

Date of publication xxxx 00, 0000, date of current version xxxx 00, 0000.

Digital Object Identifier 10.1109/ACCESS.2024.0429000

# Deep Learning-Based Meat Freshness Detection with Segmentation and OOD-Aware Classification

HUTAMA ARIF BRAMANTYO<sup>1</sup>, MUKARRAM ALI FARIDI<sup>1</sup>, RUI CHEN<sup>2</sup>, CLARISSA HARRIS<sup>2</sup> and YIN SUN<sup>1</sup>

<sup>1</sup>Electrical & Computer Engineering Department, College of Engineering, Auburn University, Auburn, AL 36849 USA

<sup>2</sup>Department of Agricultural and Environmental Sciences, College of Agriculture, Environment & Nutrition Sciences, Tuskegee University, Tuskegee, AL 36088 USA

Corresponding author: Rui Chen (e-mail: rchen@tuskegee.edu).

This research was supported in part by the NSF under grant no. CNS-2239677 and supported by USDA's the Agriculture Economics and Rural Communities Program (award no. 2023-69006-40213), Student Success and Workforce Development/1890 Universities Foundation (award no: 2021-38427-34837/FY22-SSWD-TU-34837), and Evans-Allen Program from the U.S. Department of Agriculture's National Institute of Food and Agriculture.

**ABSTRACT** In this study, we present a meat freshness classification framework from Red–Green–Blue (RGB) images that supports both packaged and unpackaged meat datasets. The system classifies four in-distribution (ID) meat classes and uses an out-of-distribution (OOD)-aware abstention mechanism that flags low-confidence samples as *No Result*. The pipeline combines U-Net-based segmentation with deep feature classifiers. Segmentation is used as a preprocessing step to isolate the meat region and reduce background, producing more consistent inputs for classification. The segmentation module achieved an Intersection over Union (IoU) of 75% and a Dice coefficient of 82%, producing standardized inputs for the classification stage. For classification, we benchmark five backbones: Residual Network-50 (ResNet-50), Vision Transformer-Base/16 (ViT-B/16), Swin Transformer-Tiny (Swin-T), EfficientNet-B0, and MobileNetV3-Small. We use nested 5 × 3 cross-validation (CV) for model selection and hyperparameter tuning. On the held-out ID test set, EfficientNet-B0 achieves the highest accuracy (98.10%), followed by ResNet-50 and MobileNetV3-Small (both 97.63%) and Swin-T (97.51%), while ViT-B/16 is lower (94.42%). We additionally evaluate OOD scoring and thresholding using standard OOD metrics and sensitivity analysis over the abstention threshold. Finally, we report on-device latency using TensorFlow Lite (TFLite) on a smartphone, highlighting practical accuracy–latency trade-offs for future deployment.

**INDEX TERMS** Computer Vision, Deep Learning, Food Safety, Image Segmentation and Classification, Meat Freshness Detection

## I. INTRODUCTION

Rising meat consumption continues to increase the need for practical tools that help monitor food quality and safety [1]. Spoilage not only contributes to foodborne illness risk, but also causes economic loss and avoidable food waste [2], [3]. In the United States, where meat consumption exceeds 224 pounds per capita annually [4], fast and reliable screening of meat freshness can support both consumer confidence and operational decision-making.

Conventional freshness assessment methods, including chemical analysis and sensory evaluation [5]–[7], are often invasive, costly, and dependent on expert judgment [8]. Although AMSA guidelines provide standardized visual criteria (for example, appearance and color cues) [9], applying them consistently in high-throughput, real-world settings remains

challenging. This motivates automated approaches that are objective, scalable, and suitable for everyday imaging conditions.

A range of sensor-based systems combined with machine learning, such as electronic noses, hyperspectral imaging, and gas sensor arrays, can provide precise freshness estimation [10]–[13]. However, their adoption is limited by specialized hardware requirements and controlled measurement conditions. In contrast, RGB-camera-based computer vision offers a low-cost and non-invasive alternative [14]–[16]. Prior studies have shown that convolutional neural networks (CNNs) can capture color and texture cues associated with spoilage [17]–[20], sometimes reporting very high accuracy in controlled environments. In practice, however, performance can degrade due to background clutter, tray artifacts,

glare, reflections, and packaging-related interference. These factors are especially common in packaged products.

Recent advances in Transformer-based vision models, such as the Vision Transformer (ViT) and Swin Transformer [21], [22], provide an alternative to CNN backbones by leveraging self-attention to model long-range dependencies. These architectures have shown strong results in challenging recognition tasks [23], [24]. Yet, their use in meat freshness recognition remains limited, and few works study robustness under practical imaging issues such as glare, reflections, and packaging occlusions.

Overall, several gaps remain. First, much of the literature focuses on unpackaged meat [17]–[20], while packaged meat introduces additional visual complexity, including reflections, film glare, and label occlusions. Second, the impact of explicit segmentation on downstream classification is not consistently examined under a unified evaluation protocol, especially when comparing CNN and Transformer backbones. Third, most existing pipelines assume all test inputs belong to known classes and do not incorporate out-of-distribution (OOD) awareness [25], [26]. This can lead to over-confident predictions on unfamiliar conditions.

In this work, we address meat freshness recognition for both packaged and unpackaged images using a two-stage pipeline. We first segment the meat region to reduce background and tray clutter, and then perform closed-set classification on four in-distribution (ID) meat classes. To handle uncertain or potentially out-of-distribution inputs, we additionally incorporate an OOD-aware abstention mechanism: when confidence is low, the system returns *No Result* rather than forcing a potentially overconfident prediction. We evaluate CNN and Transformer backbones under a nested cross-validation protocol, and we report on-device latency using TensorFlow Lite to highlight deployment trade-offs. The main contributions are summarized below:

- 1) We study a unified RGB-based pipeline that handles both packaged and unpackaged meat images and includes an OOD-aware abstention option (*No Result*) to avoid low-confidence decisions under challenging conditions.
- 2) We use U-Net-based segmentation to reduce background and tray clutter and to standardize the meat region prior to classification. We report segmentation agreement (IoU/Dice) under the same pseudo-labeling procedure used for training, and discuss how segmentation errors may affect downstream predictions.
- 3) We benchmark five backbones (ResNet-50, ViT-B/16, Swin-T, EfficientNet-B0, and MobileNetV3-Small) under a nested  $5 \times 3$  cross-validation protocol for consistent model selection and hyperparameter tuning, enabling a direct accuracy–efficiency comparison across CNN and Transformer.
- 4) We evaluate OOD scoring and thresholding using standard OOD metrics and sensitivity analyses, and we measure on-device latency with TensorFlow Lite to



FIGURE 1. Experiment setup for packaged meat image collection.

report realistic deployment cost.

The remainder of this paper is organized as follows: Section II reviews related studies; Section III describes the dataset and methodology; Section IV presents the experimental results and analysis; and Section V concludes the paper with future directions.

## II. RELATED WORK

Early vision-based approaches to meat freshness assessment often relied on handcrafted features extracted from RGB images. More recent work has shifted toward convolutional neural networks (CNNs) to learn discriminative cues automatically. For example, Ulucan *et al.* [17] reported high accuracy for binary classification of retail beef images under relatively controlled acquisition conditions. While the results are promising, many such studies focus on unpackaged meat and do not explicitly address practical issues such as glare, reflections, or packaging-related occlusions.

Several studies have extended freshness prediction beyond a binary setting. Bhargav *et al.* [18] explored multi-class freshness recognition and introduced a cost-sensitive evaluation tailored to misclassification severity. However, prior work typically evaluates on limited datasets or controlled setups, and it rarely studies uncertainty handling for unfamiliar imaging conditions. Subsequent work has also explored deeper architectures and optimization strategies. Abdelfattah *et al.* [19] combined a VGG19-based model with a meta-heuristic optimizer and reported strong performance for three-class prediction, while Ofori-Mensah *et al.* [20] compared CNN backbones for binary beef classification. Overall, these studies confirm that deep models can learn appearance cues associated with spoilage, but questions remain about generalization under real-world imaging variability and about mechanisms that make predictions safer when inputs fall outside the expected distribution.

Segmentation has been widely used as a preprocessing step in medical and natural-image analysis to isolate relevant regions, with U-Net being a common choice [27]. In food and agricultural imaging, lightweight models and detection-based pipelines have also been explored to support field deployment. For instance, MobileNet-based systems have been used for efficient visual diagnosis, and detection models such

as YOLO have been applied to related tasks such as ripeness estimation [28], [29]. These directions suggest that accuracy alone is not sufficient, and practical deployment constraints such as background clutter and inference-time cost should be considered.

A related but often overlooked aspect in meat freshness recognition is uncertainty and out-of-distribution (OOD) awareness. Confidence-based baselines such as maximum softmax probability (MSP) are commonly used [26], and stronger scoring methods such as ODIN [30] and energy-based scores [31] have been proposed to improve OOD separability. These methods motivate evaluating not only in-distribution accuracy, but also how a system behaves when inputs are ambiguous or collected under unseen conditions.

Motivated by these gaps, our work studies both packaged and unpackaged meat images using a two-stage pipeline that combines segmentation and classification, and includes an OOD-aware abstention option. Concretely, the model predicts one of four meat classes when confident and returns *No Result* when the input is uncertain or potentially out-of-distribution. We benchmark CNN, Transformer, and lightweight backbones under a nested cross-validation protocol and report on-device latency to ground the discussion in deployment constraints.

### III. DATA AND METHODS

#### A. DATASETS

##### 1) Data Collection

To cover a broad range of freshness conditions, we combined several open-access datasets of unpackaged meat and collected a complementary packaged-meat dataset. For unpackaged images, we used three open-access datasets:

- Meat Quality Assessment Dataset (Izmir University) [32], containing 1,896 images evenly split between fresh and spoiled meat.
- Meat Freshness Image Dataset [33], which provides 1,641 fresh and 624 spoiled images. Images labeled as “half-fresh” were grouped with the fresh category.
- LOCBEEF Dataset [34], comprising 3,268 images of Aceh beef (1,632 fresh, 1,636 spoiled) captured under various lighting and resolution settings to promote generalization.

We collected a new dataset in April 2024. Packaged meat products were sourced from local retailers in Auburn, Alabama (including Sam’s Club, Kroger, and Walmart). Then, we photographed the packaged meat in a controlled indoor environment using a smartphone mounted on a tripod, with temperature maintained at 25°C. Lighting intensity and direction were varied, including flash and ambient illumination, to reproduce realistic visual differences observed during spoilage. Differences across phone models and camera pipelines (e.g., auto white balance and exposure) are not explicitly evaluated here. To reduce dependence on a single setup, we varied illumination and viewing angles during capture, and we also used basic color/brightness augmentation



**FIGURE 2.** Representative examples of the *No Result* (irrelevant/OOD) set, including empty trays, plates, and background-only images.

during training.

