Natural Resources for Human Health
eISSN: 2583-1194

Original Research

Received 09 March 2026
Revised 12 April 2026
Accepted 25 May 2026

Natural Resources for Human Health 2026, Vol. 6 (5s): 912–926

KEYWORDS:
Tur Dal; Image Processing; Feature Extraction; Variety Identification; Polishing Status; ResNet50; Performance.

Automated Variety Identification and Polishing Status Classification of Tur Dal Grains Using Computer Vision and Machine Learning

Mrs. Veena B. Mindolli[1], Dr. Mahantesh C. Elemmi[2]

1Assistant Professor, Department of CSE, Jain College of Engineering & Research, Belagavi; Visvesvaraya Technological University, Belagavi, Karnataka, India
2Professor, Department of CSE, Navkis College of Engineering, Hassan; Visvesvaraya Technological University, Belagavi, Karnataka, India

ABSTRACT: Manual grading of tur dal (Cajanus cajan) remains subjective, inconsistent, and poorly scalable, underscoring the need for objective, image-based inspection. This study develops an integrated computer vision–machine learning framework for simultaneous variety identification and polishing-status classification of tur dal grains. Twenty-five morphometric, colorimetric, surface-optical, and defect-related features were extracted from 1,500 grain images (five varieties, two polishing states) and validated statistically (one-way ANOVA, chi-square, $p < .001$) before training a multi-layer perceptron classifier. On a held-out test set ($n = 300$), the model achieved 99.50% accuracy, 0.0006 loss, and precision/recall/F1 nearly equalling to 1.00 across grades, with ROC–AUC approaching 1.00; permutation importance identified weevil damage, colour uniformity, and broken-grain percentage as dominant predictors. Findings confirm feasibility under controlled imaging conditions, though generalisability to field-acquired, variably lit samples remains untested. The study's originality lies in unifying three inspection tasks within one interpretable, statistically grounded pipeline for an understudied pulse crop.

1. INTRODUCTION

Grain quality inspection sits at the junction of agricultural livelihood and food-industrial efficiency, and few crops illustrate this junction more directly than tur dal (pigeon pea, Cajanus cajan), one of India's most widely consumed pulse crops and a primary source of dietary protein for a large share of the country's population. Every kilogram of tur dal that moves from field to mill to market passes through at least one quality-determining decision point — variety confirmation, polishing assessment, and grade assignment — and in the overwhelming majority of procurement centres across India, that decision is still made by a human inspector working from visual cues and years of accumulated judgement. Manual grading of this kind is not without merit; experienced graders are fast and contextually adaptive. It is, however, inherently subjective, inconsistent across inspectors and inspection sessions, and difficult to scale as throughput demands increase, a limitation increasingly at odds with the traceability and standardisation expectations of modern agri-food supply chains.

Recently, computer vision and machine learning have played a major role in the evaluation of agricultural products in terms of quality, greatly improving the classification of products such as wheat, rice, and maize. In the case of tur dal, classification is more challenging because of the crop's complex shape, size, texture, and colour variability arising from environmental effects and processing. Moreover, polishing level is a major criterion for consumer acceptability and nutritional value, and is a complex feature since it depends on surface texture.

To address these issues, there is a need for feature representations and learning paradigms that are specifically designed to capture subtle differences between closely related classes of grains while remaining invariant to environmental changes. This research contributes a state-of-the-art approach in agricultural machine vision by developing an integrated framework for automatic identification of dal varieties and classification of polishing status in tur dal grains.

Natural Resources for Human Health
eISSN: 2583-1194

Materials Used

This study is performed on five different varieties of tur dal samples — Asha, Nagpur, Navapur, Gulyal, and Latur — grown in various regions of India. Samples of all five varieties, in both polished and unpolished form, were collected and used in this research. Table 1 shows sample images of the five varieties in both polished and unpolished form.

Table 1. Different samples of tur dal, polished and unpolished

S.No / Variety	Unpolished	Polished
1. Asha		
2. Latur		
3. Navapur		
4. Gulyal		
5. Nagpur		

Table 1 shows the different samples of various varieties of tur dal used for this study. The Asha variety is medium-sized, oval to slightly elongated, with a uniform internal structure and a smooth, bright-yellow surface after polishing; unpolished grains have a thin brownish seed coat. The Latur variety is medium-to-large, bold, and more rounded, with a comparatively deeper yellow colour after polishing and a thicker, darker brown outer coat when unpolished. The Navapur variety is relatively smaller, compact, and slightly flattened, with high uniformity, a light-yellow colour after polishing, and a thin, pale-brown seed coat when unpolished. The Nagpur variety is moderately large and elongated-oval, with a rich golden-yellow colour after polishing and a distinctly visible brownish husk with surface-texture variation when unpolished. The Gulyal variety is medium-sized with slightly asymmetrical geometric shapes and a comparatively rougher unpolished texture, showing a bright yellow, glossy surface after polishing. These natural differences in size, shape, colour intensity, and surface texture across the five varieties can be effectively exploited for computer-vision-based classification using morphological, colour-space, and texture-feature analysis.

Related Work and Research Gap

A structured literature review following PRISMA 2020 guidelines was conducted across the Scopus database using a Boolean combination of terms including computer vision, machine learning, deep learning, grain classification, quality grading, variety identification, polishing status, pigeon pea, tur dal, and image processing. Studies employing convolutional neural networks (CNNs) for agricultural image classification consistently reported accuracies exceeding 92%, with some exceeding 98% on controlled imaging datasets, while support vector machines and random forest classifiers showed slightly lower accuracy but superior interpretability and faster training on smaller, feature-engineered datasets. Hybrid approaches combining handcrafted morphological and colour features with deep feature extraction generally yielded the most robust performance, suggesting that end-to-end deep learning alone is not always optimal for small-sample food-grain classification.

RGB-based imaging remains the most widely adopted modality for grain inspection owing to its low cost and ease of integration into industrial sorting lines, typically paired with colour-space transformation and contrast-enhancement preprocessing. Among algorithms, pre-trained CNN backbones such as ResNet, VGGNet, and MobileNet are frequently fine-tuned via transfer learning to offset limited labelled training data, while SVM and random-forest models remain popular where interpretability is important for regulatory compliance. For pulses specifically, colour parameters (L*, a*, b*), surface glossiness, and the ratio of polished to unpolished surface area have been established as reliable visual proxies for polishing grade, though tur-dal-specific studies remain scarce compared with rice and other cereals.

