DEEP LEARNING FOR AGRICULTURE

Comparison of the Performance of ResNet-50 and ShuffleNet Deep Learning Algorithms in Classification of Hazelnut Fruit

Özgür Tomak, Mehmet Cem Dikbaş, Sercan Külcü

Giresun University, Computer Engineering / Electrical and Electronics Engineering, Türkiye

11th International Azerbaijan Congress on Life, Engineering, Mathematical, and Applied Sciences | June 21-22, 2025, Baku

Introduction & Motivation

- AI-based hazelnut classification improves efficiency and quality control compared to manual sorting.
- Clean, normalized image datasets and appropriate preprocessing (resizing, sharpening) improve classifier performance.
- This study proposes an automatic classification method and directly compares two CNN architectures — ResNet-50 and ShuffleNet — on the same hazelnut dataset.
- Goal: identify which model offers the better balance of accuracy and speed for real-time hazelnut sorting.

STUDY AT A GLANCE

15,770 hazelnut images (10 views × 1,577 nuts)

3 classes: good, bad, and inner hazelnuts

2 CNN architectures compared head-to-head

Related Work

1

Bayrakdar et al. (2015)

Classified Tombul, sivri, and badem hazelnut types using shape/size features (diameter, radius, area) — 84% accuracy.

2

Caner et al. (2020)

Classified hazelnut cultivars with shell using Bagging/Random Forest (83%) and DL4J deep learning (71%).

3

Keles & Taner (2021)

Used artificial neural networks and discriminant analysis on physical/mechanical/optical features of 11 cultivars — up to 92.7% accuracy.

4

Gap Addressed

Prior work targets cultivar or size classification via handcrafted features; this study directly benchmarks two modern CNNs (ResNet-50 vs. ShuffleNet) for fruit-quality classification.

Two Architectures Compared

ResNet-50

- 177 layers total — a 50-layer Residual Network.
- Pretrained on ImageNet (1M+ images, 1,000 object categories); won the 2015 ImageNet Large Scale Visual Recognition Challenge.
- 224×224 input size.
- Residual connections: block input is added to block output — H(x) = F(x) + x — easing training of very deep networks.
- Batch normalization after each convolution stabilizes training; Global Average Pooling precedes the final fully connected layer.

ShuffleNet

- Purpose-built for devices with very limited processing power (e.g., mobile).
- Reduces parameters and computation via pointwise group convolution and channel shuffle.
- Core block uses depthwise separable convolutions for computational efficiency.
- Channel shuffle mixes information across feature groups, improving overall performance.
- Stride-2 units add 3×3 average pooling on the shortcut; channel concatenation replaces addition to widen feature maps cheaply.

Dataset & Preprocessing

- Images captured with a 12-megapixel rear camera at 3024×3024 resolution.
- Each image contains a single hazelnut, expert-labeled into one of three classes.
- 519 bad hazelnuts, 535 good hazelnuts, 523 inner hazelnuts.
- 10 images per hazelnut from different angles → 15,770 total images.
- Preprocessing: cropped to 1000×1000 px, then resized to 224×224 px for network input.
- Split: 70% training, 30% testing (randomly shuffled); environment: MATLAB.

Good Hazelnuts

535 nuts

Bad Hazelnuts

519 nuts

Inner Hazelnuts

523 nuts

15,770 total images (10 views each)

