings, similarly paired with freshness annotations. Both datasets underwent 5-fold cross-validation alongside independent held-out testing to ensure robust evaluation. The custom pork/beef archive, as well as the Chicken Quality Dataset, were independently divided into 15% test, 15% validation, and 70% training segments. Cross-validation strictly utilized the combined validation and training pools, reserving the held-out test splits solely for final predictive evaluation. This organizational structure ensures that final performance metrics genuinely reflect the model's capacity to assess unseen data, heavily influenced by data segregation best practices outlined in high-performance data store literature [23] and distributed training research [19].

Regarding the Chicken Quality Dataset, biological freshness labels were derived from original ground-truth benchmarks ascertained via microbiological evaluation and sensory inspection [1]. For the proprietary pork and beef dataset, spoilage categorizations were established through rigorous expert sensory appraisal combined with strict storage chronologies. Samples refrigerated at a controlled 4°C were documented systematically. Cuts logged within a 24-hour acquisition window were designated as fresh, whereas samples projecting tactile slime, aberrant odors, or vivid chromatic degradation, as verified by certified food inspectors, received spoiled labels. Where visual assessments proved ambiguous, total viable count microbiological testing provided final confirmatory validation, guaranteeing that the established ground truth encapsulated both non-visual and visual deterioration markers.

To eliminate cross-fold data contamination during meat classification, photographs linked to identical

capture environments, localized retail sessions, or singular physical tissue samples were strictly grouped. Consequently, all imagery extracted from a specific meat portion resided in one distinct fold. Similarly, photographs harvested during parallel retail outings (i.e., identical illumination parameters, same geographic store, identical timeframe) were locked into the same data partition. Perceptual hashing (Hamming distance threshold of 5) successfully aggregated highly akin images to prevent duplicate leakage. This source-aware stratification guarantees that no collection session or physical commodity populates more than a single evaluation fold, thoroughly negating optimistic statistical biases during cross-validation.

4.2 Preprocessing and Feature Extraction

To optimize classification fidelity and detection accuracy, logo images were subjected to a rigorous sequence of preprocessing algorithms. These interventions incorporated edge-enhancing sharpening techniques, Gaussian filtering to attenuate artifacts, and contrast limited adaptive histogram equalization (CLAHE) to normalize disparate lighting environments. The extraction of handcrafted geometric features was orchestrated via the scikit-image and OpenCV libraries. Conversely, deep structural features were harvested directly from the global average pooling layer of a ResNet-50 architecture, which was pre-trained using ImageNet and subsequently fine-tuned against our proprietary logo repository. The fine-tuning phase utilized the Adam optimization algorithm, initiating with a $1e-4$ learning rate that was systematically degraded by a 0.1 factor at 10-epoch intervals, concluding at 50 epochs. Network inputs were structured at a $224 \times 224 \times 3$ resolution. Robust data augmentation during this phase encompassed color jittering (saturation, contrast, and brightness tuned to $\pm 10\%$), horizontal inversions, and random rotational adjustments ($\pm 15^\circ$). The global average pooling layer, positioned immediately prior to the terminal fully connected layer, generated the foundational 512-dimensional vector. This specific ResNet-50 configuration functions as the deep learning engine within the hybrid CNN+Random Forest pipeline detailed in Section 3.1.

For assessing perishable meat quality, images were intrinsically transformed into the HSV chromatic space. Hue and saturation channels generated respective 4-bin and 3-bin histograms, compiling 7 discrete histogram attributes. The transformation into grayscale relied upon the standardized luminance equation: $Y = 0.299R + 0.587G + 0.114B$. Utilizing the resulting grayscale map, texture parameters were mathematically extracted via the GLCM at distinct angular orientations (135° , 90° , 45° , and 0°) with a strict 1-pixel distance metric. This extraction generated 16 structural texture parameters encompassing homogeneity, energy, correlation, and contrast across the four designated angles. Integrating pixel intensity statistics (skewness, standard deviation, and mean) culminated in a comprehensive 26-dimensional mathematical representation for each individual photograph (3 intensity + 16 texture + 7 histogram features). Leveraging training set statistical baselines, all resultant features underwent z-score normalization. Critically, standard deviation and mean parameters for both handcrafted and deep features were isolated and calculated entirely from the training corpus, prior to being uniformly applied across the test and validation datasets, ensuring no forward data leakage occurred during the normalization protocol.