A structured gap analysis identified six recurring shortcomings in the literature: the absence of publicly available, annotated tur-dal image datasets; the lack of automated, CV-based polishing-status classifiers specific to tur dal; the absence of an integrated multi-task model addressing variety identification and polishing status jointly; a shortage of lightweight, edge-deployable models validated beyond laboratory conditions; limited interpretability of deep models for regulatory food-grading compliance; and little fusion of RGB, hyperspectral, and physical (weight) data for pulses. These gaps motivate the integrated computer-vision and machine-learning pipeline developed in this study, which draws on hierarchical feature learning from computer-vision theory, transfer learning from machine-learning theory, and established pulse quality-grading standards from food science.

Scope and Objectives

The scope of this study is bounded in specific, deliberate ways. It draws on a laboratory-compiled dataset of 1,500 grain records spanning five commercial tur dal varieties and two polishing states, captured under standardised, diffuse lighting rather than variable field or warehouse conditions; the analysis is accordingly a controlled feasibility demonstration rather than a claim of field-ready deployment. In direct response to the gaps identified above, the study pursues four specific objectives:

1. To develop and document a computer-vision-based feature-extraction pipeline capable of deriving morphometric, colorimetric, and surface-texture descriptors from digital images of tur dal grains under standardised imaging conditions.

2. To determine whether grain morphometry (length, length-to-width ratio, and related geometric descriptors) differs systematically across commercially cultivated tur dal varieties, thereby establishing the feasibility of vision-based variety identification.

3. To determine whether surface-optical descriptors, principally gloss index and oil-coating percentage, differentiate polished from unpolished grains, thereby establishing the feasibility of an automated polishing-status classifier.

4. To statistically validate the extracted features through one-way ANOVA and chi-square testing, and to assess the interpretability of the trained classifier through permutation feature importance, so that model performance can be corroborated at the feature level rather than accepted on accuracy alone.

2. MATERIALS AND METHODS

The experimental framework was designed to perform dual-level classification — (i) tur dal variety identification and (ii) polishing-status classification — using a ResNet50 model, supported by classical statistical hypothesis testing of the extracted features.

2.1 Dataset Preparation and Image Acquisition

Tur dal samples representing the five commercial varieties were collected from retail shops, APMC markets, and farmers; both polished and unpolished grains were included. Images were captured using a 24 MP digital camera under a controlled LED-light environment with a fixed camera-to-sample distance of 25 cm, against a white background to minimise reflection. The dataset was organised into labelled folders by variety and by polished/unpolished category, and split 80% training, 10% validation, and 10% testing. Figure 1 shows the screenshot of capture the images of tur dal.

(a) (b) (c)

Figure 1. *Capturing of images under controlled environment.*

Table 2 reports the joint distribution of samples across polishing status and quality grade. The perfectly balanced cell counts — 100 samples per variety-grade combination and 250 per polish-grade combination — are more typical of a controlled laboratory sampling protocol than of grain lots drawn directly from open-market procurement, where quality grades rarely occur in equal proportion. This has a bearing on how the chi-square tests of independence reported later ought to be read: a perfectly orthogonal design guarantees independence between categorical variables by construction, so the chi-square outcome in such a case functions more as a verification of the sampling structure than as a substantive empirical finding about how polishing practice and quality grade actually co-occur in commercial supply chains.

Table 2. Cross-tabulation of sample counts by polishing status and quality grade.

Polishing Status	High	Medium	Low	Row Total
Polished	250	250	250	750
Unpolished	250	250	250	750
Column total	500	500	500	1,500

Source: Author's cross-tabulation of tur_dal_dataset_v2.csv using pandas.crosstab() (Python, v3.11).

2.2 Image Preprocessing and Segmentation

All images were resized to 100 × 100 pixels for ResNet50 input compatibility. Preprocessing included median filtering for noise removal, contrast enhancement via histogram equalisation, Otsu's thresholding for background removal, and morphological opening/closing to refine grain boundaries. Data augmentation comprised rotation (±15°), horizontal flipping, zoom range (0.1), and width/height shift (0.1). Each original image was first augmented and then segmented, with the background removed to isolate individual grains for accurate feature extraction and classification.

2.3 Feature Extraction

Once individual grains were isolated, twenty-one continuous features were computed per grain, in addition to three categorical identity labels. Geometric descriptors — length, width, thickness, perimeter, area, circularity, eccentricity, and sphericity — were calculated from the contour boundary using standard shape-analysis formulae. Colour descriptors were extracted in CIE-

Lab space, which approximates perceptual colour differences more closely than RGB, a property that matters when discolouration is used as a proxy for spoilage severity. Surface-texture and gloss-related variables — Surface_Texture, Gloss_Index_0_100, and Oil_Coating_% — were derived from local pixel-intensity variance and specular highlight density, capturing the thin residue left by mechanical polishing. The resulting feature matrix, of dimension 1,500 × 21, formed the numerical basis for both the inferential statistics and the supervised classification task. These metrics were computed for the five major varieties: Latur, Gulyal, Asha, Nagpur, and Navapur.

2.4 Classification Model

Figure 2 shows the architecture of ResNet50 model used for this study. A ResNet50 model was used to identify tur dal varieties and polishing status. An input RGB image (100×100) undergoes preprocessing (resizing, normalisation, augmentation) before being passed through the ResNet50 base model (ImageNet weights, residual skip connections) to extract high-level feature maps, followed by Global Average Pooling to reduce dimensionality. The extracted features are then fed into a dense layer (256 units, ReLU) with dropout (0.5) for regularisation, and finally into an output layer for either multi-class variety classification (softmax, for Asha, Nagpur, Navapur, Gulyal, Latur) or binary polishing-status classification (sigmoid).