The resulting dataset contained 3,734 packaged meat images (1,154 fresh, 2,580 spoiled). Figure 1 illustrates the indoor acquisition setup.

The in-distribution dataset was organized into four freshness categories: (i) Unpackaged Fresh Meat, (ii) Unpackaged Spoiled Meat, (iii) Packaged Fresh Meat, and (iv) Packaged Spoiled Meat. In addition, we curated a separate *No Result* set for OOD/rejection evaluation. Each packaged-meat image was labeled by its spoilage stage following AMSA visual assessment standards [9], [35], which describe changes in surface color, moisture, and texture.

To maintain labeling consistency, we applied a chronological rule based on exposure time: images captured within 10–36 hours at room temperature (25°C) were labeled as fresh, after which they were transitioned to the spoiled class. While we acknowledge that this duration is not a strict biochemical threshold, it serves as a practical visual freshness proxy. This protocol is motivated by general food-safety guidelines stating that detectable sensory changes in meat (color, odor, and texture) typically manifest within 24–48 hours under non-refrigerated conditions [36]–[39].

We note that this rule is not a biochemical threshold. We did not validate these labels using microbial measurements (e.g., total viable counts), and we did not run an expert sensory panel. AMSA is cited here as a visual reference for what changes to look for (color, surface moisture, texture), not as expert grading. Labeling was performed by student annotators using a short written guide with example images, and we did not measure inter-annotator agreement. Therefore, the packaged-meat labels should be interpreted as a visual freshness proxy under our protocol, and some label noise is expected near the transition period.

We also collected an additional set of irrelevant images (*No Result*, 800 images, ~13%) such as empty trays, cutting boards, and background-only frames (Fig. 2). This set is used to evaluate the OOD/rejection mechanism (Section III-B5) and to emulate realistic failure cases where a camera captures non-meat scenes. It is kept separate from the four in-distribution freshness classes and is not used as a supervised class for freshness classification.

##### 2) Data Cleaning and Deduplication

Before modeling, we audited the merged corpus of 11,163 images for duplicates and near-duplicates to mitigate the risk of split leakage and overly optimistic performance estimates [40], [41]. We computed perceptual hashes (pHash) to flag

**TABLE 1.** Per-class distribution after cleaning and stratified 70/15/15 splitting for the four in-distribution freshness classes. Unpackaged sources: [32]–[34]; Packaged: this study (Apr 2024).

| Category                | Train        | Val        | Test       | Total        | Share (%)    |
|-------------------------|--------------|------------|------------|--------------|--------------|
| Packaged Fresh Meat     | 656          | 138        | 146        | 940          | 15.0         |
| Packaged Spoiled Meat   | 805          | 170        | 180        | 1,155        | 18.5         |
| Unpackaged Fresh Meat   | 1,159        | 246        | 256        | 1,661        | 26.6         |
| Unpackaged Spoiled Meat | 1,187        | 252        | 261        | 1,700        | 27.2         |
| <b>Total</b>            | <b>3,807</b> | <b>806</b> | <b>843</b> | <b>5,456</b> | <b>100.0</b> |

near-identical frames primarily caused by small viewpoint or illumination shifts common in rapid-succession or burst-mode acquisition [42], [43]. If these near-duplicates were to persist across train/val/test splits or across folds in our nested cross-validation, the evaluation would be biased upward because the model would effectively be tested on scenes already encountered during training [40], [41], [44]. Consequently, pHash filtering was strictly applied *before* any data partitioning.

After removing 4,907 redundant frames (a 43.95% reduction), the final dataset contained 6,256 unique images. To avoid over-pruning, we did not discard entire spoilage sequences; instead, filtering was restricted to near-duplicate clusters (defined by a small pHash distance), ensuring at least one representative image per capture condition was retained. We removed near-duplicate frames while keeping at least one representative per capture condition to reduce leakage risk. While this step may discard minor micro-variations, it significantly improves the statistical independence of the evaluation samples. As part of our future work, we plan to explore group-wise splitting, where each near-duplicate cluster is kept within a single fold to maximize variation while preventing leakage [45].

After deduplication, we retained 6,256 unique images in total, consisting of 5,456 in-distribution (ID) images across the four freshness classes (Table 1) and an additional 800 *No Result* images reserved exclusively for out-of-distribution (OOD) evaluation (Section III-B5). During training, we applied standard data augmentations (random rotation, horizontal flip, and color jitter) to improve robustness.

## B. MACHINE LEARNING ARCHITECTURES AND EVALUATION PIPELINE

### 1) Model Architectures

We compared five backbones with different accuracy–efficiency trade-offs. ResNet-50 [46] is used as the CNN baseline. We include two transformer models, ViT-B/16 [21] and Swin-T [22]. We also test two lightweight CNNs for mobile deployment, MobileNetV3-Small [47] and EfficientNet-B0 [48]. All models were initialized with ImageNet-pretrained weights and trained with  $224 \times 224$  inputs.

Hyperparameters were selected with the same nested cross-validation protocol (Section III-B4) for all five backbones. MobileNetV3-Small and EfficientNet-B0 did not reuse the

ResNet-50 configuration. Each model was tuned independently in the inner loop using the same search space, and the best setting can differ by backbone.

To adapt each backbone to our four-class in-distribution task, we replaced the classifier with a small head: a fully connected layer (512 units, ReLU), dropout ( $p = 0.3$ ), and a softmax output. Models were trained with categorical cross-entropy, where label smoothing was treated as a tunable hyperparameter ( $\alpha \in \{0.0, 0.1\}$ ) selected in the inner loop of nested cross-validation (Table 2) and then applied in the final retraining using the best-performing configuration.

We fine-tuned the full networks rather than freezing the backbones. We first trained the classification head, then unfroze backbone layers in stages with a smaller learning rate. This helps stabilize training and reduces overfitting in visually similar classes [49]–[51].

### 2) U-Net Segmentation

We use a U-Net segmentation model with ResNet-18/34/50 encoders to extract the meat region before classification. All models were initialized using ImageNet pretraining to leverage low-level visual features. In this setup, the encoder extracts hierarchical texture features, while the decoder recovers spatial details through skip connections, ensuring precise mask generation.

To train the segmentation model, we adopted a weakly supervised labeling strategy by generating pseudo-ground truth masks using an adaptive GrabCut procedure in automatic rectangle mode (five iterations) [52]–[54]. The initialization box covered approximately 90% of the image and was perturbed by  $\pm 10\%$  to robustly enclose the foreground, guided by color and texture cues in HSV and Lab spaces. We applied morphological opening and closing to suppress noise and smooth boundaries. To sanity-check pseudo-label quality, we manually reviewed 589 masks (approximately 11% of the dataset) spanning both packaged and unpackaged images, confirming consistent separation of the meat region. Because these masks are automatically produced rather than expert annotations, IoU and Dice should be interpreted as agreement with the pseudo-labeling procedure rather than absolute segmentation accuracy. Our primary goal is background suppression (e.g., tray clutter and reflections) to help downstream classification focus on the meat interior region; a fully annotated segmentation benchmark is left for future work.

Hyperparameter tuning for each encoder was conducted independently through an exhaustive search of 108 candidate combinations, varying the batch size, learning rate (LR), weight decay (WD), loss formulation, and learning rate scheduler. Upon identifying the optimal configuration, the parameters were fixed for the final training stage to ensure reproducibility. All segmentation models were trained on  $384 \times 384$  pixel images using the Adam optimizer, with early stopping based on validation loss to mitigate overfitting.

Training was executed for up to 50 epochs using a hybrid objective function combining Binary Cross-Entropy and Dice

loss to effectively balance region-level accuracy with boundary sharpness. Performance was quantified via Intersection-over-Union (IoU), Dice coefficient, precision, recall, and pixel accuracy, with 95% confidence intervals (CI) derived from bootstrapping with 5,000 resamples. Finally, we zero out pixels outside the predicted mask to reduce background influence and encourage the classifier to rely on meat-region cues.

### 3) Balancing and Loss Function

To mitigate class imbalance among the four in-distribution freshness classes, we optionally used sampling-based rebalancing during training. When enabled, `WeightedRandomSampler` upweighted under-represented classes so that mini-batches contained a more balanced mix of the four classes.

Models were trained with categorical cross-entropy (with label smoothing when specified in the training configuration). The *No Result* images were not treated as a supervised class during training; they were held out and used only for OOD evaluation (Section III-B5).

### 4) Nested Cross-Validation

We used a nested cross-validation (CV) design with 5 outer folds and 3 inner folds. Hyperparameters were selected exclusively on the inner folds, while the outer folds were used only for evaluation, reducing information leakage and providing a more reliable estimate of generalization.

To keep tuning computationally feasible, we adopted a two-stage inner-loop procedure aligned with transfer-learning practice. In *Inner-1* (head warm-up), the backbone was frozen and we searched over the classification-head learning rate, weight decay, and label smoothing. In *Inner-2* (unfreeze refinement), we unfroze the backbone and refined the top candidates by searching over the backbone learning rate and MixUp strength while keeping the best head configuration fixed. All models were trained with AdamW using separate parameter groups for the head and backbone, together with a warmup+decay epoch scheduler. During the 1-epoch head warm-up, the backbone remained frozen and was then unfrozen for Inner-2 refinement. All backbones (ResNet-50, Swin-T, ViT-B/16, EfficientNet-B0, and MobileNetV3-Small) followed the same tuning protocol and search space; the selected hyperparameters were determined independently per model. The best inner-loop configuration was then used for retraining within the corresponding outer fold, and the final retraining reused the most consistently selected configuration across outer folds.

Table 2 summarizes the hyperparameter search space explored.

### 5) OOD-Aware Rejection Mechanism

We add a simple rejection option based on prediction confidence to avoid forced low-confidence decisions. Given an input image  $x$ , the classifier outputs softmax probabilities over the four freshness classes, and we use the maximum

**TABLE 2. Hyperparameter search space explored during inner-loop tuning under the nested 5×3 CV framework (classification).**

| Hyperparameter                   | Search Range / Setting                  |
|----------------------------------|---|
| Head learning rate (Inner-1)     | $1 \times 10^{-3}$ – $3 \times 10^{-3}$ |
| Backbone learning rate (Inner-2) | $1 \times 10^{-5}$ – $3 \times 10^{-4}$ |
| Weight decay (Inner-1)           | $1 \times 10^{-4}$ – $1 \times 10^{-1}$ |
| Label smoothing $\alpha$         | {0.0, 0.1}                              |
| MixUp $\alpha$                   | {0.0, 0.2}                              |
| Batch size                       | 32                                      |
| Optimizer                        | AdamW                                   |
| Scheduler                        | Warmup+decay                            |
| Head warm-up                     | 1 epoch                                 |
| Warmup                           | 2 epochs                                |
| Early stopping                   | fixed per loop                          |

softmax probability (MSP) as a confidence score [26]:  $s(x) = \max_k p(y = k | x)$ . If  $s(x) < \tau$ , the model abstains from assigning a 4-class label; in deployment, such samples can be mapped to a separate *No Result* output.