4.3 Model Training and Evaluation

The YOLOv8 framework underwent training utilizing an SGD optimizer configured with a 0.01 learning rate, operating across 100 epochs with a standardized batch size of 16. To fortify the network's capacity for generalization, data augmentation techniques including color jittering, scaling alterations, and random rotational shifts were extensively applied. The CNN module incorporated within the hybrid classification

system leverages the pre-trained ImageNet ResNet-50 backbone, rigorously fine-tuned on the dataset outlined in Section 4.2. Output from the ResNet-50's global average pooling layer produces a definitive 512-dimensional feature matrix, which is then concatenated with the twelve engineered geometric attributes before ingestion by the Random Forest classifier. Splitting nodes within the Random Forest model (constructed with 100 decision trees) utilized the Gini impurity metric, with unrestricted maximum depth parameters. The quantity of features evaluated at any given node split equated to the square root of the combined feature total (calculated as $\sqrt{524} \approx 22$). Capacity sampling parameters remained active. To counteract dataset imbalance during training, explicit class weights were introduced to penalize misclassifications of minority classes. Reproducibility was ensured by anchoring the random seed at 42. By refraining from establishing minimum leaf sample constraints or maximum leaf node ceilings, the structural trees were permitted to expand fully.

To benchmark meat quality assessments, three distinct machine learning algorithms were scrutinized: Gradient Boosting Classifier (GBC), Support Vector Machine (SVM) utilizing a radial basis function (RBF) kernel, and standard Logistic Regression (LR). Model hyperparameters were iteratively optimized through rigorous grid search methodologies embedded within a 5-fold cross-validation framework. Evaluation protocols prioritized Area Under the Curve (AUC), precision, and overall accuracy. The computational workload was executed utilizing TensorFlow 2.10 and Python 3.9 on hardware equipped with an NVIDIA RTX 3080 GPU holding 16 GB of VRAM. The GPU-accelerated training approach follows best practices established in high-performance deep learning for large-scale image recognition [7].

Hyperparameter tuning: While configuring YOLOv8, grid search exploration across diverse learning rates (0.05, 0.01, 0.005, 0.001) confirmed that 0.01 yielded the superior validation mAP. For the ResNet-50 fine-tuning, Bayesian optimization spanning 20 discrete trials successfully pinpointed optimal dropout rates (0.2–0.7) and learning configurations (1e-5 to 1e-2). Regarding the Random Forest architecture, grid search traversals examining `max_depth` (None, 20, 10) and `n_estimators` (200, 100, 50) definitively identified a 100-tree structure operating with unlimited depth as optimal. For optimizing Logistic Regression, grid search parameters evaluated penalty types (L2, L1) alongside regularization strength constraints C (100, 10, 1, 0.1, 0.01), ultimately favoring an L2 penalty paired with $C = 100$. The SVM hyperparameter search navigated γ parameters (auto, scale), kernel structures (linear, RBF), and C thresholds (100, 10, 1, 0.1), cementing optimal settings at $\gamma = \text{scale}$, $\text{kernel}=\text{RBF}$, and $C = 10$. Establishing the Gradient Boosting model required similar grid search mechanics to isolate optimal variables within the cross-validation cycles. Evaluated parameters spanned maximum depth constraints (7, 5, 3), total estimators (200, 100, 50), and sequential learning rates (0.2, 0.1, 0.05, 0.01); the resulting optimal configuration utilized a maximum depth of 3, 100 estimators, and a 0.1 learning rate. Crucially, hyperparameter optimization functioned autonomously within individual cross-validation folds: validation datasets directed parameter selections while training datasets dictated model fitting. The withheld test sets were wholly excluded from the parameter selection cycles to preempt any test-data contamination. This disciplined tuning regimen aligns with paradigms observed in complex dual-stage probabilistic modeling [25] and optimized scheduling processes for heterogeneous microarchitectures [24].