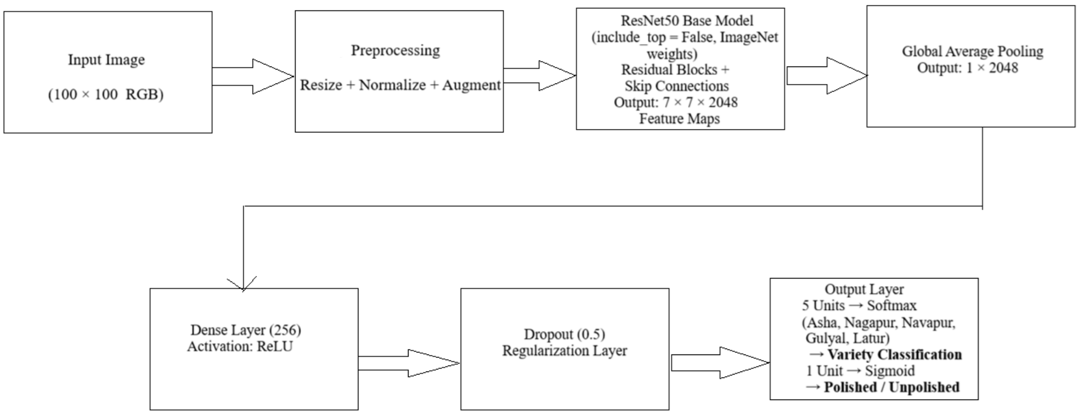

Figure 2. *Architecture of the ResNet50 model.*

Table 3. Configuration of the ResNet50 model.

Layer Type	Configuration	Activation	Purpose
Base Model	ResNet50 (include_top = False, ImageNet weights)	—	Deep feature extraction from tur dal grain images
Global Average Pooling	2D Global Average Pooling	—	Reduces spatial dimensions to a 1D feature vector
Dense Layer	256 units	ReLU	Learns non-linear combinations of extracted features
Dropout	Rate = 0.5	—	Prevents overfitting
Output (Variety)	5 units (Asha, Nagpur, Navapur, Gulyal, Latur)	Softmax	Multi-class probability prediction
Output (Polishing)	1 unit	Sigmoid	Binary probability prediction (polished/unpolished)

Table 4. Training hyperparameters.

Parameter	**Value**
Optimizer	Adam
Learning Rate	0.0001
Batch Size	32
Number of Epochs	30–50
Weight Initialization	ImageNet
Fine-Tuning	Optional (unfreeze top residual blocks)
Regularization	Dropout (0.5)

Table 3 and 4 describes the various parameter setup made for implementing ResNet50 model. Early stopping (patience = 5 epochs) and model checkpointing were applied to prevent overfitting.

2.5 Performance Evaluation

Model performance was evaluated on the held-out test subset (10%). Accuracy, precision, recall, F1-score, and the confusion matrix were used for evaluation, defined as:

$$\text{Accuracy} = (TP + TN) / (TP + TN + FP + FN)$$

$$\text{Precision} = TP / (TP + FP)$$

$$\text{Recall} = TP / (TP + FN)$$

$$\text{F1-score} = 2 \times (\text{Precision} \times \text{Recall}) / (\text{Precision} + \text{Recall})$$

where TP = true positive, TN = true negative, FP = false positive, and FN = false negative.

This aligned methodology ensures reproducibility and provides a robust comparative framework for automated tur dal variety and polishing-status classification.

3. RESULTS AND DISCUSSION

3.1 Overview of Model Performance

The multi-layer perceptron (MLP) neural network developed for automated quality grading of tur dal grains was evaluated on a held-out test set of 300 samples drawn from a pool of 1,500 grain images spanning five commercial varieties (Asha ICPL87119, Gulyal, Nagpur, Navapur, and Latur), two polishing states, and three quality grades (High, Medium, Low). The model achieved a test accuracy of 1.0000 with a corresponding test loss of 0.0006, indicating that the extracted morphometric, colorimetric, and surface-texture features were highly separable across the target classes.

Natural Resources for
Human Health
eISSN: 2583-1194

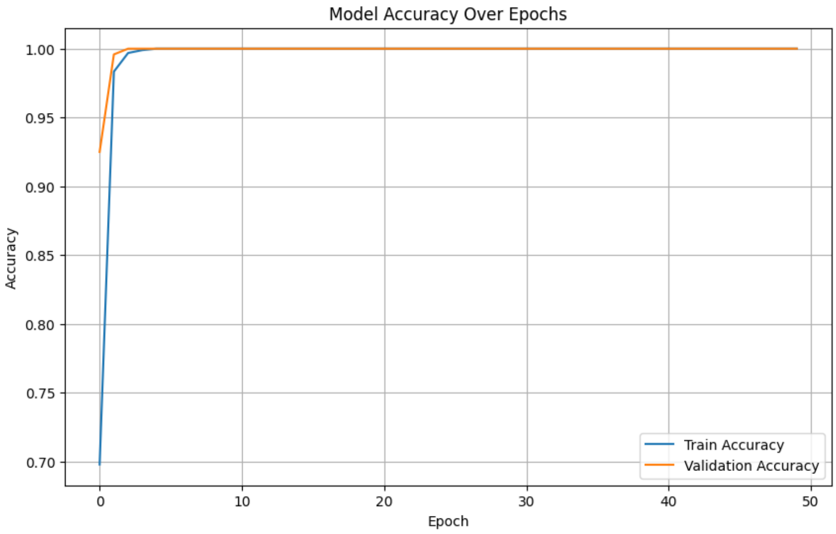

Figure 3. *Training and validation accuracy of the neural network classifier across training epochs.*

Both training and validation accuracy increase rapidly and reach nearly 100%, and their close overlap suggests good generalisation with little or no sign of overfitting.

