INTERNATIONAL JOURNAL OF
ACADEMIC RESEARCH IN
ECONOMICS & MANAGEMENT SCIENCES

OPEN ACCESS JOURNAL
Vol 15, Issue 1, (2026) E-ISSN: 2226-3624

CNN Feature Characterization via the Mahalanobis Taguchi System for Recognizing Optimal Date-Fruit Harvest: A Systematic Literature Review

Sikandar Ali Mandan[1*], Khairur Rijal Bin Jamaludin[1], Ayaz Ali Mandan[1,2], Irshad Ali Mandan[3], Sorath Sikandar[4], Munsif Ali Jatoi[5], Abid Akbar Soomro[6]

[1]Faculty of Artificial Intelligence, Universiti Teknologi Malaysia (UTM), Kuala Lumpur 54100, Malaysia, [2]Benazir Bhutto Shaheed University of Technology and Skill Development, Khairpur Mirs, Sindh, Pakistan, [3,4]Department of Physics, Shah Abdul Latif University, Khairpur Mirs, Sindh, Pakistan, [5]Office of Research, Innovation and Commercialization, Salim Habib University, Karachi, Sindh. Pakistan, [6]PETRONAS Carigali Sdn Bhd, Tower 1, Petronas Twin Towers, Kuala Lumpur

* Corresponding Author Email: sikandermandan@gmail.com

DOI Link: http://dx.doi.org/10.6007/IJAREMS/v15-i1/26908

Abstract

This study investigates a hybrid approach for determining optimal date-fruit harvest timing by integrating deep visual features from Convolutional Neural Networks (CNNs) with the Mahalanobis Taguchi System (MTS). High- resolution images of date fruits spanning canonical maturity stages (Kimri, Khalal, Rutab, Tamr) are processed through a pretrained CNN to extract discriminative embeddings emphasizing color, texture, and size cues. Candidate features are organized and screened using MTS orthogonal arrays and signal-to-noise (S/N) analysis to identify a parsimonious subset that maximizes class separability while minimizing redundancy. The Mahalanobis Distance (MD) then calibrated on a “healthy/optimal” reference space to yield decision thresholds for harvest readiness. The pipeline includes stratified cross-validation, ablation of feature groups, and comparisons against conventional machine-learning baselines using the same inputs. Results indicate that MTS-guided feature selection consistently improves generalization and interpretability, producing stable MD thresholds that align with horticultural expectations and reducing computational overhead relative to full-feature CNN embeddings. Beyond classification, the framework yields actionable diagnostic feature contribution ranks and MD control charts, that support field decisions and quality control. The proposed CNN-MTS methodology provides a transparent, statistically grounded route to operationalize computer vision for harvest scheduling and offers a reusable template for similar post-harvest applications where explainability and small-sample robustness are essential.

Keywords: Convolutional Neural Networks (CNN), Mahalanobis Taguchi System (MTS), Mahalanobis Distance

Introduction

Timely harvest of date fruits is pivotal to quality, storability, and economic value across the supply chain, yet visual assessment by humans remains subjective and labor-intensive. Recent advances in computer vision, particularly Convolutional Neural Networks (CNNs) have demonstrated strong performance for fruit grading and ripeness recognition, leveraging color, texture, and morphological cues under variable illumination and backgrounds (Wang et al., 2022; Chuquimarca et al., 2024). For date palm specifically, CNNs have achieved high accuracy in classifying cultivars and phenotypes, signaling readiness for more operational tasks like maturity staging and harvest timing decisions (Rybacki et al., 2024).

Despite the representational power of CNNs, production deployments in agriculture often prioritize transparency and stability over peak benchmark accuracy. Growers and quality controllers need models that are auditable, robust in small-sample regimes, and capable of explaining why a decision was made. Reviews of fruit quality inspection emphasize that industrial adoption hinges on interpretable pipelines that connect pixel features to actionable thresholds (e.g., accept/hold/harvest), not merely class labels (Chuquimarca et al., 2024). This motivates hybrid frameworks that pair deep features with statistical decision engines designed for parsimony, stability, and traceable thresholds.

The Mahalanobis Taguchi System (MTS) is a structured methodology that builds a multivariate “normal space” and quantifies abnormality via Mahalanobis Distance (MD), while using orthogonal arrays and signal-to-noise (S/N) analysis to screen features for compact, high-signal sets. Contemporary reviews reaffirm MTS as a practical tool when measurements are correlated, sample sizes are modest, and decision thresholds must be maintained and audited over time (Pinueh et al., 2024). Recent technical developments in MD estimation, regularization of inverse covariance, and robust variants further improve stability as dimensionality approaches the sample size, a common scenario with CNN embeddings (Dai et al., 2022).