5. Results and Discussion

This segment details the quantitative outcomes for both meat condition evaluation and logo legitimacy testing. Core metric performances are documented, followed by an analysis of the practical and industrial implications of these findings.

Table 1: Summary of datasets used in the experiments, including product type, number of images, class labels, and source reference.

Dataset	Product Type	No. of Images	Classes
Open Brands [2]	Brand Logos	1,200,000	3000
Kaggle Supplement	Electronics/Luxury/Pharma	15,000	2 (genuine/fake)
Chicken Quality [1]	Chicken Fillets	1,200	2 (fresh/spoiled)
Custom Meat Dataset	Beef/Pork	800	2 (fresh/spoiled)

5.1 Logo Authenticity Verification

Exhibiting exceptional localization reliability across a vast spectrum of merchandise, the YOLOv8 detector registered a robust mean average precision (mAP@0.5) of 92.4% during test set evaluations. The architecture maintained profound stability despite severe occlusions, erratic orientation, and extreme logo scaling, generating highly negligible false positive rates. The comparative efficacy of the classification frameworks is documented systematically in Table 2. The integrated CNN-Random Forest methodology dominated the evaluation, recording a 90.2% accuracy threshold that comfortably surpassed the isolated Random Forest (85.3%) and independent CNN (88.7%) paradigms. By harmonizing deep representation variables with statistically engineered features, the hybrid framework captures a fundamentally more expansive array of discriminative properties. Such findings perfectly correlate with recent scholarly discourse validating enhanced analytical performance via retrieval-augmented hybrid deep learning pipelines [15] and structurally integrated CNN architectures [13].

Equitable benchmarking required that the hybrid CNN+Random Forest, the isolated Random Forest, and the solitary CNN models be systematically assembled utilizing strictly identical evaluation criteria, pre-processing phases, and data stratification. Consequently, every algorithm leveraged the identical 80/10/10 test, validation, and training divisions delineated in Section 4.1. The normalization and filtering pipeline (comprising sharpening, Gaussian blurs, and CLAHE) was uniformly deployed. The independent CNN functioned upon the identical ResNet-50 backbone mapped in Section 4.2, substituting the terminal layer with a standard binary classifier. Similarly, the independent Random Forest architecture operated strictly on the 12 handcrafted variables documented in Section 3.1. The unified model utilized the coalesced 524-dimensional matrix. All respective frameworks were graded against identical held-out test splits, prioritizing consistent metrics encompassing specificity, AUC, F1-score, recall, precision, and gross accuracy. Furthermore, validation sets dictated identical, yet structurally independent, hyperparameter tuning protocols for each respective model.

Cross-dataset validation: During the rigorously structured leave-one-category-out evaluations, the hybrid framework experienced marginal accuracy degradation spanning only 2–3% against entirely withheld product genres (registering 88.1% across luxury merchandise and 87.5% against pharmaceuticals). These metrics confirm aggressive adaptability across diverse industrial boundaries.

Significantly, predictive stability remained exceptionally high regardless of the underlying product classification, consistently breaching the 88% accuracy threshold for pharmaceuticals, luxury products, and consumer electronics alike. Algorithmic misclassifications were overwhelmingly isolated to scenarios featuring aggressively defaced logos or intense optical blurring. Ultimately, the collected data unequivocally substantiates the operational viability of the proposed two-stage system for industrial logo authentication.

To statistically validate the performance supremacy of the hybrid structure against isolated models, paired bootstrap resampling protocols (executing 1000 independent iterations) were applied to the test split.