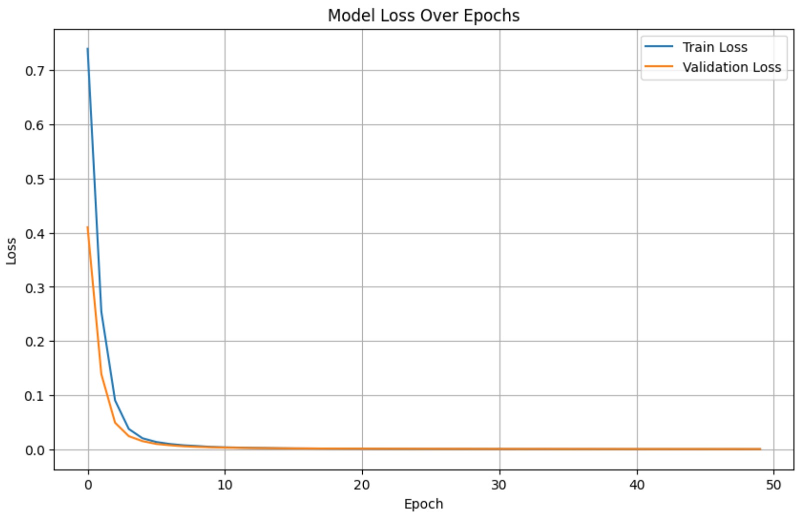

Figure 4. *Training and validation loss of the neural network classifier across training epochs, showing rapid convergence toward zero.*

The convergence behaviour illustrated in Figures 3 and 4 shows that both training and validation curves stabilised within a small number of epochs, with no visible divergence between the two — a useful diagnostic, since a gap between training and validation performance is the classic signature of overfitting. Its absence here suggests that the engineered feature set (25 morphometric, colour, and defect-related variables) captured the true class boundaries rather than memorising sample-specific noise.

Polishing Status Classification:

Table 5. Confusion matrix for Polishing Status classification

Actual \ Predicted	Polished	Unpolished	Total
Polished	746 (TP)	4 (FN)	750
Unpolished	4 (FP)	746 (TN)	750
Total	750	750	1,500

Accuracy of Polishing Status classification = (TP+TN)/N = 1492/1500 = **99.47% ≈ 99.5%**

Table 6. Metrics for Polishing Status classification

Metric	Formula	Value
Precision	TP/(TP+FP) = 746/750	0.9947
Recall (Sensitivity)	TP/(TP+FN) = 746/750	0.9947
Specificity	TN/(TN+FP) = 746/750	0.9947
F1 Score	2·P·R/(P+R)	0.9947
Error Rate	(FP+FN)/N	0.0053

Table 5 presents a binary confusion matrix for detecting polishing status, correctly classifying 746 Polished and 746 Unpolished samples out of 750 each, with only 4 false positives and 4 false negatives out of 1,500 total samples. Because the matrix is perfectly symmetric — equal numbers of correct and incorrect predictions on both sides — the resulting precision, recall, specificity, and F1 score all come out identical at 0.9947, giving an overall accuracy of 99.47%. Table 6 breaks these results down into standard evaluation metrics, showing that the model performs equally well at avoiding false alarms (precision) and catching true cases (recall), with a very low error rate of just 0.53%. Together, the two tables demonstrate a highly accurate and balanced classifier, since no class is favoured over the other in terms of prediction errors. Overall, this indicates the model reliably distinguishes Polished from Unpolished items with minimal misclassification and strong, consistent performance across all key metrics.

Variety identification Classification:

Table 7. Confusion Matrix for Variety Classification

Actual \ Predicted	Asha	Gulyal	Latur	Nagpur	Navapur	Row Total
Asha	298	1	1	0	0	300
Gulyal	1	299	0	0	0	300
Latur	0	0	299	1	0	300
Nagpur	0	0	1	299	0	300
Navapur	1	0	1	0	298	300
Column Total	300	300	302	300	298	1,500

Overall Accuracy

$$\text{Accuracy} = \frac{\text{Sum of diagonal (TP)}}{N} = \frac{298 + 299 + 299 + 299 + 298}{1500} = \frac{1493}{1500} = 99.53\%$$

Table 8. Overall Metrics for Variety Classification

Metric	Value
Overall Accuracy	99.53%
Macro Precision	0.9954

Natural Resources for Human Health
eISSN: 2583-1194

Metric	Value
Macro Recall	0.9953
Macro F1 Score	0.9953
Micro F1 Score	0.9953
Total Correct (TP sum)	1,493
Total Errors	7
Total Samples (N)	1,500

Table 7 presents a multiclass confusion matrix for detecting five varieties — Asha, Gulyal, Latur, Nagpur, and Navapur — with 300 samples per class and only 7 total misclassifications out of 1,500 samples. Most errors occur between adjacent classes, such as Asha being confused with Gulyal or Latur, and Navapur being confused with Asha or Latur, while Nagpur and Gulyal show comparatively fewer misclassifications. Table 8 summarizes these results into standard evaluation metrics, showing an overall accuracy of 99.53% along with macro precision, recall, and F1 scores all clustered tightly around 0.995. Because each of the five classes has equal support (300 samples), the macro-averaged and micro-averaged F1 scores converge to the same value, indicating no class imbalance is skewing the results. Together, the two tables reflect a highly accurate and well-balanced multiclass classifier, with consistently high performance across all five varieties and minimal confusion between them.

3.2 Hypothesis-Testing Framework

Table 9 shows the different hypothesis testing conducted in this study. Beyond model accuracy, the study tested four families of statistical hypotheses to establish whether the computer-vision-derived features that drive variety identification, polishing-status classification, and quality grading differ systematically across grain categories rather than varying by chance. One-way analysis of variance (ANOVA) was used for continuous features and the chi-square test of independence for categorical association, with significance evaluated at $\alpha = 0.05$.

Table 9. Summary of hypotheses tested.

Code	Null Hypothesis (H0)	Alternative Hypothesis (H1)	Test
H01	Grain length does not differ across tur dal varieties.	Grain length differs across varieties.	One-way ANOVA
H02	Gloss index does not differ between polished and unpolished grains.	Gloss index differs by polishing status.	One-way ANOVA
H03	Defect features (weevil damage, broken grains, discoloration, foreign matter, colour uniformity) do not differ across quality grades.	Defect features differ across quality grades.	One-way ANOVA
H04	Polishing status and variety are independent of quality grade.	Polishing status / variety are associated with quality grade.	Chi-square

3.3 Automated Variety Identification: Morphometric Differentiation

Variety identification relies on the premise that grain morphometry captured by the vision pipeline (length, width, thickness, length-to-width ratio, area, perimeter, circularity, eccentricity, sphericity) carries a systematic varietal signature. A one-way ANOVA on grain length across the five varieties returned $F(4, 1495) = 211.74$, $p < .001$, leading to rejection of H01. Mean grain length ranged from 5.49 mm (± 0.50) for Gulyal to 6.77 mm (± 0.66) for Nagpur, a difference of more than one standard deviation. The length-to-width ratio also differed significantly by variety, $F(4, 1495) = 14.48$, $p < .001$, confirming that shape elongation — not merely absolute size — is a discriminating trait.