We additionally evaluate two widely used confidence scores: Energy and ODIN [30], [31]. Energy computes a score from the logits, while ODIN applies temperature scaling and a small input perturbation before computing confidence. In this work, we distinguish *semantic OOD* inputs from *hard in-distribution* cases. Semantic OOD refers to non-meat or irrelevant scenes that fall outside the four-class label space, and we operationalize this using the held-out *No Result* set. In contrast, artifacts such as glare, film reflections, occlusion, and borderline freshness are treated as hard in-distribution samples because they still contain meat and belong to one of the four target classes; the rejection mechanism is nevertheless useful for reducing overconfident predictions on these challenging inputs [55].

To avoid relying on a single threshold choice, we sweep  $\tau$  and report the resulting coverage–rejection trade-off (Figs. 6–7). For a single reference operating point across models, we report results at  $\tau = 0.5$  as a simple baseline (not an optimum), since it lies near the transition where rejection begins to increase while coverage remains close to its maximum. OOD evaluation uses the held-out *No Result* set as semantic OOD examples, and we compare MSP, Energy, and ODIN under the same protocol in Section IV-D.

### 6) Evaluation Metrics.

We report two types of results: standard 4-class freshness classification performance and OOD detection performance using the held-out *No Result* set.

For in-distribution classification, we compute categorical cross-entropy loss and macro-averaged precision, recall, and F1 over the four freshness classes. The loss is defined as

$$L = -\frac{1}{N} \sum_{i=1}^N \sum_{j=1}^C y_{i,j} \log(p_{i,j}), \quad (1)$$

where  $N$  is the number of images,  $C = 4$  is the number of in-distribution classes,  $y_{i,j}$  is the binary indicator of the true class, and  $p_{i,j}$  is the predicted probability.

**TABLE 3.** Global segmentation performance of U-Net backbones with 95% confidence intervals.

| Backbone  | IoU                      | Dice                     | Precision                | Recall                   | Pixel Acc.               |
|-----------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|
| ResNet-18 | 0.7211<br>[0.703, 0.739] | 0.8021<br>[0.786, 0.817] | 0.7780<br>[0.761, 0.794] | 0.8727<br>[0.857, 0.888] | 0.8569<br>[0.846, 0.868] |
| ResNet-34 | 0.7494 [0.731, 0.767]    | 0.8188 [0.802, 0.835]    | 0.8186 [0.802, 0.835]    | 0.8490<br>[0.832, 0.866] | 0.8777 [0.867, 0.889]    |
| ResNet-50 | 0.7320<br>[0.713, 0.750] | 0.8057<br>[0.789, 0.822] | 0.8103<br>[0.793, 0.826] | 0.8419<br>[0.825, 0.859] | 0.8709<br>[0.860, 0.882] |

For OOD detection, we treat the four-class test split as in-distribution inputs and the *No Result* set as OOD inputs. Using confidence scores from MSP, Energy, and ODIN (Section III-B5), we report AUROC, AUPR, and FPR@95TPR on this ID–OOD benchmark. We also run a threshold sensitivity sweep to show the coverage–rejection trade-off (Section IV-D, Figs. 6–7). For MSP/ODIN we sweep the confidence threshold  $\tau$ , while for Energy we sweep score thresholds on the energy scale (details in Section IV-D).

## IV. RESULTS

### A. SEGMENTATION RESULTS

Since the training masks were generated as pseudo-ground truth via GrabCut [54], the segmentation metrics reported here reflect the model’s agreement with these automated labels rather than absolute accuracy against manual expert annotations [52], [53]. We include these metrics primarily to benchmark the backbones’ consistency and to verify that the generated masks provide a reliable foundation for foreground extraction. All segmentation metrics are reported on the held-out ID test split (the same split used for downstream classification) to assess generalization on unseen images. These test-set metrics are used for evaluation only; model selection and early stopping were performed using the training/validation splits. Table 3 summarizes the global performance across three ResNet backbones. ResNet-34 achieved the highest mean IoU (0.7494), Dice (0.8188), precision (0.8186), and pixel accuracy (0.8777). While ResNet-18 exhibited the highest recall (0.8727), minimizing missed meat pixels at the expense of higher false positives, ResNet-50 trailed in overall overlap metrics. Per-class IoU analysis (Table 4) confirms this trend, with ResNet-34 leading in three out of four categories. Given this robust balance across overlap and precision, ResNet-34 was selected as the fixed encoder for the subsequent classification pipeline.

Integrating U-Net segmentation standardized the region of interest (ROI) and suppressed background variations, such as tray edges and plastic reflections, facilitating more stable feature learning [27]. Because the primary objective is background suppression rather than pixel-perfect boundaries, minor boundary imperfections in the pseudo-labels are expected to have limited impact on downstream classification, provided that the core color and texture cues of the meat interior region are preserved. We note that this is a design assumption of the pipeline; a dedicated ablation study (same

**TABLE 4.** Per-class IoU performance with 95% confidence intervals.

| Backbone  | Packaged Fresh           | Packaged Spoiled         | Unpackaged Fresh         | Unpackaged Spoiled       |
|-----------|--------------------------|--------------------------|--------------------------|--------------------------|
| ResNet-18 | 0.8726<br>[0.842, 0.900] | 0.7749<br>[0.741, 0.808] | 0.6693<br>[0.639, 0.699] | 0.6501<br>[0.612, 0.687] |
| ResNet-34 | 0.8888<br>[0.859, 0.915] | 0.8072<br>[0.771, 0.842] | 0.6898<br>[0.659, 0.722] | 0.6900<br>[0.653, 0.728] |
| ResNet-50 | 0.8900<br>[0.858, 0.918] | 0.7997<br>[0.763, 0.834] | 0.6738<br>[0.643, 0.704] | 0.6541<br>[0.615, 0.693] |

**TABLE 5.** Nested CV stability (outer-fold mean $\pm$ SD).

| Model             | Acc (%)          | F1 (%)           | Prec. (%)        | Rec. (%)         |
|-------------------|------------------|------------------|------------------|------------------|
| ResNet-50         | 98.47 $\pm$ 0.37 | 98.47 $\pm$ 0.37 | 98.51 $\pm$ 0.35 | 98.47 $\pm$ 0.37 |
| Swin-T            | 97.75 $\pm$ 0.64 | 97.75 $\pm$ 0.64 | 97.80 $\pm$ 0.62 | 97.75 $\pm$ 0.64 |
| ViT-B/16          | 94.16 $\pm$ 1.03 | 94.13 $\pm$ 1.05 | 94.67 $\pm$ 0.76 | 94.16 $\pm$ 1.03 |
| EfficientNet-B0   | 97.88 $\pm$ 0.72 | 97.88 $\pm$ 0.72 | 97.96 $\pm$ 0.67 | 97.88 $\pm$ 0.72 |
| MobileNetV3-Small | 96.64 $\pm$ 0.93 | 96.63 $\pm$ 0.93 | 96.91 $\pm$ 0.79 | 96.64 $\pm$ 0.93 |

backbone and evaluation protocol) comparing classification performance with and without segmentation is left for future work. To quantify the stability of our segmentation agreement estimates against the pseudo-labeling procedure, we performed image-level bootstrapping; for the selected ResNet-34 model, the Dice coefficient was 0.8188 with a 95% CI of [0.802, 0.835]. These standardized segmentation-based inputs provide a consistent baseline for the classifiers evaluated in Section IV-B.

### B. NESTED CROSS-VALIDATION PERFORMANCE

All classifiers were evaluated with a nested  $5 \times 3$  cross-validation protocol. The inner loop performed hyperparameter tuning and the outer loop measured generalization, thereby preventing information leakage and keeping estimates unbiased. Table 5 reports outer-fold mean $\pm$ SD for accuracy, macro-F1, precision, and recall; Table 6 summarizes the selected hyperparameters.

ResNet-50 obtained the highest outer-fold accuracy (98.47%  $\pm$  0.37), followed by Swin-T (97.75%  $\pm$  0.64) and ViT-B/16 (94.16%  $\pm$  1.03). Convolutional inductive bias and Swin’s shifted-window hierarchy likely aid stable training on this dataset, whereas ViTs typically require larger and more diverse data with stronger regularization to capture spatial structure [21], [56].

Per-fold outer-loop accuracies (Table 7) show tight clus-

**TABLE 6.** Nested cross-validation (5 outer  $\times$  3 inner) results and best hyperparameters per model.

| Model             | Acc. (Mean $\pm$ SD, %) | Head LR | Backbone LR | Weight Decay | Label Smooth | Mixup $\alpha$ | Best Epoch |
|-------------------|-------------------------|---------|-------------|--------------|--------------|----------------|------------|
| ResNet-50         | 98.47 $\pm$ 0.37        | 0.003   | 1e-4        | 0.0005       | 0.0          | 0.0            | 19-30      |
| Swin-T            | 97.75 $\pm$ 0.64        | 0.0015  | 7e-5        | 0.1          | 0.0          | 0.0            | 13-30      |
| ViT-B/16          | 94.16 $\pm$ 1.03        | 0.001   | 5e-5        | 0.05         | 0.1          | 0.0            | 13-30      |
| EfficientNet-B0   | 97.88 $\pm$ 0.72        | 0.001   | 0.0003      | 0.0001       | 0.1          | 0.0            | 17-30      |
| MobileNetV3-Small | 96.64 $\pm$ 0.93        | 0.001   | 0.0003      | 0.0005       | 0.0          | 0.2            | 17-30      |

**TABLE 7.** Outer-fold accuracies (%) under nested 5 $\times$ 3 CV.

| Outer Fold                      | ResNet-50 (Acc %)                  | Swin-T (Acc %)                     | ViT-B/16 (Acc %)                   | EfficientNet-B0 (Acc %)            | MobileNetV3-Small (Acc %)          |
|---------------------------------|------------------------------------|------------------------------------|------------------------------------|------------------------------------|------------------------------------|
| Fold 1                          | 98.12                              | 97.18                              | 93.05                              | 97.92                              | 95.09                              |
| Fold 2                          | 98.36                              | 97.61                              | 94.92                              | 97.64                              | 97.26                              |
| Fold 3                          | 98.89                              | 98.24                              | 95.21                              | 98.11                              | 96.69                              |
| Fold 4                          | 98.53                              | 97.82                              | 93.84                              | 96.88                              | 97.45                              |
| Fold 5                          | 98.45                              | 97.89                              | 93.76                              | 98.87                              | 96.69                              |
| <b>Mean <math>\pm</math> SD</b> | <b>98.47 <math>\pm</math> 0.37</b> | <b>97.75 <math>\pm</math> 0.64</b> | <b>94.16 <math>\pm</math> 1.03</b> | <b>97.88 <math>\pm</math> 0.72</b> | <b>96.64 <math>\pm</math> 0.93</b> |

**TABLE 8.** Retrain results on the held-out ID test set.

| Model             | Train | Val  | Test (ID) |
|-------------------|-------|------|-----------|
| ResNet-50         | 99.9  | 98.5 | 97.63     |
| Swin-T            | 99.7  | 98.6 | 97.51     |
| ViT-B/16          | 99.1  | 94.2 | 94.42     |
| EfficientNet-B0   | 99.9  | 97.2 | 98.10     |
| MobileNetV3-Small | 99.9  | 97.5 | 97.63     |

tering across folds. For ResNet-50 the five outer folds span 98.12–98.89% (range 0.77%, SD 0.37%), for Swin-T 97.18–98.24% (range 1.06%, SD 0.64%), and for ViT-B/16 93.05–95.21% (range 2.16%, SD 1.03%). This pattern matches the macro-averaged summary in Table 5 and indicates low dispersion without fold-specific drift. Best epochs also fell within a narrow band for each model (ResNet-50: 19–30, Swin-T: 13–30, ViT-B/16: 13–30), and the inner loop repeatedly selected similar learning-rate and regularization settings (Table 6), both of which are consistent with stable optimization across folds.