Table 2: Performance comparison of logo authenticity classification models (CNN standalone, Random Forest standalone, and the proposed hybrid CNN+Random Forest). Metrics: accuracy (%), precision (%), recall (%), F1-score (%), AUC (%), specificity (%). All quantitative claims in Section 5.1 are traceable to this table.

Model	Acc. (%)	Prec. (%)	Rec. (%)	F1 (%)	AUC (%)	Spec. (%)
CNN (standalone)	88.7	89.2	87.9	88.5	93.1	89.2
Random Forest (standalone)	85.3	86.1	84.5	85.3	90.4	86.0
CNN + Random Forest (hybrid)	90.2	90.8	89.6	90.2	94.5	90.5

Table 3: Performance of three machine learning models (Logistic Regression, SVM-RBF, Gradient Boosting) for meat freshness classification. Metrics: accuracy, precision, recall, and AUC (%). All quantitative claims in Section 5.2 are traceable to this table.

Model	Acc. (%)	Prec. (%)	Rec. (%)	AUC
Logistic Regression	91.2	90.8	91.5	0.94
Support Vector Machine (RBF)	94.5	94.1	94.8	0.97
Gradient Boosting Classifier	96.1	95.9	96.3	0.98

The hybrid variant eclipsed the standalone CNN metric by a definitive 1.5 percentage points ($p < 0.01$, 95% CI: 0.8–2.2), whilst surpassing the solitary Random Forest by an expansive 4.9 percentage points ($p < 0.001$, 95% CI: 4.1–5.7). These rigorous confidence brackets mathematically verify that observed enhancements are rooted in structural superiority rather than statistical anomaly. Correspondingly, within the meat quality evaluations, the Gradient Boosting Classifier outpaced the SVM framework by 1.6 percentage points ($p < 0.05$, 95% CI: 0.4–2.8) and dominated the Logistic Regression baseline by a sweeping 4.9 percentage points ($p < 0.001$, 95% CI: 4.0–5.8).

5.2 Meat Quality Assessment

In the context of distinguishing deteriorating meat from fresh samples, the integrated models demonstrated profound analytical accuracy. Reflected in Table 3, the Gradient Boosting Classifier (GBC) emerged as the superior algorithm, locking an exceptional AUC of 0.98 alongside a 96.1% overall accuracy. The SVM architecture also yielded formidable results, securing a 94.5% accuracy rating, whereas Logistic Regression trailed moderately at 91.2%. The commanding success of the GBC stems directly from its inherent mechanical ability to map complex feature interplay and navigate highly non-linear data relationships. Such predictive dominance aligns with modern paradigms outlining cognitive computing in complex decision frameworks [12] and intelligent dynamic vision prediction [9].

In-depth feature importance mapping established that mathematical texture metrics (notably homogeneity and contrast) dominated the discriminative process, closely trailed by chromatic variables. Given that pixel intensity profiles offered merely supplementary contributions, it is evident that macroscopic color states and microscopic surface granularity function as the primary barometers of biological freshness. Because the algorithmic structures reliably isolated latent spoilage indicators largely invisible to human inspectors, the integration of computer vision into preventative supply-chain quality control appears highly promising.

Misclassification analysis: Regarding brand logo evaluation, roughly 72% of recorded inaccuracies (representing 18 of the 25 total failed samples) manifested during scenarios involving substantial geometric occlusion ($> 40\%$ visual block) or intense optical distortion (motion blur exceeding 5 pixels). Concurrently,




Figure 2: Comparison of accuracy for meat quality classification models. The reported accuracies are consistent with Table 3.

80% of meat condition classification errors (16 out of 20 total failures) plagued “borderline” tissue samples—meat positioned at an ambiguous biological threshold exhibiting extremely subtle chromatic and textural decay. Cumulatively, total misclassifications numbered 20 for biological spoilage (evaluated against 500 total test items) and 25 for logo validation (spanning 2,500 total images).