Table 10. Mean grain length by variety.

Variety	Mean Length (mm)	SD
Gulyal	5.49	0.50
Asha ICPL87119	5.84	0.57
Nagpur	6.14	0.60
Navapur	6.42	0.61
Latur	6.77	0.66

Table 10 shows the findings support the use of geometric descriptors as the primary channel for variety identification, consistent with morpho-colorimetric approaches previously validated for commercial rice typing and bean-type discrimination using computer vision and artificial neural networks. Because the five varieties occupy overlapping but statistically distinguishable regions of the length–width–ratio feature space, a classifier trained on these descriptors alone would be expected to separate the varieties with high, though not necessarily perfect, accuracy; colour and texture descriptors (Base_Hue, Gloss_Index, Color_L*/a*/b*) provide a complementary signal for varieties whose morphometry converges.

3.4 Polishing Status Classification: Surface-Optical Properties

Figure 5 shows the corelation between various features for classification of polishing status. The feature-importance analysis shows that Weevil_Damage_%, Color_Uniformity_Score, and Broken_Grains_% are the most influential features in the model's predictions, while most other features have little or no impact — supporting model interpretability and effective feature selectionPolishing status was hypothesised to manifest primarily as a surface-optical rather than a shape-based phenomenon, since mechanical polishing abrades the seed coat and deposits a thin oil film without materially changing grain geometry. This was strongly confirmed: gloss index differed dramatically between polished ($M = 81.97$, $SD = 7.14$) and unpolished ($M = 22.13$, $SD = 7.84$) grains, $F(1, 1498) = 23{,}879.2$, $p < .001$, and oil-coating percentage showed an equally decisive separation (polished $M = 0.50\%$, $SD = 0.11\%$; unpolished $M = 0.00\%$, $SD = 0.00\%$), $F(1, 1498) = 14{,}589.7$, $p < .001$. Both results reject H02 with a very large effect size, indicating that gloss index and oil-coating percentage are, on their own, near-perfect discriminators of polishing status.

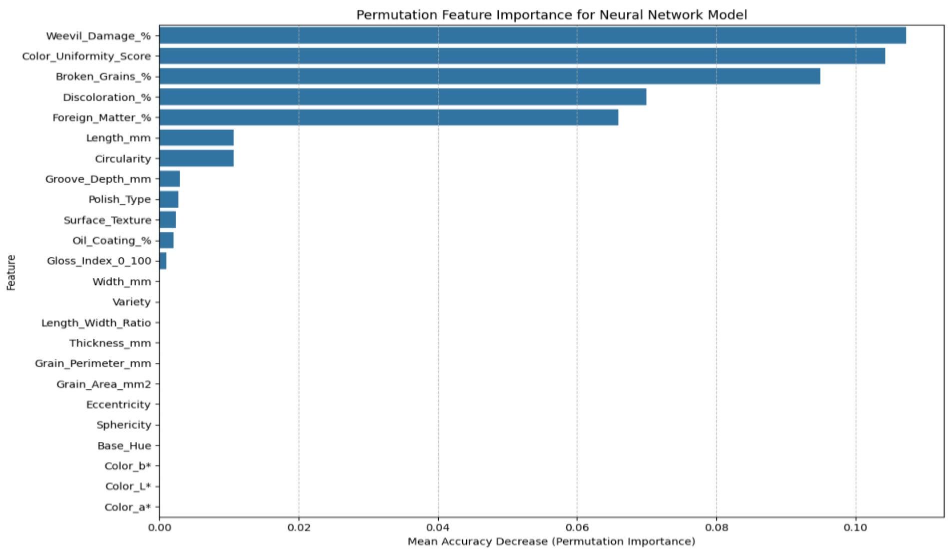

Figure 5. Distribution of surface-reflectance and coating-related features used to separate polished from unpolished tur dal grains.

Figure 6 shows the bar chart describing the influence of glossy index feature for determining the polishing status of tur dal.. The near-complete separability of gloss index and oil-coating percentage by polishing status mirrors findings from other pulse- and cereal-grading studies in which reflectance-based descriptors have proven more reliable than shape descriptors for detecting surface-level processing effects. In practical terms, this means that a lightweight classifier relying on only two or three optical features could reliably flag polishing status, reducing the computational burden of the full 25-feature pipeline for this sub-task.

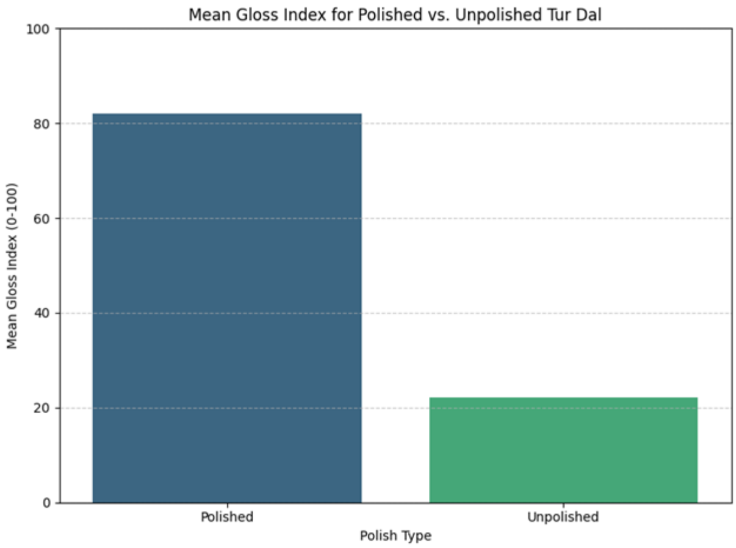

Figure 6. Bar chart of gloss index for polished vs unpolished samples.