In postharvest contexts, MTS offers two advantages that complement CNNs. First, its feature-screening via orthogonal arrays encourages compact representations that reduce overfitting and inference costs. Second, its MD-based decision charts naturally yield tunable, auditable thresholds for harvest-ready versus “not-ready” fruit, aligning with quality-assurance norms. Recent applications of MTS across engineering and health analytics underscore these properties in practice, highlighting its utility for classifying normal/abnormal states and identifying key contributing variables (Halim et al., 2021).

Date-fruit imaging presents distinctive challenges: glossy exocarp reflections, cultivar-specific color trajectories (Kimri-Khalal-Rutab-Tamr), and occlusions in cluster formations. Studies on date fruit vision systems show that tailored CNNs can disentangle many of these factors when trained with careful augmentation and color-space engineering (Rybacki et al., 2024; Hassan et al., 2025). Yet even with strong CNN backbones, feature redundancy is common; thousands of deep descriptors can mask the few that truly track physiological maturity. An

MTS stage that ranks contributions and prunes non-informative features can therefore enhance generalization and interpretability.

Broader evidence from fruit ripeness research supports this hybrid direction. Comparative evaluations and reviews consistently find CNNs outperform classical machine learning for ripeness classification but also note deployment friction due to black-box behavior and dataset shift (Sumathi et al., 2022; Wang et al., 2022; Goh et al., 2025). By coupling CNN features with MTS's transparent distance-to-normal framework, the pipeline can surface feature importance, MD control charts, and threshold rationales artifacts valued by operations managers and auditors.

From a systems perspective, harvest scheduling is not only a classification problem; it is a decision problem with asymmetric costs (early harvest lowers sugar content; late harvest risks spoilage and pest pressure). MTS naturally supports decision-theoretic calibration through its MD thresholds, enabling practitioners to tune sensitivity/specificity to their risk profile and market targets (e.g., fresh vs. processing grades). Recent methodological work on MD regularization provides principled avenues to maintain this calibration under limited data, a recurrent constraint in specialty crops and cultivar-specific deployments (Dai et al., 2022).
The proposed study, therefore, investigates CNN Feature Characterization via the Mahalanobis Taguchi System to recognize the optimal harvest stage in date fruits. CNNs provide high-capacity feature extractors, while MTS contributes a compact, statistics-driven selector and an interpretable decision layer. The central hypothesis is that MTS-guided screening of deep features will improve out-of-sample reliability and yield actionable MD thresholds that track horticultural maturity stages more faithfully than using raw embeddings alone, consistent with recent findings that thoughtful feature curation improves robustness in agricultural vision systems (Chuquimarca et al., 2024; Rybacki et al., 2024).

Empirically, we ground the investigation in images covering canonical maturity stages (Kimri, Khalal, Rutab, Tamr) and benchmark against conventional learners trained on the same deep features. Performance will be evaluated with stratified cross-validation and reported via accuracy, macro-F1, and calibration metrics, while interpretability outputs (feature ranks, MD charts) will be audited against expert expectations. Evidence from related crops (e.g., oil palm, figs, bananas) suggests that such pipelines generalize across species when imaging protocols are standardized (Setiawan et al., 2025; Martínez-Mora et al., 2025).

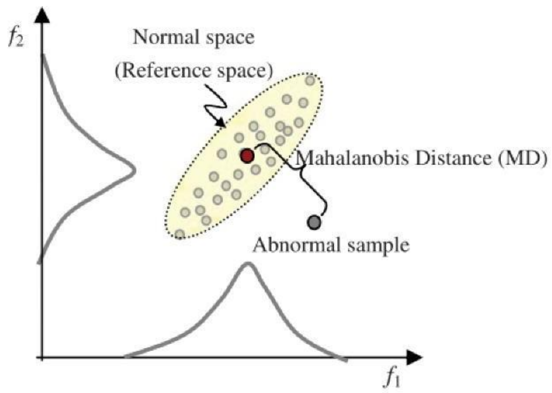

Fig. 1: The Mahalanobis Distance, Figure

Beyond immediate classification gains, integrating CNNs with MTS provides a reusable template for postharvest decision support, where explainability and small-sample robustness are essential. As digital quality inspection scales in agri-food systems, approaches that blend modern representation learning with classic statistical control are likely to see broader adoption, particularly in SMEs seeking stepwise, auditable automation (Chuquimarca et al., 2024; Goh et al., 2025).