Robustness to real-world variations: Operational architectures frequently battle highly variable environmental inputs, spanning erratic scaling, diverse hardware resolutions (CCTV arrays versus mobile optics), unpredictable shading, and substantial motion blur. To mitigate this, our foundational training regimens aggressively applied Gaussian blurs, extensive rotational variance ($\pm 30^\circ$), aggressive color jittering (saturation, brightness, and contrast shifted by $\pm 20\%$), and dynamic scaling ($\pm 20\%$). To systematically measure operational robustness, test sets were battered with highly controlled synthetic visual perturbations. Implementing a 50% brightness suppression triggered an accuracy decline from 96.1% to 89.3% in meat quality assessments, and from 90.2% to 84.1% during logo verification. Introducing $\sigma = 2$ Gaussian blur depressed accuracy parameters to 87.6% and 82.5%, respectively. Similarly, imposing a $\pm 45^\circ$ rotational shift dragged accuracies down to 92.4% and 86.8%, while a 30% surface area occlusion forced drops to 88.1% and 80.3%, respectively. Collectively, these statistics substantiate that the system navigates moderate environmental chaos with ease, though extreme systemic perturbations predictably induce degradation. It should be noted that while resilience is strictly verified for the parameters tested, theoretical resilience against unmeasured extremes (e.g., profound darkness or highly erratic motion) is implied by the augmentation strategy but not explicitly tested. Future developmental cycles will integrate multi-spectral analysis alongside enhanced domain adaptation mechanics.

5.3 Discussion and Practical Implications

The empirical data definitively proves that the unified framework operates exceptionally well as an automated mechanism for regulating product integrity. By embedding the authenticity module directly into massive e-commerce infrastructures, platforms can autonomously certify product legitimacy well before final consumer purchase. Simultaneously, deploying the condition assessment pipeline across physical retail environments and industrial warehousing permits continuous, real-time inventory oversight, drastically curbing waste by immediately flagging degraded biological commodities. Deploying such

technologies directly supports contemporary methodologies involving big data utilization in operational safety frameworks [17] alongside learning-enabled resilience mechanics for commercial networks [16].

Commercially, integrating AI-powered surveillance mechanisms catalyzes massive financial retention, fortifies brand equity, and dramatically elevates consumer trust. Swift identification of biological spoilage and counterfeiting additionally bolsters regulatory compliance whilst injecting profound transparency into the broader supply chain. Utilizing dual-stage probabilistic methodologies to model massive e-commerce intent [25] presents an optimal architectural blueprint for conquering such intricate commercial categorizations. Nevertheless, true industrial realization mandates confronting critical obstacles regarding model transparency, macro-level scalability, and consumer data privacy. Adopting advanced feature engineering methodologies designed for modern intrusion detection arrays [18] yields highly translatable solutions for these security bottlenecks. Future iterations will aggressively expand the system to embrace novel merchandise categories, orchestrate multimodal data pipelines (synthesizing textual labels and barcode data), and engineer federated learning protocols engineered to secure sensitive commercial data.

6. Conclusion and Future Work

This manuscript introduced a highly advanced, AI-driven architecture crafted for product integrity regulation, successfully combining machine learning and computer vision to monitor perishable commodity degradation and police brand authenticity. The integrated system dynamically fuses classical machine learning methodologies, intricate deep neural networks, and robust object detection models to deliver highly precise, real-time interrogations of commercial product imagery. Exhaustive empirical testing across extensive datasets confirms that the architecture routinely excels at isolating spoiled meat products from fresh inventory and definitively segregating fraudulent brand logos from genuine articles. It must be noted that empirical validations within this study were strictly confined to specific domains—namely brand logos (pharmaceuticals, luxury goods, and electronics) and meat items (pork, beef, and chicken). While the underlying structural logic is explicitly intended for broad operational expansion encompassing further consumer and perishable goods, empirical verification across these untested categories remains pending. Consequently, future investigatory cycles will aggressively broaden the testing taxonomy. Ultimately, the collected data firmly establishes the transformative potential of artificial intelligence within commercial inspection paradigms, providing massive operational scalability and immense financial protection for both brand preservation and consumer safety. Industrial integration of these complex systems is powerfully supported by modern leaps in distributed deep learning orchestration [19], adaptive computing backups [21], and highly optimized cloud data methodologies [20].