Results: Training & Test Accuracy

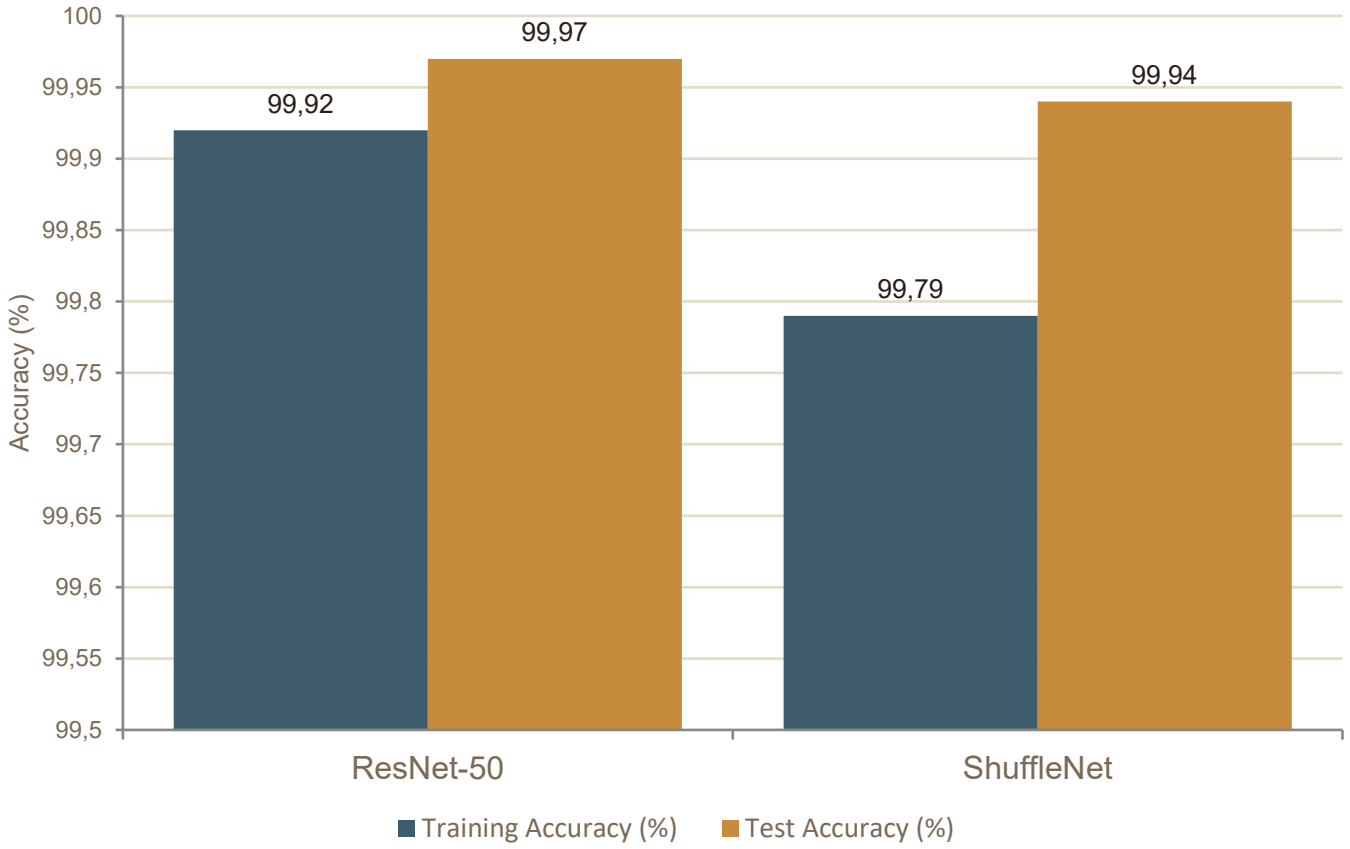

HIGHEST ACCURACY

ResNet-50

99.97%

Test Accuracy

ShuffleNet trails by just 0.03 pts (99.94%) — but with a far smaller, faster model, better suited to real-time use.

Confusion Matrix Metrics

Model	Precision (Sensitivity)	Specificity	F1-Score
ResNet-50	**0.9997**	**1.0000**	**0.9999**
ShuffleNet	0.9997	0.9994	0.9997

Table 1 — Complexity-matrix metrics for ResNet-50 and ShuffleNet.

ResNet-50 edges out ShuffleNet on every metric, but the margins are marginal — both models classify hazelnut quality with near-perfect reliability.

THE TRADE-OFF

177

ResNet-50 layers — a large, accuracy-first architecture.

ShuffleNet: fewer parameters, lower computational load — purpose-built for fast, resource-limited inference.

Why Sorting Matters

1

Harvest & Capture

Hazelnuts are photographed from multiple angles before storage.

2

Classify

CNN model (ResNet-50 or ShuffleNet) labels each nut as good, bad, or inner.

3

Separate

Bad and inner hazelnuts are removed from the batch.

4

Store

Only good-quality hazelnuts proceed to storage, preventing spoilage and yield loss.

Unsorted bad or inner hazelnuts reduce yield and can spoil during storage — automated classification prevents both.

Discussion

1 ResNet-50: Accuracy Leader

Residual connections and depth (177 layers) yield the highest test accuracy (99.97%) and perfect specificity (1.0).

2 ShuffleNet: Speed Leader

Lightweight design with far fewer parameters trains and infers faster, at a near-identical 99.94% test accuracy.

3 Practical Recommendation

Given the closeness of results and the importance of fast inference in real-time sorting lines, ShuffleNet is recommended as the better overall choice.

4 Limitation & Future Work

Both models were evaluated on a single dataset (15,770 images, 3 classes); broader validation across cultivars and lighting conditions would strengthen the comparison.

Conclusion

- ResNet-50 achieves the highest test accuracy (99.97%), narrowly ahead of ShuffleNet (99.94%).
- Confusion-matrix metrics confirm both models classify hazelnut quality with near-perfect reliability.
- ShuffleNet trains and infers faster due to its lightweight design — a decisive advantage for real-time sorting applications.
- Given the closeness in accuracy and the importance of speed in production lines, ShuffleNet is recommended as the better overall choice for hazelnut fruit classification.

ShuffleNet

Recommended Choice

99.94%
test accuracy

Faster inference,
smaller footprint

Ideal for real-time
production sorting

Contact: sercan.kulcu@giresun.edu.tr | ORCID: 0000-0002-4871-709X