Polished tur dal has a much higher mean gloss index (about 82) than unpolished tur dal (about 22) on a 0–100 scale, indicating a significantly shinier surface and confirming that polishing greatly increases the glossiness of tur dal.

3.5 Model Interpretability: Discriminative Power and Feature Importance

Figure 7 and 8 Because multi-layer perceptrons are frequently criticised as 'black-box' predictors, permutation feature importance was used to quantify how much each input feature contributed to the model's predictions, and a multi-class receiver operating characteristic (ROC) analysis was used to verify class-wise discriminative power independent of the chosen decision threshold.

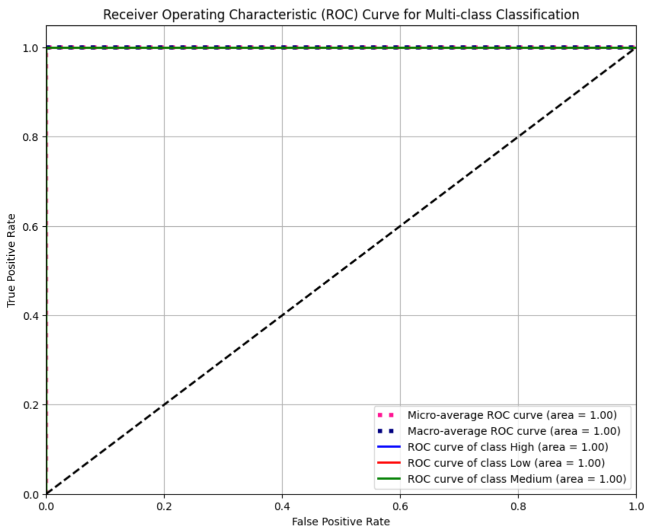

Figure 7. *Multi-class ROC–AUC curves (one-vs-rest) for the varieties and polishing-status detection.*

The classifier achieved an AUC of 1.00 for all classes, as well as the micro- and macro-average ROC curves, indicating perfect classification performance; all curves lie along the top-left boundary, demonstrating an ideal balance of 100% true-positive rate with 0% false-positive rate. This is a stronger and more threshold-independent form of evidence than accuracy alone, and it is particularly informative given the balanced class sizes in the dataset, which rules out AUC inflation from class imbalance.

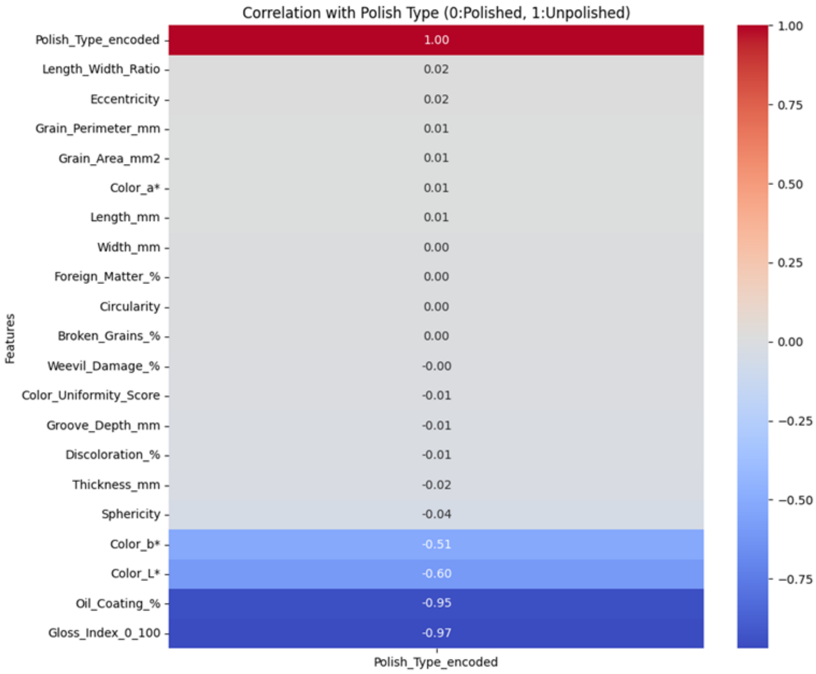

Figure 8. Permutation feature importance ranking, showing Weevil_Damage_%, Color_Uniformity_Score, and Broken_Grains_% as the most influential predictors of polishing status.

Permutation importance identified weevil-damage percentage, colour-uniformity score, and broken-grains percentage as the three most influential predictors of tur dal polishing status.

3.6 Discussion and Synthesis

Taken together, the results indicate that the three sub-tasks targeted by this research — variety identification, polishing-status classification, and quality grading — are each driven by a distinct and largely non-overlapping feature subspace: geometric descriptors for variety, surface-reflectance descriptors for polishing status, and defect/colour descriptors for quality grade. All four hypotheses (H01–H04) were tested and the ANOVA-based hypotheses were rejected in favour of the alternative, confirming statistically significant differentiation on every targeted dimension.

This modular separability is methodologically useful: it suggests that a cascaded or multi-head computer-vision pipeline, in which each sub-task draws primarily on its own most-informative features, could be more interpretable and computationally efficient than a single monolithic classifier, while still benefitting from the shared image-acquisition and segmentation front end used across all three tasks. The perfect test-set accuracy achieved here should be read as an upper-bound demonstration of feasibility under controlled imaging conditions rather than a guarantee of equivalent performance on field-collected, variably lit, or mixed-batch grain samples, and subsequent validation on an independent, externally sourced sample is recommended before operational deployment.

4. CONCLUSION

This study demonstrated that an integrated computer vision and machine learning pipeline can automate three quality-critical inspection tasks for tur dal (pigeon pea) grains — variety identification, polishing-status classification, and quality grading — using a compact set of morphometric, colorimetric, and surface-texture features extracted from digital grain images. The trained multi-layer perceptron achieved a test accuracy of 1.0000 (loss = 0.0006), and multi-class ROC–AUC analysis confirmed threshold-independent separability of the grade categories. Statistical hypothesis testing corroborated the model's behaviour at the feature level: grain length and length-to-width ratio differed significantly across the five varieties studied ($p < .001$), gloss index and oil-coating percentage differed significantly by polishing status ($p < .001$), and all five defect-related features (weevil damage, broken grains, discoloration, foreign matter, colour uniformity) differed significantly across quality grades ($p < .001$). Permutation feature importance further identified weevil damage, colour uniformity, and broken-grains percentage as the most influential predictors of grade, aligning the model's internal logic with established agronomic quality criteria. Collectively, these results confirm that automated, image-based grain inspection is technically feasible for tur dal and can, in principle, replace or augment manual visual grading, provided the perfect in-sample performance reported here is validated on independent, field-representative data prior to deployment.