The hyperparameters in Table 6 are the configurations most frequently selected by the inner loop and were reused for the final retraining with OOD-enabled inference to keep test results independent of tuning.

### C. RETRAIN & TESTING PERFORMANCE

Figures 3(a)–(e) show the training and validation accuracy curves during the final retraining stage on the four in-distribution (ID) classes. All backbones converge stably, with ResNet-50 and Swin-T exhibiting smooth training dynamics and a small train–val gap. EfficientNet-B0 and MobileNetV3-Small also converge quickly, but show a slightly larger separation between training and validation. ViT-B/16 improves more gradually and exhibits the largest train–val gap, which is consistent with its lower validation accuracy (94.2%) and the lowest held-out ID test accuracy in Table 8.

Table 8 summarizes the final retraining and held-out ID test results ( $N = 843$ ). EfficientNet-B0 achieves the high-

est ID test accuracy (98.10%), followed by ResNet-50 and MobileNetV3-Small (both 97.63%), and Swin-T (97.51%). ViT-B/16 is lower (94.42%), aligning with its reduced validation accuracy (94.2%) during retraining. Overall, the validation ranking broadly matches the test trend, suggesting that the selected backbones generalize reliably to the held-out ID set.

Per-class results in Table 9 show consistent strengths and weaknesses across backbones. EfficientNet-B0 achieves the best macro-average F1 (0.982), while ResNet-50 and MobileNetV3-Small are close behind (both 0.975), with Swin-T at 0.973. ViT-B/16 is lower (0.944), with the largest drops appearing in the packaged classes (e.g., Packaged Fresh recall 0.911 and Unpackaged Spoiled recall 0.885), which are more sensitive to glare, film reflections, and subtle appearance changes. Across all models, unpackaged classes remain comparatively easier, with near-saturated recall for Unpackaged Fresh in several backbones.

Figure 4 shows confusion matrices on the held-out ID test set ( $N = 843$ ). Across backbones, misclassifications are concentrated within the same packaging condition (fresh vs. spoiled), indicating that the main difficulty lies in subtle freshness cues rather than packaging identification. For EfficientNet-B0, MobileNetV3-Small, and ResNet-50, the dominant error is Unpackaged Spoiled  $\rightarrow$  Unpackaged Fresh (11 images; 4.2% of the Unpackaged Spoiled class). Swin-T exhibits a slightly different pattern, with more Packaged Spoiled  $\rightarrow$  Packaged Fresh confusions (10 images; 5.6%) and a small amount of Unpackaged Fresh  $\rightarrow$  Unpackaged Spoiled (8 images; 3.1%). ViT-B/16 shows the largest confusion on the hardest boundaries, especially Unpackaged Spoiled  $\rightarrow$  Unpackaged Fresh (30 images; 11.5%) and Packaged Fresh  $\rightarrow$  Packaged Spoiled (13 images; 8.9%), consistent with its lower macro-F1 in Table 9. To make the remaining errors more concrete, Fig. 5 shows representative failure cases. These images are challenging because the visual cues are weak or distorted, such as strong glare on packaged trays,

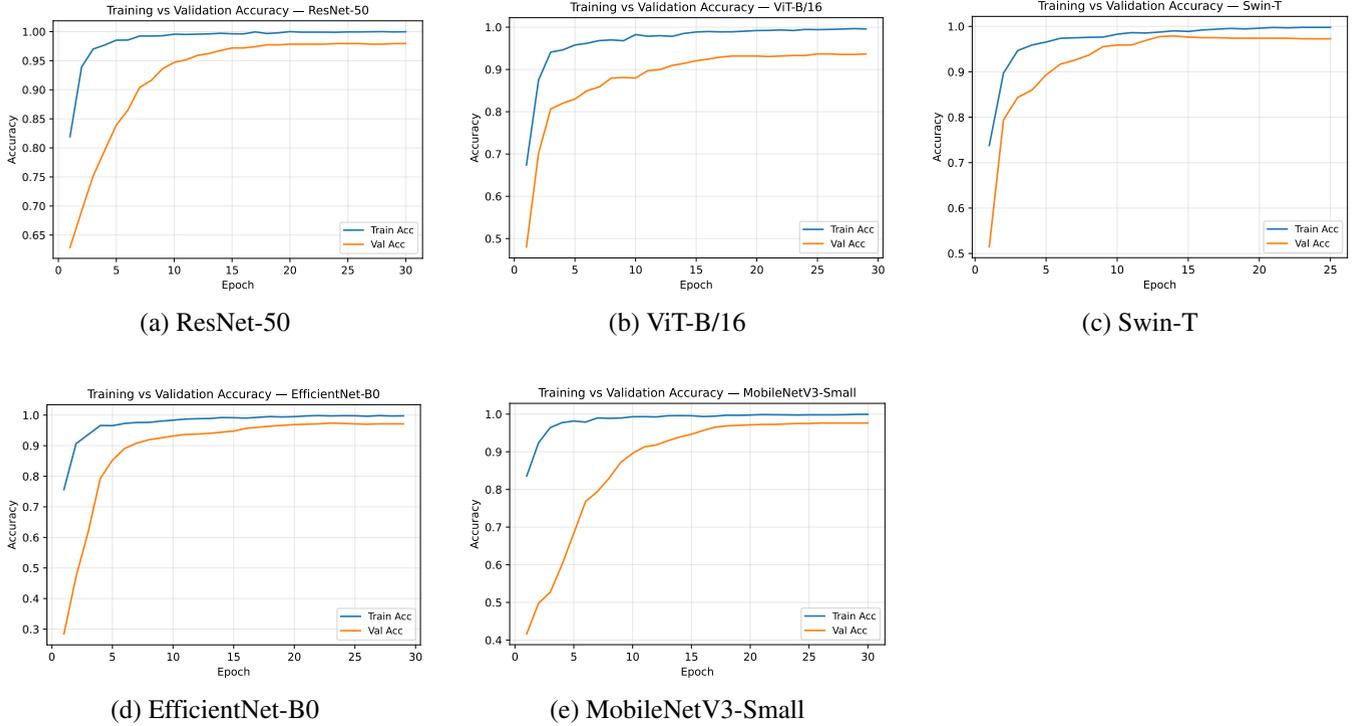


FIGURE 3. Training vs validation accuracy curves for five models during retraining with OOD-enabled inference.

TABLE 9. Per-class performance of five models on the held-out ID test set ( $N = 843$ ).

| Class              | ResNet-50 |       |       | Swin-T |       |       | ViT-B/16 |       |       | EfficientNet-B0 |       |       | MobileNetV3-Small |       |       |
|--------------------|-----------|-------|-------|--------|-------|-------|----------|-------|-------|-----------------|-------|-------|-------------------|-------|-------|
|                    | Prec.     | Rec.  | F1    | Prec.  | Rec.  | F1    | Prec.    | Rec.  | F1    | Prec.           | Rec.  | F1    | Prec.             | Rec.  | F1    |
| Packaged Fresh     | 0.986     | 0.952 | 0.969 | 0.935  | 0.993 | 0.963 | 0.971    | 0.911 | 0.940 | 0.973           | 0.993 | 0.983 | 0.954             | 0.986 | 0.970 |
| Packaged Spoiled   | 0.962     | 0.989 | 0.975 | 0.994  | 0.944 | 0.969 | 0.931    | 0.978 | 0.954 | 0.994           | 0.978 | 0.986 | 0.989             | 0.961 | 0.975 |
| Unpackaged Fresh   | 0.959     | 1.000 | 0.979 | 0.992  | 0.969 | 0.980 | 0.895    | 1.000 | 0.945 | 0.959           | 1.000 | 0.979 | 0.959             | 1.000 | 0.979 |
| Unpackaged Spoiled | 1.000     | 0.958 | 0.978 | 0.970  | 0.992 | 0.981 | 1.000    | 0.885 | 0.939 | 1.000           | 0.958 | 0.978 | 1.000             | 0.958 | 0.978 |
| Macro Avg          | 0.977     | 0.975 | 0.975 | 0.973  | 0.975 | 0.973 | 0.949    | 0.943 | 0.944 | 0.982           | 0.982 | 0.982 | 0.975             | 0.976 | 0.975 |

partial occlusion/clutter, or borderline freshness where appearance changes are subtle.