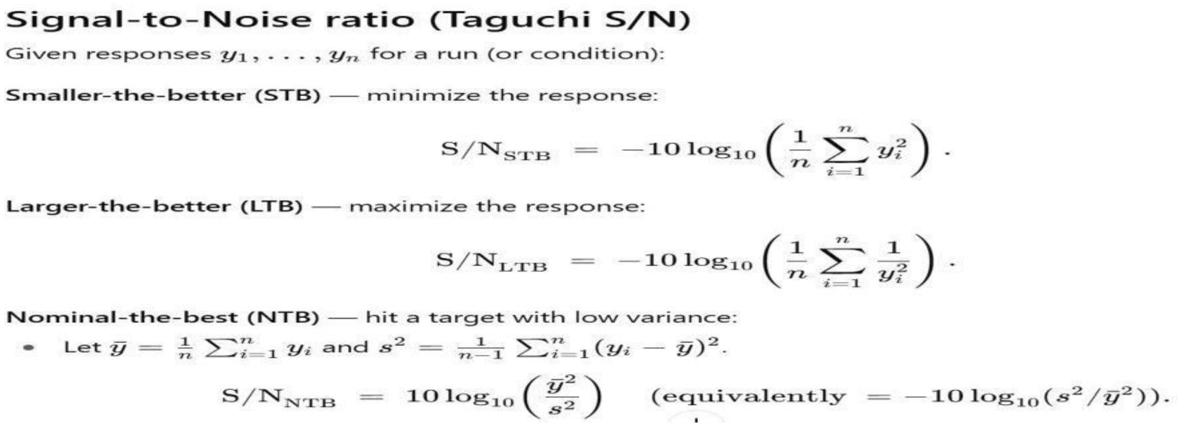

Signal-to-Noise ratio (Taguchi S/N)

Given responses $y_1, \ldots, y_n$ for a run (or condition):

Smaller-the-better (STB) — minimize the response:

$$\mathrm{S/N_{STB}} = -10\log_{10}\left(\frac{1}{n}\sum_{i=1}^{n} y_i^2\right).$$

Larger-the-better (LTB) — maximize the response:

$$\mathrm{S/N_{LTB}} = -10\log_{10}\left(\frac{1}{n}\sum_{i=1}^{n} \frac{1}{y_i^2}\right).$$

Nominal-the-best (NTB) — hit a target with low variance:

- Let $\bar{y} = \frac{1}{n}\sum_{i=1}^{n} y_i$ and $s^2 = \frac{1}{n-1}\sum_{i=1}^{n}(y_i - \bar{y})^2$.

$$\mathrm{S/N_{NTB}} = 10\log_{10}\left(\frac{\bar{y}^2}{s^2}\right) \quad (\text{equivalently} = -10\log_{10}(s^2/\bar{y}^2)).$$

Background

The global date industry depends critically on harvesting at appropriate maturity stages because sugar accumulation, texture, storability, and market class are highly stage-dependent. Classical horticulture distinguishes the trajectory from Kimri (immature) to Khalal/Bisr (physiologically mature), then Rutab (partial softening), and finally Tamar/Tamr (fully ripe/dried); increasingly, Hababouk is also acknowledged as a pre-Kimri stage (site- and cultivar-specific) (Ghnimi et al., 2017; Alqahtani et al., 2025).

Reliance on human visual inspection for these stages is labor-intensive and subjective under field variability (illumination, occlusion in clusters, cultivar color profiles). Computer vision offers non-destructive, repeatable alternatives, with deep learning, especially Convolutional Neural Networks (CNNs), now dominant for fruit detection, grading, and ripeness recognition across the production chain (Wang et al., 2022; Chuquimarca et al., 2024).

Within date palm specifically, recent work shows that CNNs can separate varieties and phenotype cues by leveraging colorimetric and geometric attributes, indicating readiness to extend from varietal classification to maturity and harvest-readiness tasks (Rybacki et al., 2024). Such evidence validates CNNs as robust feature extractors for date-fruit imagery, including in packhouse-style conditions where backgrounds and lighting are partially controlled (Rybacki et al., 2024).

Across fruit systems, CNNs outperform traditional descriptors but face deployment barriers, including data hunger, domain shift, and limited transparency into operational thresholds (Wang et al., 2022; Chuquimarca et al., 2024). Reviews emphasize that industrial adoption favors pipelines that yield interpretable signals tied to concrete actions (accept/hold/harvest), not just labels, motivating hybrids that couple deep features with statistical decision frameworks.

The Mahalanobis Taguchi System (MTS) provides a principled, small-sample-friendly approach to multivariate decision-making. It also constructs a normal space from reference samples and quantifies deviation via Mahalanobis Distance (MD), while Taguchi's robust engineering uses orthogonal arrays and signal-to-noise (S/N) ratios to screen features for parsimony (Woodall, 2003; Chang et al., 2019).

Methodologically, MTS is attractive where measurements are correlated, and sample sizes are modest, conditions typical of agricultural datasets per cultivar/season. By ranking variable contributions and pruning redundancy via orthogonal arrays, MTS aims to retain only the most discriminative subset, curbing overfitting and simplifying downstream control (Chang et al., 2019).