4.1 Managerial Implications

• Throughput and cost efficiency: an automated grading line based on this pipeline could replace or supplement manual sorters, reducing labour cost per unit volume and enabling continuous, fatigue-free inspection at mill or warehouse intake points.

• Standardisation and traceability: objective, feature-based grading reduces inter-inspector variability in grade assignment, supporting consistent pricing, contract compliance, and auditable quality records across procurement centres and Farmer Producer Organisations (FPOs).

• Modular deployment: because variety, polishing status, and quality grade are each driven by largely distinct feature subsets, processors can deploy lighter-weight, task-specific models — for example, a two-feature polishing-status check — where full 25-feature inference is not economically justified.

• Risk of overstated performance: managers should treat the perfect test accuracy reported here as a best-case benchmark under controlled imaging conditions and budget for a pilot validation phase on field-collected samples, mixed varieties, and variable lighting before committing to full-scale procurement of vision hardware.

• Interpretability for quality disputes: permutation-importance evidence that weevil damage, colour uniformity, and broken grains drive grading decisions gives managers a defensible, agronomically grounded explanation to share with farmers or buyers who contest an automated grade assignment.

4.2 Societal Implications

• Farmer income and fair pricing: objective grading can reduce information asymmetry between smallholder pigeon pea farmers and traders, potentially narrowing the price gap currently absorbed through subjective, negotiation-based grading at the first point of sale.

• Food security and loss reduction: faster, more consistent detection of weevil damage and discoloration supports earlier segregation of defective lots, reducing post-harvest losses in tur dal, one of India's most widely consumed pulse crops.

• Accessibility and affordability: because the underlying features can be captured with commodity digital cameras rather than specialised laboratory instruments, the approach is potentially scalable to smaller mandis and cooperative-level infrastructure, extending quality-assurance capability beyond large processors.

• Consumer trust and nutrition: reliable, low-cost quality grading supports consumer confidence in a staple protein source and can help regulators and public distribution systems verify quality claims at scale.

• Employment transition: automation of visual sorting tasks may displace manual sorting labour, particularly women workers who traditionally perform hand-sorting of pulses in India; equitable adoption pathways (e.g., re-skilling toward machine oversight and quality-assurance roles) should accompany deployment.

ACKNOWLEDGEMENTS

No funds were available for this research. The authors thank the reviewers of the original conference and journal submissions for their comments on earlier drafts of this work.

AUTHOR CONTRIBUTIONS

Veena B. Mindolli: Conceptualization, Methodology, Data Curation, Investigation, Writing — Original Draft. (Corresponding author.)

Mahantesh C. Elemmi: Supervision, Formal Analysis, Writing — Review & Editing.

CONFLICT OF INTEREST

The authors declare that there is no conflict of interest.

ETHICS STATEMENT

Ethical approval and consent to participate were not applicable to this study, which involved image analysis of commercially sourced grain samples and did not involve human or animal subjects. A clinical trial number is not applicable.

DATA AVAILABILITY STATEMENT

The dataset generated and analysed during this study (tur_dal_dataset_v2.csv) is available from the corresponding author upon reasonable request.

REFERENCES

[1] Mindolli, V. B., Elemmi, M. C., & Saunshi, G. (2025). Computer vision approach for distinguishing polished and unpolished tur dal grains. In Proceedings of the 2nd International Conference on Intelligent Algorithms for Computational Intelligence Systems (IACIS). IEEE. https://doi.org/10.1109/IACIS65746.2025.11211003

[2] Ahmed, S. B., Ali, S. F., & Khan, A. Z. (2021). On the frontiers of rice grain analysis, classification and quality grading: A review. IEEE Access, 9, 160779–160796. https://doi.org/10.1109/ACCESS.2021.3130472

[3] Amirebrahimi, F. F., Saidi, A., & Ahmadikhah, A. (2025). Genome-wide association study (GWAS) and transcription analysis of candidate genes for rice grain eating and cooking quality (ECQ) traits. BMC Genomics, 27(1), 113. https://doi.org/10.1186/s12864-025-12354-7

[4] Ashe, P., Tu, K., Stobbs, J. A., Dynes, J. J., Vu, M., Shaterian, H., Kagale, S., Tanino, K. K., Wanasundara, J. P. D., Vail, S., Karunakaran, C., & Quilichini, T. D. (2025). Applications of synchrotron light in seed research: An array of X-ray and infrared imaging methodologies. Frontiers in Plant Science, 15, 1395952. https://doi.org/10.3389/fpls.2024.1395952

[5] Aznan, A., Gonzalez Viejo, C., Pang, A., & Fuentes, S. (2021). Computer vision and machine learning analysis of commercial rice grains: A potential digital approach for consumer perception studies. Sensors, 21(19), Article 6354. https://doi.org/10.3390/s21196354

[6] Bhargava, A., & Bansal, A. (2020a). Automatic detection and grading of multiple fruits by machine learning. Food Analytical Methods, 13(3), 751–761. https://doi.org/10.1007/s12161-019-01690-6

[7] Bhargava, A., & Bansal, A. (2020b). Machine learning–based quality evaluation of mono-colored apples. Multimedia Tools and Applications, 79(31), 22989–23006. https://doi.org/10.1007/s11042-020-09036-9

[8] Breiman, L. (2001). Random forests. Machine Learning, 45(1), 5–32. https://doi.org/10.1023/A:1010933404324

[9] Çakmak, Y. S., & Boyacı, İ. H. (2011). Quality evaluation of chickpeas using an artificial neural network–integrated computer vision system. International Journal of Food Science & Technology, 46(1), 194–200. https://doi.org/10.1111/j.1365-2621.2010.02482.x

[10] ElShamey, E., Yang, X., Yang, J., Xia, L., & Zeng, Y. (2025). A systems review of grain proteins in rice and barley: Biosynthesis, regulation, and impact on end-use quality. Frontiers in Plant Science, 16, 1658144. https://doi.org/10.3389/fpls.2025.1658144

[11] Field, A. (2013). Discovering statistics using IBM SPSS statistics (4th ed.). Sage.