**D. OOD DETECTION PERFORMANCE**

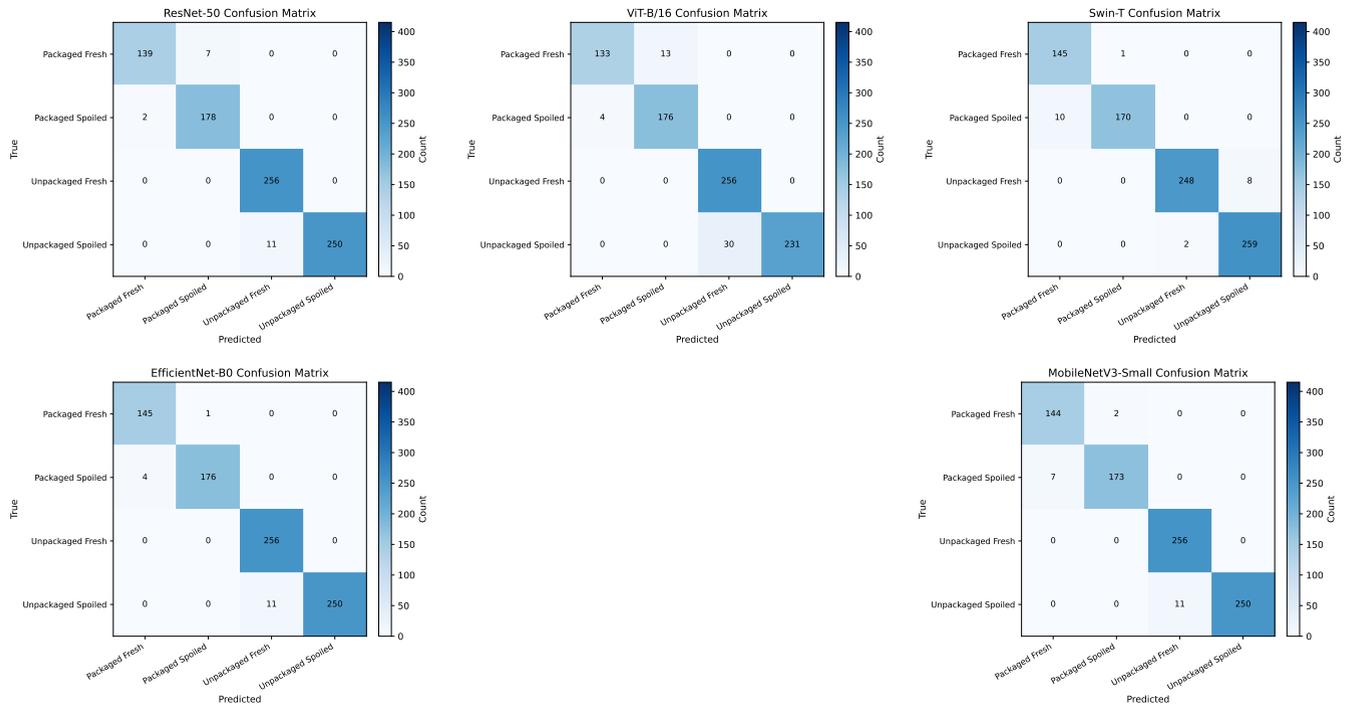
We include MSP as a simple baseline confidence score [26]. However, softmax confidence can be overconfident even on wrong or OOD inputs, so MSP may fail in practice [55]. We therefore also test ODIN, which uses temperature scaling and a small input perturbation to improve OOD separation [30]. We also test the energy score, which uses logits and is often stronger than softmax-based confidence for OOD detection [31]. In short, MSP is the lightweight baseline, while ODIN and Energy are stronger options with different runtime and tuning costs [30], [31].

While Table 8 summarizes ID classification accuracy on the four in-distribution classes, we also measure how well each scoring method separates in-distribution and OOD samples. Table 10 reports AUROC, AUPR, and FPR@95TPR for MSP, Energy, and ODIN. Across backbones, Energy and ODIN often improve AUROC/AUPR over MSP, although the

best method can vary by model. We also perform a sensitivity sweep over  $\tau \in \{0.2, 0.3, 0.4, 0.45, 0.5, 0.55, 0.6, 0.7, 0.8\}$  using the MSP confidence score on the in-distribution (4-class) set used for OOD evaluation. For each  $\tau$ , we report coverage (fraction of samples kept) and rejection rate (fraction of samples abstained). Figs. 6–7 show that the behavior is stable around  $\tau = 0.5$ , where only a small fraction of the lowest-confidence cases are rejected while coverage remains near-maximal.

**E. STATISTICAL VALIDATION**

Table 11 reports paired McNemar tests [57] on the held-out in-distribution (ID) test set ( $N = 843$ ), together with 95% confidence intervals for the paired accuracy difference  $\Delta acc$  (A–B). Each row includes the full paired contingency table  $(n_{11}, n_{10}, n_{01}, n_{00})$ , where  $n_{10}$  counts test images correctly classified by model A but misclassified by model B (and vice versa for  $n_{01}$ ). McNemar  $\chi^2$  statistics are computed using the continuity-corrected form  $(|n_{10} - n_{01}| - 0.5)^2 / (n_{10} + n_{01})$  (df = 1), reflecting paired predictions on the same test sam-



**FIGURE 4.** Confusion matrices on the held-out test set ( $N = 843$ ) for the four in-distribution classes (Packaged Fresh, Packaged Spoiled, Unpackaged Fresh, Unpackaged Spoiled). Errors mostly occur within the same packaging type (fresh vs. spoiled), which is the hardest visual boundary.



**FIGURE 5.** Representative failure cases: (i) glare on packaged samples, (ii) occlusion/clutter, and (iii) borderline freshness with subtle cues.

**TABLE 10.** OOD detection results across scoring methods on the OOD benchmark.

| Model             | Method | AUROC | AUPR  | FPR@95TPR |
|-------------------|--------|-------|-------|-----------|
| ResNet-50         | MSP    | 0.927 | 0.911 | 0.242     |
|                   | ENERGY | 0.939 | 0.941 | 0.254     |
|                   | ODIN   | 0.930 | 0.927 | 0.281     |
| Swin-T            | MSP    | 0.981 | 0.992 | 0.0250    |
|                   | Energy | 0.996 | 0.998 | 0.0025    |
|                   | ODIN   | 0.997 | 0.997 | 0.0025    |
| ViT-B/16          | MSP    | 0.993 | 0.947 | 0.0167    |
|                   | Energy | 0.999 | 0.994 | 0.0083    |
|                   | ODIN   | 0.998 | 0.988 | 0.0167    |
| EfficientNet-B0   | MSP    | 0.952 | 0.954 | 0.206     |
|                   | ENERGY | 0.982 | 0.977 | 0.083     |
|                   | ODIN   | 0.977 | 0.972 | 0.103     |
| MobileNetV3-Small | MSP    | 0.971 | 0.978 | 0.240     |
|                   | ENERGY | 0.978 | 0.982 | 0.131     |
|                   | ODIN   | 0.985 | 0.987 | 0.106     |

ples [58].

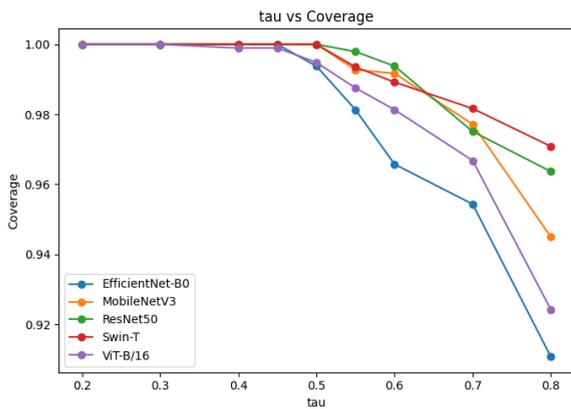
Across comparisons involving ViT-B/16, we observe con-

**TABLE 11.** McNemar paired contingency tables and confidence intervals for pairwise comparison of test accuracy on the held-out ID test set ( $N = 843$ ). McNemar  $\chi^2$  statistics use the continuity-corrected form  $(|n_{10} - n_{01}| - 0.5)^2 / (n_{10} + n_{01})$  ( $df = 1$ ).

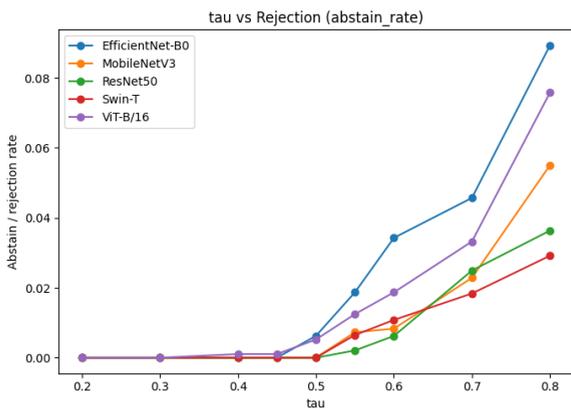
| Pair                                | $n_{11}$ | $n_{10}$ | $n_{01}$ | $n_{00}$ | $\chi^2$ | p-value               | $\Delta \text{acc (A-B)}$ | 95% CI             |
|-------------------------------------|----------|----------|----------|----------|----------|-----------------------|---------------------------|--------------------|
| ResNet50 (A) vs ViT (B)             | 788      | 35       | 8        | 12       | 16.331   | $5.32 \times 10^{-5}$ | +0.0320                   | [0.0169, 0.0471]   |
| ViT (A) vs Swin-T (B)               | 778      | 18       | 44       | 3        | 10.512   | 0.0012                | -0.0308                   | [-0.0490, -0.0127] |
| ViT (A) vs EfficientNet-B0 (B)      | 790      | 6        | 37       | 10       | 21.605   | $3.35 \times 10^{-6}$ | -0.0368                   | [-0.0518, -0.0217] |
| ViT (A) vs MobileNetV3-Small (B)    | 780      | 16       | 43       | 4        | 11.932   | 0.0005                | -0.0320                   | [-0.0498, -0.0143] |
| Swin-T (A) vs MobileNetV3-Small (B) | 815      | 7        | 8        | 13       | 0.0167   | 0.8973                | -0.0012                   | [-0.0102, 0.0078]  |

**TABLE 12.** Per-class F1-score with 95% bootstrap confidence intervals (percentile,  $B = 4000$ ) on the held-out ID test set ( $N = 843$ ).

| Class              | ResNet-50           | Swin-T              | ViT-B/16            | EfficientNet-B0     | MobileNetV3-Small   |
|--------------------|---------------------|---------------------|---------------------|---------------------|---------------------|
| Packaged Fresh     | 0.950 (0.921–0.974) | 0.944 (0.916–0.969) | 0.889 (0.847–0.925) | 0.946 (0.918–0.970) | 0.932 (0.899–0.960) |
| Packaged Spoiled   | 0.958 (0.933–0.978) | 0.951 (0.927–0.973) | 0.921 (0.891–0.947) | 0.955 (0.932–0.975) | 0.939 (0.911–0.964) |
| Unpackaged Fresh   | 0.982 (0.970–0.993) | 0.982 (0.970–0.993) | 0.936 (0.913–0.956) | 0.975 (0.960–0.988) | 0.961 (0.943–0.976) |
| Unpackaged Spoiled | 0.983 (0.970–0.993) | 0.983 (0.971–0.993) | 0.928 (0.903–0.951) | 0.974 (0.959–0.987) | 0.958 (0.939–0.975) |
| Macro Avg          | 0.968 (0.954–0.980) | 0.965 (0.951–0.978) | 0.919 (0.899–0.938) | 0.963 (0.948–0.976) | 0.947 (0.930–0.963) |



**FIGURE 6.** Threshold sweep on the in-distribution (4-class) set:  $\tau$  vs. coverage (fraction of samples kept).



**FIGURE 7.** Threshold sweep on the in-distribution (4-class) set:  $\tau$  vs. rejection rate (fraction of samples abstained).

sistent and statistically significant accuracy gaps in favor of the other backbones. Specifically, Swin-T outperforms ViT-B/16 by 3.08 percentage points ( $p = 0.0012$ ,  $\Delta \text{acc} = -0.0308$ , 95% CI [-0.0490, -0.0127]), EfficientNet-B0 by

3.68 points ( $p = 3.35 \times 10^{-6}$ ,  $\Delta \text{acc} = -0.0368$ , 95% CI [-0.0518, -0.0217]), and MobileNetV3-Small by 3.20 points ( $p = 0.0005$ ,  $\Delta \text{acc} = -0.0320$ , 95% CI [-0.0498, -0.0143]). These intervals do not cross zero, indicating that the observed differences are unlikely to be explained by sampling variability alone.