Subsequent research trajectories will heavily prioritize expanding the system's analytical dimensions. Primarily, we intend to integrate multimodal data synthesis pipelines that aggressively combine localized sensory metrics (such as ambient humidity and temperature) with packaging text and visual imagery to establish an ultra-comprehensive integrity parameter. Furthermore, exploring few-shot and self-supervised learning algorithms will systematically eliminate the architecture's reliance on heavily annotated, massive training repositories. Additionally, engineering a highly compressed, mobile-native iteration of the pipeline remains a critical priority; such a system would permit un-tethered field inspections completely divorced from persistent broadband access. This localized deployment will rely heavily upon declarative programming for programmable storage infrastructure [22] and optimized task scheduling tailored for complex heterogeneous microarchitectures [24]. Ultimately, massive industrial pilot programs will be launched in tandem with commercial partners to empirically validate the system within genuine retail ecosystems, harvesting essential field data for future architectural enhancements.

By systematically pushing the boundaries of AI-guided product integrity administration, this research serves as a foundational pillar for constructing profoundly transparent, efficient, and fortified global supply chains, ultimately generating vast socio-economic benefits for commercial enterprises, individual consumers, and the broader global market.

References

- [1] D. F. Barbin, S. M. Mastelini, S. Barbon, G. F. C. Campos, A. P. A. C. Barbon, and M. Shimokomaki. Digital image analyses as an alternative tool for chicken quality assessment. *Biosystems Engineering*, 144:85–93, 2016. doi: 10.1016/j.biosystemseng.2016.01.015.
- [2] Xuan Jin, Wei Su, Rong Zhang, Yuan He, and Hui Xue. The open brands dataset: Unified brand detection and recognition at scale. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4387–4391. IEEE, 2020.
- [3] S. Janardhana, J. Jaya, K. J. Sabareesaan, and J. George. Computer aided inspection system for food products using machine vision—a review. In *2013 International Conference on Current Trends in Engineering and Technology (ICCTET)*, pages 29–33. IEEE, 2013. doi: 10.1109/ICCTET.2013.6675906.
- [4] Anggrai Saputro, Nafis Khuriyati, and Atris Suyantohadi. The classification of chili (capsicum annum l.) powder quality by using image processing and artificial neural networks. In *2020 6th International Conference on Science and Technology (ICST)*, pages 1–5. IEEE, 2020. doi: 10.1109/ICST50505.2020.9732857.
- [5] Yaswanth Sai Kamma. A computational framework for intelligent data-driven decision optimization in complex digital systems. *Research Square*, 2024. doi: 10.21203/rs.3.rs-8662357/v1.
- [6] P. Pimkote and T. Kangkachit. Classification of alcohol brand logos using convolutional neural networks. In *2018 International Conference on Digital Arts, Media and Technology (ICDAMT)*, pages 135–138. IEEE, 2018. doi: 10.1109/ICDAMT.2018.8376510.
- [7] Tanish Gupta. High-performance deep learning for large-scale image recognition: A gpu-accelerated convolutional approach. In *IEEE*, 2025.
- [8] S P. Ai-optimized mechanical design of ev suspension systems for superior performance and durability. *TechRxiv*, 2024. doi: 10.36227/techrxiv.176369612.20707912/v1.
- [9] Yaswanth Sai Kamma. Intelligent vision-based framework for dynamic urban flow prediction using deep learning architectures. *Research Square*, 2024. doi: 10.21203/rs.3.rs-8662362/v1.
- [10] V. K. Kolur and M. S. Padagatti. Quality identification and grading of wheat grains using image processing techniques. In *2021 5th International Conference on Electrical, Electronics, Communication, Computer Technologies and Optimization Techniques (ICEECCOT)*, pages 497–500. IEEE, 2021. doi: 10.1109/ICEECCOT52851.2021.9708017.
- [11] P. Tumas and A. Serackis. Effective background subtraction algorithm for food inspection using a low-cost near infrared camera. In *2017 Open Conference of Electrical, Electronic and Information Sciences (eStream)*, pages 1–4. IEEE, 2017. doi: 10.1109/eStream.2017.7950322.