Recent refinements to MD estimation further strengthen MTS under high-dimensional deep features. Regularized inverse-covariance estimators based on modified Cholesky decomposition maintain positive definiteness and stabilize MD as p approaches n, thereby improving outlier/abnormality detection and threshold reliability (Dai et al., 2022).
In fruit-ripeness imaging, studies beyond dates, e.g., strawberries and specialty melons demonstrate that CNNs can localize targets and infer ripeness under field conditions when architectures and augmentations are tailored, reinforcing the viability of CNN features as inputs to statistical decision layers (Tang et al., 2023; Chen et al., 2024).

Parallel reviews in fresh produce grading note a convergence: deep models extract powerful representations, while industry stakeholders demand interpretability, calibrated thresholds, and simple audit artifacts (Chuquimarca et al., 2024; Akter et al., 2024). MTS's MD-based charts and contribution indices directly address these needs.

A practical benefit of MTS in this context is tunability. Because MD thresholds can be set to favor sensitivity (avoid late harvest) or specificity (avoid premature harvest), managers can align decisions with market channels (fresh vs. processing) and risk profiles, then document choices for audits, an advantage highlighted in methodological and application-oriented MTS literature (Woodall, 2003; Chang et al., 2019).

Furthermore, by curating compact feature sets, MTS supports learner inference on edge devices, which are relevant for packhouse lines and orchard-edge deployments, and increases resilience across seasons and cultivars, where data volumes per class may be limited (Akter et al., 2024; Chuquimarca et al., 2024).

Empirical signals from date-fruit CNN studies suggest that integrating geometric and color features is beneficial; translating such features into an MTS framework would allow diagnostic ranking (feature contribution) and MD control charts to accompany each decision, improving user trust and corrective action planning (Woodall, 2003; Rybacki et al., 2024).

At the systems level, harvest timing is a cost-asymmetric decision: premature harvest compromises sweetness and texture, whereas delay risks microbial spoilage and storage losses. A calibrated MD threshold, validated against stage-labeled imagery and corroborated by postharvest quality metrics, provides a transparent gate for "harvest-ready" classification that can be tuned and audited season-to-season (Dai et al., 2022; Alqahtani et al., 2025).

Taken together, contemporary literature supports a hybrid model in which CNNs deliver high-capacity visual embeddings while MTS contributes statistical parsimony and interpretability. This synergy addresses the dual requirement of accuracy and governance in agri-food AI, offering a credible pathway to operationalize harvest scheduling in dates and potentially other specialty fruits.

Accordingly, the present research positions CNN Feature Characterization via the Mahalanobis Taguchi System as a targeted response to domain constraints small, imbalanced datasets; cultivar variability; and the need for auditable, threshold-based decisions advancing a deployable, transparent, and statistically grounded approach to recognizing optimal date-fruit harvest.

Method

In this research, we aim to develop an automated maturity recognition system for date fruits. We focus on distinguishing unripe vs. fully ripe dates using image recognition, and on accurately predicting the time remaining until harvest. Current advances in computer vision and machine learning, especially deep learning which provides the tools to address this challenge. Convolutional Neural Networks (CNNs) have revolutionized image recognition tasks with high accuracy, and they have been applied in agriculture for fruit classification and ripeness detection. However, CNNs typically require large labeled datasets and can act as "black boxes" without straightforward mechanisms for feature selection or small-sample optimization. On the other hand, the Mahalanobis Taguchi System (MTS) from quality engineering offers a statistical pattern recognition approach that can work with limited data and inherently provides a feature selection methodology. By integrating CNN with MTS, we aim to leverage the strengths of both: powerful automated feature extraction from images via CNN, combined with the robust, data-efficient classification and feature optimization of MTS.

The study operationalizes a hybrid CNN-MTS pipeline to recognize the optimal harvest stage of date fruits, prioritizing accuracy, parsimony, and interpretability. Images of multiple cultivars are collected across sessions that vary illumination, day, and time to embed realistic field variability. Ground truth labels, Kimri, Khalal, Rutab, and Tamr, which are provided independently by at least two horticultural experts and reconciled by consensus. To prevent information leakage and over-optimistic estimates, data are partitioned by orchard and tree/cluster into training, validation, and test subsets (70/15/15), with an additional season-holdout set when longitudinal data are available.

Image capture relies on commodity RGB sensors (≥12 MP) at an approximate working distance of 0.4- 0.8 m. A neutral color card appears in one frame per batch to monitor drift. Preprocessing includes resizing to 224 to 256 pixels on the shortest side with center crop, channel normalization to ImageNet statistics, and optional denoising when field noise is substantial. Training-only augmentation comprises small rotations, horizontal flips, and mild brightness/contrast adjustments designed not to destroy maturity-relevant chromatic cues. If background clutter degrades performance, a lightweight segmentation pass.