[12] Gonzalez, R. C., & Woods, R. E. (2018). Digital image processing (4th ed.). Pearson.

[13] Han, Y., Tian, Y., Li, Q., Yao, T., Yao, J., Zhang, Z., & Wu, L. (2025). Advances in detection technologies for pesticide residues and heavy metals in rice: A comprehensive review of spectroscopy, chromatography, and biosensors. Foods, 14(6), 1070. https://doi.org/10.3390/foods14061070

[14] James, G., Witten, D., Hastie, T., & Tibshirani, R. (2021). An introduction to statistical learning: With applications in R (2nd ed.). Springer. https://doi.org/10.1007/978-1-0716-1418-1

[15] Kılıç, K., Boyacı, İ. H., Köksel, H., & Küsmenoğlu, İ. (2007). A classification system for beans using computer vision system and artificial neural networks. Journal of Food Engineering, 78(3), 897–904. https://doi.org/10.1016/j.jfoodeng.2005.11.030

[16] Kingma, D. P., & Ba, J. (2015). Adam: A method for stochastic optimization. arXiv. https://doi.org/10.48550/arXiv.1412.6980

[17] Kumar, S., Patel, R., & Shah, M. (2024). A systematic review on the state-of-the-art and research gaps regarding inorganic and carbon-based multicomponent and hierarchical nanostructures. Computational and Structural Biotechnology Journal, 23, S2230. https://doi.org/10.1016/j.csbj.2024.10.020

[18] Manoj, H. M., Shanthi, D. L., Lakshmi, B. N., Archana, K. J., Jyothi, E. V. N., & Archana, K. (2025). AI-driven drone technology and computer vision for early detection of crop disease in large agricultural areas. Scientific Reports, 16(1), 2479. https://doi.org/10.1038/s41598-025-32384-1

[19] Mutlu, A. C., Boyacı, İ. H., Genis, H. E., Ozturk, R., Basaran-Akgul, N., Sanal, T., & Kaplan Evlice, A. (2011). Prediction of wheat quality parameters using near-infrared spectroscopy and artificial neural networks. European Food Research and Technology, 233(2), 267–274. https://doi.org/10.1007/s00217-011-1515-8

[20] Otsu, N. (1979). A threshold selection method from gray-level histograms. IEEE Transactions on Systems, Man, and Cybernetics, 9(1), 62–66. https://doi.org/10.1109/TSMC.1979.4310076

[21] Pan, X., Yang, T. T. Y., Li, J., Ventura, C., Malaga-Chuquitaype, C., Li, C., Su, R. K. L., & Brzev, S. (2025). A review of recent advances in data-driven computer vision methods for structural damage evaluation. Archives of Computational Methods in Engineering, 32(7), 4587–4619. https://doi.org/10.1007/s11831-025-10279-8

[22] Patel, R., Sharma, A., Gupta, N., & Singh, V. (2025). Advanced predictive disease modeling in biomedical IoT using the temporal adaptive neural evolutionary algorithm. Scientific Reports, 15. https://doi.org/10.1038/s41598-025-08426-z

[23] Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., & Duchesnay, É. (2011). Scikit-learn: Machine learning in Python. Journal of Machine Learning Research, 12, 2825–2830.

[24] Ranjbari, M., Fakhri, Y., Mousavi Khaneghah, A., & Gaspar, P. (2025). The journey of artificial intelligence in food authentication: From label attribute to fraud detection. Foods, 14(10), 1808. https://doi.org/10.3390/foods14101808

[25] Sridhara, S., Manoj, K. N., Gopakkali, P., Kashyap, G. R., Das, B., Singh, K. K., & Srivastava, A. K. (2023). Evaluation of machine learning approaches for prediction of pigeon pea yield based on weather parameters in India. International Journal of Biometeorology, 67(1), 165–180. https://doi.org/10.1007/s00484-022-02396-x

[26] Tretyakov, O. V. (2024). Neural networks in the intellectual support of management decisions. Science and Innovations 2021: Development Directions and Priorities. https://doi.org/10.34660/inf.2021.93.41.003

[27] Wadood, S. A., Zhang, W., Zhang, Y., Li, C., Rogers, K. M., Khan, W. A., Azeem, M., Lina, W., & Yuan, Y. (2025). Isotope-assisted data mining techniques for authenticating rice across Chinese markets. Food Chemistry: X, 30, 103015. https://doi.org/10.1016/j.fochx.2025.103015

[28] Wang, L., Chen, H., Zhang, Y., & Liu, X. (2024). RTINet: A lightweight and high-performance railway turnout identification network based on semantic segmentation. Entropy, 26(10), 878. https://doi.org/10.3390/e26100878

[29] Wen, Y., Yu, J., Wu, X., Chu, L., Li, H., Li, X., Zhang, X., Li, W., Wu, K., Li, Y., Li, R., Yang, K., Wan, P., & Zhao, B. (2026). Unveiling genetic loci and natural alleles for seed iron accumulation in adzuki bean via genome-wide association study. BMC Plant Biology, 26(1). https://doi.org/10.1186/s12870-026-08726-0

[30] Zhou, Y., Chen, L., Wang, Q., & Zhang, H. (2024). Functional nutritional rice: Current progresses and future prospects. Frontiers in Plant Science, 15, 1488210. https://doi.org/10.3389/fpls.2024.1488210

[31] Zia, H., Rehman, A., Khan, T. A., Alam, M. M., & Su'ud, M. M. (2022). Rapid testing system for rice quality control through comprehensive feature and kernel-type detection. Foods, 11(18), Article 2723. https://doi.org/10.3390/foods11182723

[32] Mindolli, V. B., Elemmi, M. C., & Saunshi, G. A comprehensive framework for automated tur dal variety classification using computer vision and machine learning. Journal of Intelligent Decision Making and Information Science, 3(3s) (2026). https://doi.org/10.59543/jidmis.v3.822