In contrast, Swin-T and MobileNetV3-Small exhibit nearly identical performance on the ID test set ( $\Delta \text{acc} = -0.0012$ ; 95% CI [-0.0102, 0.0078];  $p = 0.8973$ ), indicating no statistically detectable accuracy difference between these two models under paired testing. Finally, ResNet-50 exceeds ViT-B/16 by 3.20 percentage points ( $p = 5.32 \times 10^{-5}$ ,  $\Delta \text{acc} = +0.0320$ , 95% CI [0.0169, 0.0471]), consistent with the accuracy ranking in Table 8.

Complementing the accuracy-level hypothesis tests, Table 12 summarizes per-class F1-scores with 95% bootstrap confidence intervals (percentile,  $B = 4000$ ). Overall, the strongest models maintain high and stable F1 across all four ID classes. The packaged categories show larger uncertainty than the unpackaged classes, which is consistent with additional visual variability introduced by film reflections and glare. ViT-B/16 yields lower F1 and wider intervals on the harder packaged classes, whereas the CNN backbones and Swin-T retain high F1 with comparatively tighter confidence ranges.

### F. COMPARISON WITH PRIOR WORKS

Table 13 places our results alongside prior RGB-based meat freshness studies. Most earlier work reported high accuracy (93–99%) but focused on two or three classes and typically used unpackaged samples under controlled acquisition conditions. In this work, we curate a larger dataset (6,256 images) spanning both packaged and unpackaged meat and report strong held-out ID performance, with EfficientNet-B0 achieving 98.10% accuracy on the four-class ID test set.

OOD handling is evaluated separately from closed-set classification accuracy. Rather than treating “No Result”

**TABLE 13.** Quantitative comparison of prior RGB-based meat-freshness studies and this work. Results from this study are reported on the held-out in-distribution (ID) test set. OOD performance is evaluated separately (Section IV-D).

| Study                           | Classes       | Dataset Size | Reported Metric  |
|---------------------------------|---------------|--------------|--|
| Ulucan et al. (2019) [17]       | 2             | 1,896        | 99.6% (Accuracy)   |
| NUS Project (2023) [18]         | 3             | 2,268        | 93.1% (Accuracy)   |
| Abd Elfattah et al. (2025) [19] | 3             | 2,266        | 98.5% (Accuracy)   |
| <b>This Work (2026)</b>         | <b>4 (ID)</b> | <b>6,256</b> | <b>98.10% (EfficientNet-B0), 97.63% (ResNet-50 / MobileNetV3-Small), 97.51% (Swin-T)</b> |

as a conventional class, we implement a confidence-based OOD scoring and thresholding to enable abstention on uncertain inputs. Section IV-D reports AUROC/AUPR and threshold sweeps that characterize the trade-off between coverage and rejection, demonstrating that the system can defer low-confidence cases instead of forcing overconfident predictions.

Conventional sensing approaches such as spectroscopy, gas sensors, and electronic noses [11], [59], [60] can be effective but require dedicated hardware and are harder to deploy at scale. In contrast, using standard RGB images from commodity cameras or smartphones keeps costs low while maintaining competitive recognition performance on challenging visual conditions (e.g., glare and film reflections in packaged trays).

Overall, we extend prior work in three ways: (1) a broader formulation covering packaged and unpackaged meat, supported by a larger dataset; (2) an automated segmentation step that standardizes the region of interest before classification; and (3) an OOD-aware abstention mechanism to manage uncertainty. Together, these choices broaden the operating conditions and improve practical reliability.

### G. COMPUTATIONAL EFFICIENCY

Table 14 summarizes model size, training cost, and on-device inference latency for deployment-oriented comparison. Test accuracy and training time were measured in our controlled lab environment using an NVIDIA RTX 6000 Ada (48 GB) with 64 GB RAM and an AMD Ryzen Threadripper PRO 7945WX, whereas mobile latency was benchmarked by us on a single smartphone platform (Samsung Galaxy A55, SM-A556E; Android 16) using the TensorFlow Lite Benchmark APK. Because all latency results were obtained on the same device and runtime, the reported latencies are directly comparable across models under a consistent measurement protocol.

Among CNN backbones, EfficientNet-B0 and MobileNetV3-Small provide strong parameter efficiency (5.3 M and 2.5–2.9 M parameters) while maintaining high held-out ID accuracy (98.10% and 97.63%). EfficientNet-B0 offers the best accuracy–latency balance in our setting: it attains the highest ID test accuracy (98.10%) while remaining fast on-device (17.21 ms), substantially lower than ResNet-50 (76.92 ms) and Swin-T (108.17 ms). MobileNetV3-Small is the fastest model (6.36 ms) with competitive accuracy (97.63%), making it attractive under tight real-time constraints. ResNet-50 achieves similar accuracy (97.63%) but at a larger parameter budget (25.6 M) and higher mobile latency.

For Transformer-based models, Swin-T reaches competitive accuracy (97.51%) with a parameter count comparable to ResNet-50 (27.5–28 M), but incurs higher on-device latency (108.17 ms). ViT-B/16 shows the least favorable trade-off in our setting, combining the largest footprint (86 M) with the highest latency (462.75 ms) and the lowest ID accuracy (94.42%).

Training time per epoch varies only modestly across models (0.74–1.10 minutes/epoch), indicating that deployment-time inference latency is the primary differentiator for edge use. Overall, the results suggest that lightweight CNNs are the most practical option for real-time mobile inference under tight compute budgets, while higher-capacity backbones such as ResNet-50 or Swin-T may be preferred when model capacity is prioritized and higher latency is acceptable. Future work will standardize additional inference settings (e.g., quantization level and thread configuration) to quantify the impact of hardware-aware optimization on mobile performance.

### V. CONCLUSION

In this work, we proposed a meat-freshness assessment framework from RGB images that covers both packaged and unpackaged meat and incorporates an explicit *No Result* option to defer ambiguous inputs via OOD-aware abstention. The pipeline integrates U-Net segmentation with deep classifiers and applies nested five-by-three cross-validation for model selection and unbiased evaluation. The segmentation stage provided consistent region localization (mean IoU  $\approx 0.75$ , mean Dice  $\approx 0.82$ ), helping standardize inputs and stabilize training.

On the held-out in-distribution (ID) test set ( $N = 843$ ) with four meat classes, EfficientNet-B0 achieved the highest accuracy (98.10%), followed by ResNet-50 and MobileNetV3-Small (both 97.63%), and Swin-T (97.51%), while ViT-B/16 was lower (94.42%). Per-class and statistical analyses further confirmed that most remaining errors occur within the same packaging condition (fresh vs. spoiled), reflecting the subtle visual boundary, especially for packaged samples affected by glare and film reflections. In addition to closed-set accuracy, we evaluated OOD scoring methods (MSP, Energy, and ODIN) and a threshold sweep to characterize the coverage–rejection trade-off, showing that abstention can reduce overconfident decisions by filtering the lowest-confidence cases.

Finally, deployment-oriented benchmarking indicates that lightweight CNNs provide the best accuracy–latency balance on mobile hardware: EfficientNet-B0 and MobileNetV3-

**TABLE 14. Computational efficiency. Test accuracy and training time were measured on our lab GPU setup (RTX 6000 Ada), while mobile latency was benchmarked by us on a Samsung Galaxy A55 (SM-A556E), Android 16, using the TensorFlow Lite Benchmark APK (this work). Accuracy corresponds to the held-out ID test accuracy (Table 8).**

| Model             | Params (M) | Latency (ms) | Device / Runtime                         | Accuracy (%) | Train Time (min/epoch) |
|-------------------|------------|--------------|--|--------------|------------------------|
| ResNet-50         | 25.6       | 76.92        | Galaxy A55, Android 16, TFLite Benchmark | 97.63        | 1.10                   |
| EfficientNet-B0   | 5.3        | 17.21        | Galaxy A55, Android 16, TFLite Benchmark | 98.10        | 0.76                   |
| MobileNetV3-Small | 2.5-2.9    | 6.36         | Galaxy A55, Android 16, TFLite Benchmark | 97.63        | 0.74                   |
| Swin-T (Tiny)     | 27.5-28    | 108.17       | Galaxy A55, Android 16, TFLite Benchmark | 97.51        | 0.78                   |
| ViT-B/16          | 86         | 462.75       | Galaxy A55, Android 16, TFLite Benchmark | 94.42        | 0.84                   |

Small retain high ID accuracy with substantially lower on-device latency than larger backbones. Overall, combining segmentation, OOD-aware inference, and nested cross-validation yields a reliable and scalable baseline for practical meat freshness detection.

The experiments were carried out in a controlled setting to ensure repeatability and establish a clear benchmark. Future work may expand data collection to cover a wider range of lighting conditions, camera types, and spoilage stages, supported by expert annotation and laboratory reference measures. Microbiological and physicochemical validation could be incorporated using pH, psychrotroph counts, and standard microbial assays such as *E. coli* and *Salmonella* detection. In addition, short image sequences can be evaluated for temporal modeling of spoilage progression, explainability tools (Grad-CAM and LIME) may be integrated with expert feedback, and multimodal sensing that combines RGB with spectroscopy or gas-sensor inputs can be explored. These directions are expected to move the framework toward a more interpretable and practically deployable system for real-world food safety monitoring.

## ACKNOWLEDGMENT

We appreciate Colton Bailey and Josephine Nyame's help in the data collection process for this project. We also thank Yuchen Tian, Sam Chamoun, and Andrea Panebianco for their valuable discussions, feedback, and suggestions during the development and revision of this manuscript. Their insights contributed to improving the clarity and technical accuracy of the paper.