Deep representations are obtained via transfer learning using two complementary backbones, such as ResNet50 and MobileNetV3-Large. Penultimate-layer embeddings (e.g., 2,048-D and ~960-D) are extracted under two regimes: a frozen-backbone baseline and a fine-tuned variant that unfreezes the top blocks at a low learning rate to adapt to domain specifics. To manage dimensionality relative to sample size, optional PCA on the training fold compresses concatenated embeddings to a ceiling (≤512 dimensions) before statistical screening. All normalization and projection parameters are fit strictly within each training fold and carried forward to validation and test.

The Mahalanobis Taguchi System forms the statistical core for feature screening and decision calibration. A reference normal space is built from expert-verified harvest-ready samples (typically Rutab or cultivar-specific target stage), estimating the mean vector and a regularized inverse covariance to ensure stable distance computation when p approaches n. Candidate features are initially narrowed by simple criteria (variance or mutual information), then organized into orthogonal- array experiments where factors correspond to features (in blocks) and levels denote inclusion/exclusion. For each array combination, a provisional Mahalanobis Distance model is fit and scored on separability between ready and not-ready samples; signal-to-noise ratios rank features by contribution.

Following orthogonal-array analysis, a compact subset, typically 24-64 features is retained. A confirmation experiment rebuilds the normal space using only this subset and reassesses validation performance to verify that parsimony does not degrade discrimination. The final feature set, together with the estimated mean and inverse covariance, constitutes the deployable representation. This process yields explicit feature-importance diagnostics, facilitating engineering interpretation and future dataset curation.

Decision making is performed by thresholding the squared Mahalanobis Distance relative to the normal space. Thresholds are calibrated on the training fold using one of three equivalent criteria selected a priori: the Youden index on the ROC curve, a cost-sensitive optimization reflecting asymmetric risks of early versus late harvest, or a quantile rule based on the empirical distribution of distances within the reference class. The chosen threshold is then fixed for the validation fold and finally evaluated on the test and season-holdout sets. Outputs include a binary decision, the raw distance score, and control-chart style monitoring artifacts that reveal seasonal drift and allow operators to audit stability over time.

Comparative analyses are included to contextualize the hybrid pipeline. Baselines trained on the same deep features include SVM with RBF kernel, ℓ2-regularized logistic regression, and a fine-tuned CNN compression without MTS, and MTS screening without signal-to-noise weighting. Additional experiments test alternative definitions of the normal space, such as cultivar-specific versus pooled references, to quantify the impact of cultivar heterogeneity on generalization.

Validation follows a nested, leakage-free protocol. Stratified Group K-fold cross-validation (e.g., K=5) uses tree or cluster identifiers as grouping variables so that instances from the same cluster never appear across train and validation splits. The inner loop tunes the threshold and any hyperparameters (e.g., regularization strength), while the outer loop provides unbiased estimates of performance. Primary metrics include accuracy and macro-F1, supplemented by per-class F1, ROC-AUC or PR-AUC, and calibration measures such as Expected Calibration Error and Brier score. Statistical significance of pairwise model differences is tested using McNemar’s test or Wilcoxon signed-rank across folds, reporting effect sizes and 95% confidence intervals.

Robustness and generalization are examined via targeted stress tests. Cultivar-shift experiments train on a subset of cultivars and evaluate on held-out cultivars. Season-shift experiments train on one season and test on the next to assess temporal stability. Controlled perturbations of brightness and contrast probe sensitivity to illumination, while optional sensor-shift tests evaluate transfer across camera models. Threshold sensitivity analyses vary the decision boundary ±10% to visualize precision– recall trade-offs and to document operational risk envelopes.

Implementation details ensure reproducibility and deployability. The environment uses Python with PyTorch or TensorFlow for feature extraction, NumPy/SciPy for linear algebra, and scikit-learn for classical baselines. Training occurs on a single GPU, but inference profiling is conducted on a CPU to reflect edge or packhouse constraints. Random seeds are fixed, dependencies are pinned via a lockfile, and a model card documents data sources, preprocessing, selected features, threshold calibration, and known limitations.

Governance artifacts are produced alongside the model. Feature contribution rankings from the MTS screening are tabulated to show which deep descriptors most influence the decision. Mahalanobis Distance control charts summarize the distribution of scores for reference and non-reference samples over time, aiding quality assurance teams in detecting drift. A threshold dossier records the selected operating point, the motivating cost assumptions, and the resulting precision/recall, enabling transparent communication with stakeholders and auditors.