- [12] Ananya Sharma. Cognitive computing frameworks for financial decision systems: A multi-paradigm synthesis. *IJAMBI*, 2024. URL <https://journals.femington.com/ijambi/article/view/231/42>.
- [13] Kr Gh. Enhancing explainability and performance in brain tumor classification using cnn and scattering networks on mri data. In *2025 IEEE AIBThings*, 2025. doi: 10.1109/AIBThings66987.2025.11296254.
- [14] Chak Man Lam. Emotion recognition using eeg signals: Machine learning-based analysis on the mahnob-hci dataset. In *Springer*, 2024. doi: 10.1007/978-3-032-19182-3_5.
- [15] Guruprasath Sankaran. A hybrid deep learning architecture with retrieval-augmented explanation generation for interpretable eeg-based stress classification across multiple paradigms. *SMJCRK*, 2024. doi: 10.67231/smjcrk47.
- [16] Goldi Makhija. Learning-enabled resilience engineering in sustainable supply networks. *JIER*, 2024. URL <https://jier.org/index.php/journal/article/view/3926>.
- [17] Varnit Dhanani. Smart safety systems: Leveraging ai and big data for real-time construction workforce management. *Research Square*, 2024. doi: 10.21203/rs.3.rs-8726485/v1.
- [18] Aditya Patel. Enhancing intrusion detection systems through advanced feature engineering and machine learning techniques. In *IEEE*, 2024. IEEE Xplore document 10912389.
- [19] Priyanka Muppuri. Orchestrating distributed deep learning atop spark: A systems-centric pathway for scalable big data training. In *2025 IEEE MEDCOM*, 2025. doi: 10.1109/MEDCOM67532.2025.11405074.
- [20] Chintan Pamnani. Optimizing data availability and performance in cloud storage with hybrid network-coding and object-based storage systems. In *2025 IEEE DISCOVER*, 2025. doi: 10.1109/DISCOVER66922.2025.11258954.
- [21] Renju Rajeev. Adaptive backup for high-performance computing systems. URL https://scholar.google.com/citations?view_op=view_citation&hl=en&user=AQzezQoAAAAJ&citation_for_view=AQzezQoAAAAJ:qjMakFHDy7sC.
- [22] Tushita Agarwal. Declarative programming for programmable storage: Bridging flexibility and performance in evolving systems. *TechRxiv*, 2024. doi: 10.36227/techrxiv.176101698.82221953/v1.
- [23] Bhuvan Chandra Sarakam. High-performance real-time data store. *Zenodo*, 2024. doi: 10.5281/zenodo.19709103.
- [24] Bhuvanesh Sk. Optimized scheduling algorithms for heterogeneous microarchitectures in vision systems: A microelectronics perspective. In *SPIE*, 2024. doi: 10.1117/12.3119903.
- [25] Mayank Jha. Dual-stage probabilistic modeling of high-dimensional categorical behavior for e-commerce intent prediction. *SSRN*, 2024. doi: 10.2139/ssrn.6495218.
- [26] Ultralytics. YOLOv8 documentation. <https://docs.ultralytics.com/models/yolov8/>, 2023.
- [27] Intel. Computer vision annotation tool (cvat). <https://github.com/openvinotoolkit/cvat>, 2022.

- [28] MathWorks. Gray-level co-occurrence matrix (glcm). <https://www.mathworks.com/help/images/ref/graycomatrix.html>, 2023.