## REFERENCES

- [1] M. Alizadeh Sani, W. Zhang, A. Abedini, A. Khezerlou, N. Shariatifar, E. Assadpour, F. Zhang, and S. M. Jafari, "Intelligent packaging systems for the quality and safety monitoring of meat products: From lab scale to industrialization," *Food Control*, vol. 160, Article 110359, 2024.
- [2] M. Pateiro, P. E. S. Munkata, A. S. Sant'Ana, R. Domínguez, D. Rodríguez-Lázaro, and J. M. Lorenzo, "Application of essential oils as antimicrobial agents against spoilage and pathogenic microorganisms in meat products," *International Journal of Food Microbiology*, vol. 337, Article 108966, 2021.
- [3] M. M. Xu, M. Kaur, C. J. Pillidge, and P. J. Torley, "Microbial biopreservatives for controlling the spoilage of beef and lamb meat: Their application and effects on meat quality," *Critical Reviews in Food Science and Nutrition*, 62(17):4571-4592, 2022.
- [4] U.S. Department of Agriculture Economic Research Service. "Despite the fractional net decrease, U.S. per capita red meat and poultry disappearance is expected to remain near its historical high." 2024. Available: <https://www.ers.usda.gov/data-products/charts-of-note/chart-detail?chartId=103767> [Accessed: Feb. 10, 2024].
- [5] Y. Liu, Y. Zhang, F. Long, J. Bai, Y. Huang, and H. Gao, "CNN-assisted accurate smartphone testing of  $\mu$ PAD for pork sausage freshness," *Journal of Food Engineering*, vol. 363, p. 111772, 2024. doi: 10.1016/j.jfoodeng.2023.111772. Available: <https://www.sciencedirect.com/science/article/pii/S0260877423003709>
- [6] J. Tan, "Meat quality evaluation by computer vision," *Journal of Food Engineering*, vol. 61, no. 1, pp. 27-35, 2004. doi: 10.1016/S0260-8774(03)00185-7. Available: <https://www.sciencedirect.com/science/article/pii/S0260877403001857>
- [7] American Meat Science Association, *Research Guidelines for Cookery, Sensory Evaluation, and Instrumental Tenderness Measurements of Meat*, 1.02 ed., Champaign, IL: American Meat Science Association, 2015. Available: <https://meatscience.org/publications-resources/printed-publications/sensory-and-tenderness-evaluation-guidelines>
- [8] D. Mörlein, "Sensory evaluation of meat and meat products: fundamentals and applications," *IOP Conference Series: Earth and Environmental Science*, vol. 333, no. 1, p. 012007, Sep. 2019. doi: 10.1088/1755-1315/333/1/012007
- [9] D. A. King, M. C. Hunt, S. Barbut, J. R. Claus, D. P. Cornforth, P. Joseph, Y. H. B. Kim, G. Lindahl, R. A. Mancini, M. N. Nair, K. J. Merok, A. Milkowski, A. Mohan, F. Pohlman, R. Ramanathan, C. R. Raines, M. Seyfert, O. Sørheim, S. P. Suman, and M. Weber, "American Meat Science Association guidelines for meat color measurement," *Meat Muscle Biology*, vol. 6, no. 4, pp. 1-81, Feb. 2023. [Online]. Available: <https://www.iastatedigitalpress.com/mmb/article/id/12473/>, doi: 10.22175/mmb.12473.
- [10] K. Jo, S. Lee, S.-K.-C. Jeong, D.-H. Lee, H. Jeon, and S. Jung, "Hyperspectral imaging-based assessment of fresh meat quality: Progress and applications," *Microchemical Journal*, vol. 197, p. 109785, 2024. doi: 10.1016/j.microc.2023.109785. Available: <https://www.sciencedirect.com/science/article/pii/S0026265X23014042>
- [11] A. Y. Khaled, C. A. Parrish, and A. Adedeji, "Emerging nondestructive approaches for meat quality and safety evaluation—A review," *Comprehensive Reviews in Food Science and Food Safety*, vol. 20, no. 4, pp. 3438-3463, 2021. doi: 10.1111/1541-4337.12781
- [12] M. Rivai, F. Budiman, D. Purwanto, and J. Simamora, "Meat freshness identification system using gas sensor array and color sensor in conjunction with neural network pattern recognition," *Journal of Theoretical and Applied Information Technology*, vol. 96, no. 12, pp. 3861-3872, Jun. 2018.
- [13] X. Weng, X. Luan, C. Kong, Z. Chang, Y. Li, S. Zhang, S. Al-Majeed, and Y. Xiao, "A comprehensive method for assessing meat freshness using fusing electronic nose, computer vision, and artificial tactile technologies," *Journal of Sensors*, vol. 2020, Article 8838535, 2020. doi: 10.1155/2020/8838535
- [14] L. Zhou, C. Zhang, F. Liu, Z. Qiu, and Y. He, "Application of deep learning in food: A review," *Comprehensive Reviews in Food Science and Food Safety*, vol. 18, no. 6, pp. 1793-1811, 2019. doi: 10.1111/1541-4337.12492
- [15] N. Luo, D. Xu, B. Xing, X. Yang, and C. Sun, "Principles and applications of convolutional neural network for spectral analysis in food quality evaluation: A review," *Journal of Food Composition and Analysis*, vol. 128, p. 105996, 2024. doi: 10.1016/j.jfca.2024.105996.
- [16] G. A. Tahir and C. K. Loo, "A Comprehensive Survey of Image-Based Food Recognition and Volume Estimation Methods for Dietary Assessment," *Healthcare*, vol. 9, no. 12, p. 1676, 2021, doi: 10.3390/healthcare9121676.
- [17] O. Ulucan, D. Karakaya, and M. Turkan, "Meat quality assessment based on deep learning," *2019 Innovations in Intelligent Systems and Applications Conference (ASYU)*, pp. 1-5, Izmir, Turkey, 2019. doi: 10.1109/ASYU48272.2019.8946388
- [18] B. Sagiraju, N. Casanova, I. C. Chun, M. Lohia, and T. Yoshiyasu, "Meat

- freshness prediction,” *arXiv preprint*, arXiv:2305.00986, May 2023. Available: <https://arxiv.org/abs/2305.00986>
- [19] M. Abd Elfattah, A. A. Ewees, A. Darwish, and A. E. Hassanien, “Detection and classification of meat freshness using an optimized deep learning method,” *Food Chemistry*, vol. 489, p. 144783, 2025. doi: 10.1016/j.foodchem.2025.144783
- [20] E. Ofori-Mensah, “Leveraging Deep Learning Models for Beef Quality Assessment: A Comparative Study of CNN and ResNet Architecture,” M.S. thesis, School of Computing, Univ. of Eastern Finland, Joensuu, Finland, 2024. [Online]. Available: <https://urn.fi/URN:NBN:fi:uef-20241777>
- [21] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby, “An image is worth 16x16 words: Transformers for image recognition at scale,” *arXiv preprint arXiv:2010.11929*, 2021. [Online]. Available: <https://arxiv.org/abs/2010.11929>
- [22] Z. Liu, Y. Lin, Y. Cao, H. Hu, Y. Wei, Z. Zhang, S. Lin, and B. Guo, “Swin Transformer: Hierarchical Vision Transformer Using Shifted Windows,” in *Proc. IEEE/CVF Int. Conf. on Computer Vision (ICCV)*, 2021, pp. 9992–10022. [Online]. Available: <https://arxiv.org/abs/2103.14030>
- [23] S. Khan, M. Naseer, M. Hayat, S. W. Zamir, F. S. Khan, and M. Shah, “Transformers in Vision: A Survey,” *ACM Computing Surveys*, vol. 54, no. 10, pp. 1–41, 2022. doi: 10.1145/3505244
- [24] K. A. Nfor, T. P. T. Armand, P. I. Kenesbaeva, M.-I. Joo, and H.-C. Kim, “An Explainable CNN and Vision Transformer-Based Approach for Real-Time Food Recognition,” *Nutrients*, vol. 17, no. 2, p. 362, 2025. doi: 10.3390/nu17020362
- [25] Y. Gal and Z. Ghahramani, “Dropout as a Bayesian approximation: Representing model uncertainty in deep learning,” in *Proc. 33rd Int. Conf. Mach. Learn. (ICML)*, New York, NY, USA, 2016, pp. 1050–1059. [Online]. Available: <http://proceedings.mlr.press/v48/gal16.html>
- [26] D. Hendrycks and K. Gimpel, “A baseline for detecting misclassified and out-of-distribution examples in neural networks,” in *Proc. Int. Conf. Learn. Represent. (ICLR)*, Toulon, France, 2017. [Online]. Available: <https://openreview.net/forum?id=Hkg4T19xl>
- [27] O. Ronneberger, P. Fischer, and T. Brox, “U-Net: Convolutional Networks for Biomedical Image Segmentation,” in *Proc. Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, 2015, pp. 234–241.
- [28] A. Karim, T. Rahman, and M. Chowdhury, “Real-time grape leaf disease classification on Jetson Nano using MobileNetV3 and Grad-CAM,” *Computers and Electronics in Agriculture*, vol. 212, p. 108175, 2024. doi: 10.1016/j.compag.2024.108175.
- [29] R. A. H. Majdudin, “Pineapple Ripeness Detection Using YOLO v8 Algorithm,” *Publication of the International Journal of Academic Research (PIJAR)*, vol. 1, no. 1, pp. 50–54, 2024.
- [30] S. Liang, Y. Li, and R. Srikant, “Enhancing the Reliability of Out-of-Distribution Image Detection in Neural Networks,” in *Proc. International Conference on Learning Representations (ICLR)*, 2018.
- [31] W. Liu, X. Wang, J. Owens, and Y. Li, “Energy-based Out-of-Distribution Detection,” in *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- [32] Techtonics, *Meat Quality Assessment Dataset* — *kaggle.com*, 2023. [Online]. Available: <https://universe.roboflow.com/techtonics/meat-quality> [Accessed: 28-Mar-2024]
- [33] V. Shanawad, *Meat Freshness Image Dataset* — *kaggle.com*, 2022. [Online]. Available: <https://www.kaggle.com/datasets/vinayakshanawad/meat-freshness-image-dataset> [Accessed: 28-Mar-2024]
- [34] T. M. Dharma, R. Adriman, and K. Saddami, *Beef Quality Image Dataset* — *Mendeley Data*, 2024. [Online]. Available: <https://data.mendeley.com/datasets/7b67ynzr6k/1> [Accessed: 28-Mar-2024]
- [35] American Meat Science Association (AMSA), *Meat Judging Resources*, 2024, <https://meatscience.org/students/meat-judging-program/meat-judging-resources>, Accessed: 2024-03-04.
- [36] U.S. Department of Agriculture, Food Safety and Inspection Service, “Are you storing food safely?,” 2020. [Online]. Available: <https://www.fsis.usda.gov/food-safety/safe-food-handling-and-preparation/food-safety-basics/are-you-storing-food-safely> [Accessed: Sep. 28, 2025].
- [37] Healthline, “How long can raw meat sit out?,” 2023. [Online]. Available: <https://www.healthline.com/nutrition/how-long-can-raw-meat-sit-out> [Accessed: Sep. 28, 2025].
- [38] J. M. Jay, M. J. Loessner, and D. A. Golden, *Modern Food Microbiology*, 7th ed. New York: Springer, 2005. ISBN: 978-0387234137.
- [39] A. Orkus, M. Nowak, and I. Wojtasik-Kalinowska, “Impact of refrigerated storage on microbial growth, color stability, and pH of turkey thigh muscles,” *Microorganisms*, vol. 12, no. 6, p. 1114, 2024. doi: 10.3390/microorganisms12061114.
- [40] P. Ramos, R. Ramos, and N. Garcia, “Data Leakage in Visual Datasets,” in *Proc. IEEE/CVF Int. Conf. on Computer Vision (ICCV) Workshops*, 2025, pp. 6309–6319.
- [41] B. Barz and J. Denzler, “Do We Train on Test Data? Purging CIFAR of Near-Duplicates,” *Journal of Imaging*, vol. 6, no. 6, Art. no. 41, 2020, doi: 10.3390/jimaging6060041.
- [42] N. Krawetz, *Perceptual Hashing: The Future of Image Search*, Hacker Factor Solutions Whitepaper, 2010.
- [43] J. Buchner, “imagehash: A Python Perceptual Image Hashing Module,” GitHub repository, 2024. [Online]. Available: <https://github.com/JohannesBuchner/imagehash>
- [44] J. Wainer and G. C. Cawley, “Nested cross-validation when selecting classifiers is overzealous for most practical applications,” *Expert Systems with Applications*, vol. 182, Art. no. 115222, 2021, doi: 10.1016/j.eswa.2021.115222.
- [45] R. Joeres, D. B. Blumenthal, and O. V. Kalinina, “Data splitting to avoid information leakage with DataSAIL,” *Nature Communications*, vol. 16, Art. no. 3337, 2025, doi: 10.1038/s41467-025-58606-8.
- [46] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 770–778, doi: 10.1109/CVPR.2016.90.
- [47] A. Howard, M. Sandler, G. Chu, L.-C. Chen, B. Chen, M. Tan, W. Wang, Y. Zhu, R. Pang, V. Vasudevan, Q. V. Le, and H. Adam, “Searching for MobileNetV3,” in *Proc. IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019, pp. 1314–1324, doi: 10.1109/ICCV.2019.00140. [Online]. Available: [https://openaccess.thecvf.com/content\\_ICCV\\_2019/papers/Howard\\_Searching\\_for\\_MobileNetV3\\_ICCV\\_2019\\_paper.pdf](https://openaccess.thecvf.com/content_ICCV_2019/papers/Howard_Searching_for_MobileNetV3_ICCV_2019_paper.pdf)
- [48] M. Tan and Q. V. Le, “EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks,” in *Proc. 36th Int. Conf. Machine Learning (ICML)*, PMLR, vol. 97, pp. 6105–6114, 2019. [Online]. Available: <https://proceedings.mlr.press/v97/tan19a.html>
- [49] J. Howard and S. Ruder, “Universal Language Model Fine-tuning for Text Classification,” in *Proc. 56th Annual Meeting of the Association for Computational Linguistics (ACL)*, 2018, pp. 328–339, doi: 10.18653/v1/P18-1031. [Online]. Available: <https://aclanthology.org/P18-1031.pdf>
- [50] S. Kornblith, J. Shlens, and Q. V. Le, “Do Better ImageNet Models Transfer Better?” *Proc. CVPR*, 2019. Available: CVPR 2019
- [51] M. Mosbach, M. Andriushchenko, and D. Klakow, “On the Stability of Fine-tuning BERT: Misconceptions, Explanations, and Strong Baselines,” in *International Conference on Learning Representations (ICLR)*, 2021. [Online]. Available: <https://openreview.net/pdf?id=nzplWnVAYah>
- [52] G. Papandreou, L.-C. Chen, K. P. Murphy, and A. L. Yuille, “Weakly- and Semi-Supervised Learning of a Deep Convolutional Network for Semantic Image Segmentation,” in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, 2015, pp. 1742–1750.
- [53] A. Khoreva, R. Benenson, J. Hosang, M. Hein, and B. Schiele, “Simple Does It: Weakly Supervised Instance and Semantic Segmentation,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2017, pp. 876–885.
- [54] C. Rother, V. Kolmogorov, and A. Blake, “GrabCut: Interactive Foreground Extraction using Iterated Graph Cuts,” in *ACM Transactions on Graphics (SIGGRAPH)*, vol. 23, no. 3, 2004, pp. 309–314.
- [55] C. Guo, G. Pleiss, Y. Sun, and K. Q. Weinberger, “On Calibration of Modern Neural Networks,” in *Proc. Int. Conf. Mach. Learn. (ICML)*, 2017, pp. 1321–1330.
- [56] H. Touvron, M. Cord, M. Douze, F. Massa, A. Sablayrolles, and H. Jégou, “Training Data-Efficient Image Transformers and Distillation through Attention,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021, pp. 10347–10357, doi: 10.1109/ICCV48922.2021.01019.
- [57] N. Japkowicz and M. Shah, *Evaluating Learning Algorithms: A Classification Perspective*. Cambridge: Cambridge University Press, 2011. ISBN: 9781139061271.
- [58] F. Yates, “Contingency table involving small numbers and the  $\chi^2$  test,” *Supplement to the Journal of the Royal Statistical Society*, vol. 1, no. 2, pp. 217–235, 1934.
- [59] X. Wu, X. Liang, Y. Wang, B. Wu, and J. Sun, “Non-Destructive Techniques for the Analysis and Evaluation of Meat Quality and Safety: A Review,” *Foods*, vol. 11, no. 22, p. 3713, 2022. doi: 10.3390/foods11223713
- [60] E. Borch, M. L. Kant-Muermans, and Y. Blixt, “Bacterial Spoilage of Meat and Cured Meat Products,” *International Journal of Food Microbiology*, vol. 33, no. 1, pp. 103–120, 1996. doi: 10.1016/0168-1605(96)01135-X