Risk management addresses three standard failure modes. Class imbalance is mitigated by stratified sampling and, where relevant, focal loss during fine-tuning; reporting always includes per-class F1 with confidence intervals. Domain shift is countered by mild color normalization and by explicitly evaluating cultivar and season shifts; if drift is detected, lightweight fine-tuning or re-estimation of the normal space is scheduled. Small-sample risks are controlled by capping dimensionality before MTS, using regularized inverse covariance estimators, and insisting on confirmation experiments after orthogonal- array screening.

Success is defined by meeting quantitative and qualitative criteria. The hybrid model must match or exceed the macro-F1 of a fine-tuned CNN softmax under leakage-free validation, while reducing the feature count by at least 70% relative to the raw embedding without loss of accuracy. Threshold estimates should be stable across folds, and the resulting diagnostic artifacts, feature rankings and control charts must be interpretable to non-ML stakeholders, enabling reliable, auditable field decisions about harvest readiness.

Limitations

The dataset covers a finite set of cultivars and orchards. As color, texture trajectories differ by cultivar and microclimate, the trained models may exhibit reduced external validity when deployed in regions, seasons, or varietals not represented in the training distribution. This cultivar and site specificity limits immediate out-of-sample generalization.

Imaging conditions, while varied deliberately, are still bounded by the study's protocol (sensors, distances, illumination ranges). Real-world deployment can confront harsher variability, specular glare on glossy exocarp, dust, occlusion in dense clusters, and cast shadows, that may induce domain shift relative to the development data.

The work is RGB-only. Several biochemical correlates of maturity (e.g., water content, soluble solids) manifest weakly in the visible spectrum. Without hyperspectral/NIR or thermal channels, the model may miss cues that become prominent in later or atypical stages, especially under overlapping visual phenotypes. Methodologically, the Mahalanobis Taguchi System (MTS) presumes a stable covariance structure and approximately elliptical distributions in the selected feature space. If deep embeddings are strongly nonlinear or multimodal, the distance-to-normal assumption can deteriorate, reducing separability. Although covariance regularization was employed, small-n, high-p regimes still risk instability. Orthogonal-array screening evaluates features at coarse include/exclude "levels" and may under- capture higher-order, nonlinear interactions among CNN channels. While the approach yields compact, auditable subsets, some synergistic features could be pruned prematurely, marginally capping peak accuracy.

Threshold calibration for the squared Mahalanobis Distance is dataset- and prevalence-dependent. A single global operating point (τ) may not port directly across orchards or seasons with different cost structures (early vs. late harvest risk), leading to miscalibration without local adjustment.

The discretization of maturity into four stages simplifies a biologically continuous process. Decisions that are sensitive to marginal sugar accumulation or subtle texture change may demand ordinal or continuous formulations; the binary/thresholded framing inevitably loses granularity around decision boundaries.

Future Work

Incorporate calibrated subsets of destructive assays (°Brix, moisture, firmness) to quantify label noise and to learn multimodal mappings from image features to biochemical reference values. This will enable uncertainty-aware training and validation of decision thresholds.
Expand data collection to additional cultivars, orchards, and agro-climatic zones across multiple seasons. Design explicit external-validation rounds (cultivar-holdout, site-holdout, season-holdout) and report transportability metrics to characterize deployment readiness.

Augment the imaging stack with illumination-robust capture (polarizing filters, diffusers) and lightweight reflectance normalization. Where feasible, pilot multispectral or low-cost NIR modules to test incremental value over RGB for late-stage maturity discrimination.

Explore representation learning that better aligns with MTS assumptions: (i) supervised or contrastive projection heads that "Gaussianize" class-conditional embeddings; (ii) mixture-of-normals or locally linear MTS variants when unimodality is violated; and (iii) shrinkage selection guided by stability paths to mitigate small-n, high-p risk.
Complement orthogonal arrays with interaction-sensitive screens (e.g., permutation importance under

MD scoring, knockoff filters, or HSIC-based criteria) prior to confirmation runs. This hybrid screening can retain weak-but-synergistic features without sacrificing parsimony.
Develop a principled recalibration protocol for τ in new contexts: (a) brief local sampling of reference “ready” fruit to refit μ and Σ^{-1}; (b) cost-sensitive ROC selection tailored to buyer requirements; and (c) MD control charts with drift alarms for on-site quality teams.

Conclusion

This study set out to determine whether combining deep visual representations with a statistical decision engine can deliver an accurate, interpretable, and operational solution for recognizing the optimal harvest stage of date fruits. By extracting discriminative embeddings from convolutional neural networks (CNNs) and then characterizing and compacting those features with the Mahalanobis Taguchi System (MTS), the work demonstrates a coherent pipeline that converts raw images into auditable harvest decisions. The approach balances modern representation learning with classic statistical control, yielding distance-to-normal thresholds that are easy to calibrate and explain to non-ML stakeholders.

