

Segmentation of Fruit Surface Defects using Semi Supervised Learning

Saarukesan Premkumar*, Siyamalan Manivannan

Department of Computer Science, University of Jaffna, Sri Lanka

*saaru27kesan@gmail.com

Automated detection and segmentation of fruit-surface defects are essential for improving product quality, reducing dependence on manual inspection, and enabling scalable agricultural automation. Fruit defects vary widely in shape, texture, size, and severity, limiting the generalization of deep learning models. This study evaluates semantic segmentation of orange surface defects using a U-Net encoder–decoder with a ResNet34 backbone, trained on the Orange Navel dataset comprising 1,448 images annotated for three defect categories: sunburn, ulcer, and wind scarring. The model achieves strong quantitative performance: mean Pixel Accuracy (mPA) of 90.57%, mean Intersection over Union (mIoU) of 84.68%, and Dice coefficient of 91.37%. Class-wise IoU highlights near-perfect background segmentation (99.13%) and robust performance on sunburn (85.67%), while ulcer (83.61%) and wind scarring (70.29%) show reduced fidelity-likely due to subtler texture cues, lower contrast, and more irregular lesion boundaries. These results position U-Net with a ResNet34 encoder as a reliable baseline for fruit-defect segmentation. To strengthen generalization under limited annotations, a semi-supervised learning strategy based on the FixMatch paradigm is employed. In this approach, pseudo labels are generated for unlabeled samples, and only high-confidence predictions are used as supervisory signals. Training is guided by enforcing consistency between the model’s predictions and these pseudo labels, allowing effective use of both labeled and unlabeled data and improving segmentation performance across all defect categories.

Keywords: Fruit Defect Segmentation, Image Segmentation, U-Net, ResNet34, Semi Supervised Learning