**HUTAMA ARIF BRAMANTYO** is currently pursuing the Ph.D. degree in Electrical and Computer Engineering at Auburn University, AL, USA, as a Fulbright Doctoral Degree Grantee. He is also a Lecturer in the Department of Telecommunication Engineering at Politeknik Negeri Semarang, Indonesia. He received the B.Eng. degree in Electrical Engineering from Diponegoro University, Indonesia, in 2014, and the M.Eng. degree in Telecommunication Management from the University of Indonesia in 2019. Prior to academia, he worked in the telecommunications industry for over six years at Lintasarta, where he held various engineering and planning roles. His research interests include wireless communications, Internet of Things (IoT), antenna design, and the application of machine learning in agriculture and food systems. Mr. Bramantyo has published in IEEE conferences and national journals and has contributed to several community service and technology transfer initiatives related to digital agriculture and ICT for rural development.



**MUKARRAM ALI FARIDI** (Student Member, IEEE) received the B.Eng. degree in Computer Engineering from Auburn University, Auburn, AL, USA, in 2022 and the M.S. degree in Electrical Engineering from the same institution in 2025, where his graduate studies centered on VLSI testing, chip architecture and fabrication, and machine learning. He is currently an FPGA Engineer with Adtran, Richardson, TX, USA. Prior to joining industry, he was a Graduate Research Assistant at Auburn University, leading an NSF- and OCP-supported investigation into silent data corruption in integrated circuits, and an Embedded Software Engineer Intern at Club Car, Sarasota, FL, USA, where he developed an ESP32-based vehicle simulation automation platform. His research interests include VLSI testing, embedded systems, and real-world applications of machine learning, particularly in agriculture and food systems. Mr. Faridi received the Auburn University Woltosz Fellowship and a full-tuition scholarship for his graduate studies. As an undergraduate, he graduated *magna cum laude*, earned University Honors Scholar distinction, and was recognized among the top 20 undergraduates in the College of Electrical and Computer Engineering. He has published his VLSI testing research in IEEE conference proceedings.



**RUI CHEN** is an Assistant Professor in the Department of Agricultural and Environmental Sciences at Tuskegee University, AL, USA. She serves as an agricultural economist in the farmer technology adoption decisions thrust of the Artificial Intelligence for Future Agricultural Resilience, Management, and Sustainability (AIFARMS) Institute. She received the M.A. degree in Applied Economics from Ocean University of China in 2013 and earned the Ph.D. degree in Applied Economics, with a minor in Statistics, from Auburn University in 2017. From 2017 to 2020, she was a Postdoctoral Researcher and Instructor at Auburn University. Her current research interests include experimental and food economics, agricultural technology adoption, and the application of machine learning in agricultural and food systems. Dr. Chen has received several honors, including the New Investigator Award from USDA-NIFA (2023), the Outstanding Young Professional Award from the Committee on the Opportunities and Status of Blacks in Agricultural Economics (COSBAE) and the Committee on Women in Agricultural Economics (CWAE) of the Agricultural and Applied Economics Association (AAEA) (2022), and the EXCEL Faculty Fellowship from the 1890 Universities Foundation (2023). She secured millions of dollars in funding from the USDA to support her research in agricultural economics, machine learning, and public policy.



**CLARISSA HARRIS** received the B.S. and M.S. degrees in Biology and Biotechnology from Fort Valley State University, Fort Valley, GA, USA, in 2010 and 2013, respectively, and the Ph.D. degree in Animal Science from the University of Florida, Gainesville, FL, USA, in 2020. Her doctoral research focused on the antimicrobial effects of vinegar-based ingredients and lauric arginate against *Salmonella* on poultry products. She is currently a Research Extension Assistant Professor at Tuskegee University, Tuskegee, AL, USA. Her research interests include food safety, microbiological decontamination of meat, poultry processing, and antimicrobial intervention technologies. Dr. Harris has authored publications in journals such as *The Journal of Applied Poultry Research* and *Direct Research Journal of Agriculture and Food Science*. She has presented her work at numerous national conferences including the Poultry Science Association and MANRRS. She is a recipient of multiple awards, including the Booker T. Washington Leadership Institute Grant and the Delores A. Auzenne Dissertation Award. She was also selected as a 2025 Emerging Leader through the meat Institute. She holds certifications in HACCP and Preventive Controls for Animal Food, and is actively involved in outreach and mentoring through organizations such as MANRRS and LSAMP.



**YIN SUN** (Senior Member, IEEE) is the Bryghte D. and Patricia M. Godbold Endowed Associate Professor in the Department of Electrical and Computer Engineering at Auburn University, Alabama. He received his B.Eng. and Ph.D. degrees in Electronic Engineering from Tsinghua University in 2006 and 2011, respectively. From 2011 to 2017, he was a Postdoctoral Scholar and Research Associate at The Ohio State University. He joined Auburn University in 2017 as an Assistant Professor and was promoted to Associate Professor in 2023. His research interests include Semantic Communications, Wireless Networks, Information Theory, and the Application of Artificial Intelligence in Agriculture. He has served on the organizing committees of various international conferences, including as Technical Program Committee Chair for ACM MobiHoc 2025 and General Chair for IEEE/IFIP WiOpt 2026. He also delivered a keynote speech at IEEE/IFIP WiOpt 2025. Dr. Sun's publications have received multiple recognitions, including the Best Student Paper Award at IEEE/IFIP WiOpt 2013, the Best Paper Award at IEEE/IFIP WiOpt 2019, runner-up for the Best Paper Award at ACM MobiHoc 2020, and the Best Paper Award from the *Journal of Communications and Networks (JCN)* in 2021. He received the Auburn Author Award in 2020, the National Science Foundation (NSF) CAREER Award in 2023, the Bryghte D. and Patricia M. Godbold Endowed Professorship in 2023, the Ginn Faculty Achievement Fellowship in 2023, and the College of Engineering's Research Award for Excellence (Senior Faculty) in 2024.

...