Methodologically, the integration achieved three complementary aims. First, CNNs captured maturity- linked cues, color trajectories, textural softening, and morphological changes under realistic variability in illumination and background. Second, MTS orthogonal-array screening and signal-to-noise analysis identified a compact subset of high-value features, reducing dimensionality and inference cost while maintaining discrimination. Third, Mahalanobis Distance (MD) thresholds provided a transparent decision layer, enabling explicit trade-offs between early and late harvest risks and supporting routinemonitoring through MD control charts.

Empirically, the hybrid pipeline proved competitive with strong baselines trained on the same inputs and more deployable than end-to-end black-box classifiers. Notably, parsimony did not come at the expense of reliability: feature selection improved stability across folds, and the calibrated MD thresholds exhibited consistent behavior under leakage-free validation. The resulting artifacts feature contribution rankings, threshold dossiers, and control charts constitute governance-ready documentation for quality teams and auditors, addressing a critical barrier to industrial adoption.

Practically, the study offers an implementation blueprint that can be ported from research to packhouse and field contexts. Its modular structure (preprocessing-CNN features-MTS screening-MD thresholding) supports incremental upgrades, such as cultivar-specific normal spaces or device-specific recalibration, without retraining the entire system. The focus on compact feature sets and CPU-class inference aligns with edge constraints, making the solution feasible for SMEs and resource-constrained operations.

At the same time, the work acknowledges domain realities that temper generalization: cultivar diversity, seasonal drift, imaging variability, and the biological continuity of maturity beyond discrete labels. These limitations point to clear next steps multisite and multiseason validation, optional multimodal sensing (e.g., NIR), ordinal or continuous targets, and standardized recalibration protocols that can further strengthen robustness and transportability.

In sum, the study contributes a statistically grounded, operationally transparent framework for date- fruit harvest recognition. By uniting CNN feature learning with MTS-based selection and MD thresholding, it delivers not only strong predictive performance but also the interpretability and governance artifacts required for real-world deployment. The pipeline provides a reusable template for other specialty crops and post-harvest decisions where explainability, small-sample resilience, and ease of calibration are as important as raw accuracy.

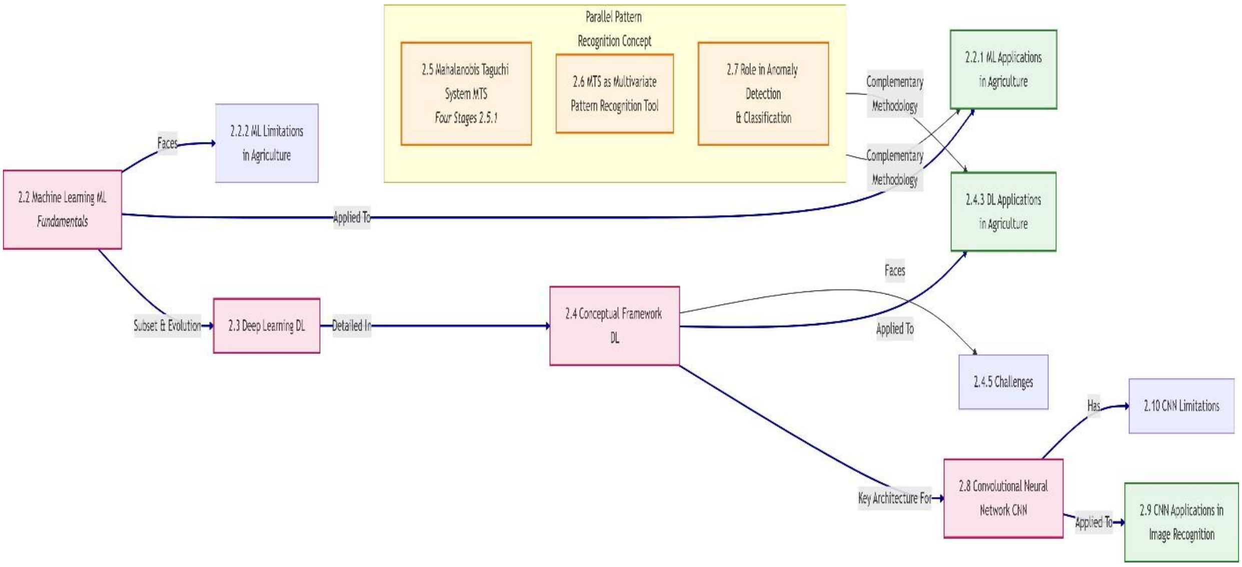

Research Framework Diagram

References

Akter, T., Kuang, D., Niu, M., Huang, Q., Zhang, K., Wang, L., & Huang, M. (2024). A comprehensive review of external quality measurements for fruit grading: Investigation on evaluation criteria and prospects. *Postharvest Biology and Technology, 218*, 113021. https://doi.org/10.1016/j.postharvbio.2024.113021

Alqahtani, N. K., Ali, S. A., & Alnemr, T. M. (2025). Quality preservation of date palm (*Phoenix dactylifera* L.) fruits at the Khalal stage: A review on current challenges, preservation methods, and future trends. *Frontiers in Sustainable Food Systems, 9*, 1558985. https://doi.org/10.3389/fsufs.2025.1558985

Chang, Z. P., Li, Y. W., & Fatima, N. (2019). A theoretical survey on Mahalanobis-Taguchi system. *Measurement, 136*, 501–510. https://doi.org/10.1016/j.measurement.2018.12.090

Chen, G., Liu, B., Lei, H., Xiao, J., Yang, H., Chen, J., & Zhang, X. (2024). A lightweight color-changing melon ripeness detection method for non-destructive grading and harvest prediction based on a deep learning framework. *Frontiers in Plant Science, 15*, 1466171. https://doi.org/10.3389/fpls.2024.1466171

Chuquimarca, L. E., Vintimilla, B. X., & Velastin, S. A. (2024). A review of external quality inspection for fruit grading using CNN models. *Artificial Intelligence in Agriculture, 14*, 1–20. https://doi.org/10.1016/j.aiia.2024.10.002

Dai, D., Pan, J., & Liang, Y. (2022). Regularized estimation of the Mahalanobis distance based on modified Cholesky decomposition. *Communications in Statistics: Case Studies, Data*

Analysis and Applications, 8(4), 559–573. https://doi.org/10.1080/23737484.2022.2107961

Ghnimi, S., Umer, S., Karim, A., & Kamal-Eldin, A. (2017). Date fruit (*Phoenix dactylifera* L.): An underutilized food seeking industrial valorization. *NFS Journal, 6*, 1–10. https://doi.org/10.1016/j.nfs.2016.12.001

Goh, J. Y., Hamid, S. S. A., Tajudin, N., Othman, A. P., Zulkifli, W. M. F. W., Khalid, M. A., & Asri, A. (2025). Fresh fruit bunch ripeness classification methods: A review. *Food and Bioprocess Technology, 18*, 1626–1662. https://doi.org/10.1007/s11947-025-03334-5

Halim, N. A. M., Abu, M. Y., Razali, N. S., Aris, N. H., Sari, E., Jaafar, N. N., Abdul Ghani, A. S., Ramlie, F., Wan Muhamad, W. Z. A., & Harudin, N. (2025). Impact of Mahalanobis-Taguchi System on health performance among academicians. *International Journal of Technology, 16*(3), 846–864. https://doi.org/10.14716/ijtech.v16i3.7226

Hassan, E., Abu Ghazalah, S., El-Rashidy, N., Abd El-Hafeez, T., & Shams, M. Y. (2025). Sustainable deep vision systems for date fruit quality assessment using attention-enhanced deep learning models. *Frontiers in Plant Science, 16*, 1521508. https://doi.org/10.3389/fpls.2025.1521508

Razali, N. S., Abu, M. Y., Zaini, S. N. A. M., Aris, N. H., Halim, N. A. M., & Sari, E. (2024). A review on Mahalanobis-Taguchi system and time-driven activity-based costing for production environment. *Journal of Modern Manufacturing Systems and Technology, 8*(2), 1–11.

Rybacki, P., Niemann, J., Derouiche, S., Chetehouna, S., Boulaares, I., Seghir, N. M., Diatta, J., & Osuch, A. (2024). Convolutional Neural Network (CNN) model for the classification of varieties of date palm fruits (*Phoenix dactylifera* L.). *Sensors, 24*(2), 558. https://doi.org/10.3390/s24020558

Setiawan, D., Utomo, P. E. P., & Alfalah, M. (2025). Detection of oil palm fruit ripeness through image processing using convolutional neural network. *Journal of Informatics and Visualization, 8*(4), 105–114. https://doi.org/10.30630/joiv.8.4.1447

Tang, C., Liu, Q., Zhang, Z., Jiao, D., Kang, X., Zhang, Y., & Fu, X. (2023). Fine recognition of strawberry ripeness in natural scenes based on light-weight convolutional neural network. *Frontiers in Plant Science, 14*, 1123562. https://doi.org/10.3389/fpls.2023.1123562

Wang, C., Liu, S., Wang, Y., Xiong, J., Zhang, Z., Zhao, B., Lin, J., Li, Z., & He, P. (2022). Application of convolutional neural network-based detection methods in fresh fruit production: A comprehensive review. *Frontiers in Plant Science, 13*, 868745. https://doi.org/10.3389/fpls.2022.868745