Computer Vision Integration

Seyed Alireza Vaezi
sv22900@uga.edu

Hunter David Halloran
hdh20267@uga.edu

Jiajia Xie
jx25706@uga.edu

UNIVERSITY OF GEORGIA
College of Engineering
Manufacturing Living Labs

Overview & Where We Left Off

Motivation and Roadmap

Maximize or Control Yield

- Determine how machine parameters and pecan conditions influence crack dynamics and final yield.

Computer Vision helps extract detailed metrics to help better understand crack dynamics

Overview & Where We Left Off

Vision System Hardware Modification & Integration

- Installed and integrated **GigE network cameras**
- Dedicated high-intensity lighting

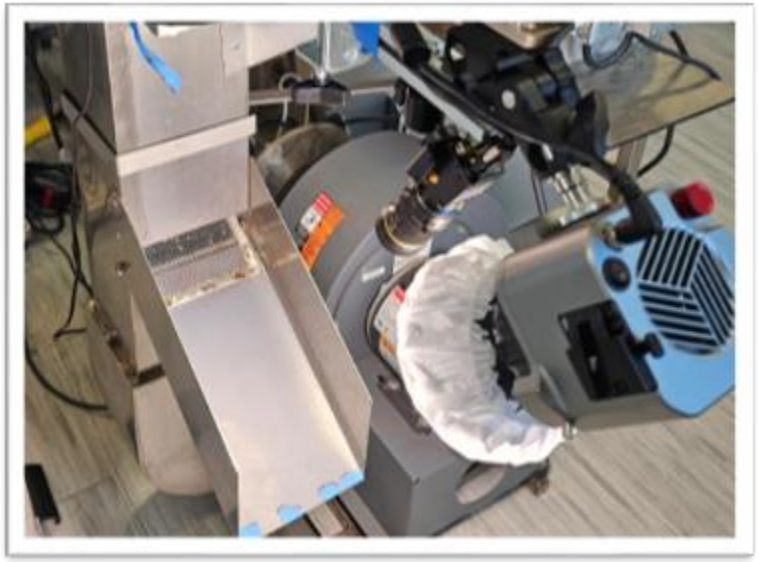

Overview & Where We Left Off

Vision System Hardware Modification & Integration

- Installed and integrated **GigE network cameras**
- Dedicated high-intensity lighting

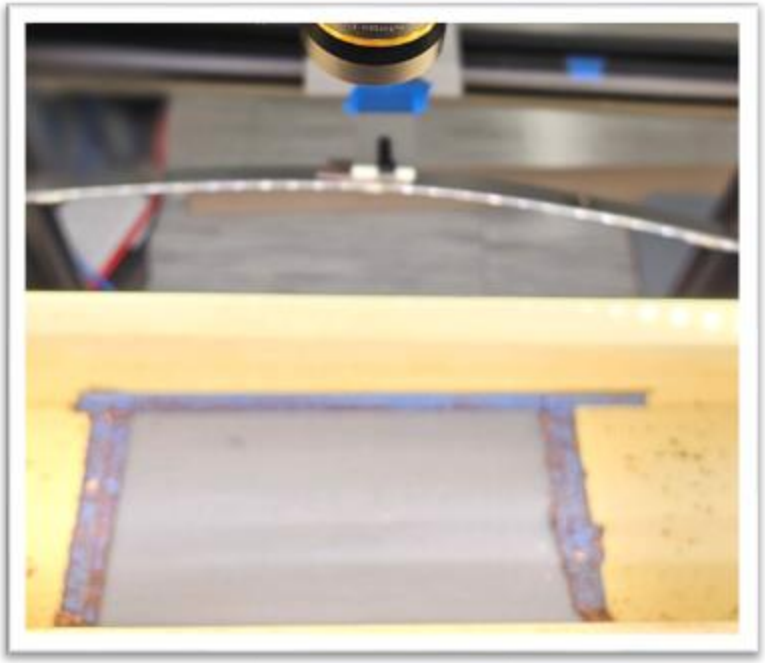

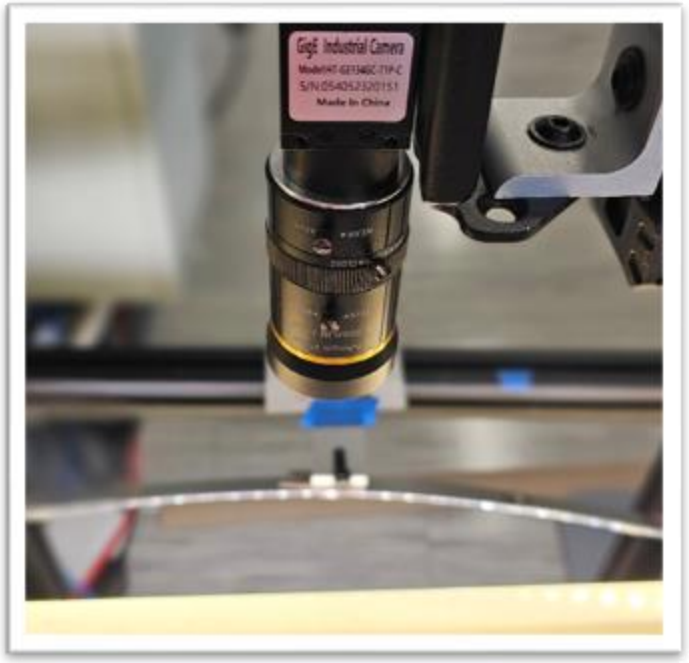

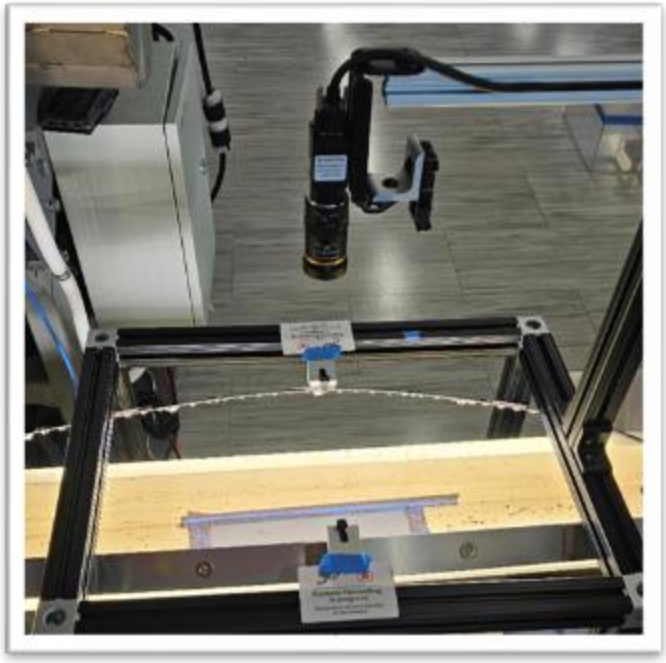

Blower separator exit

Vibratory conveyor post cracking

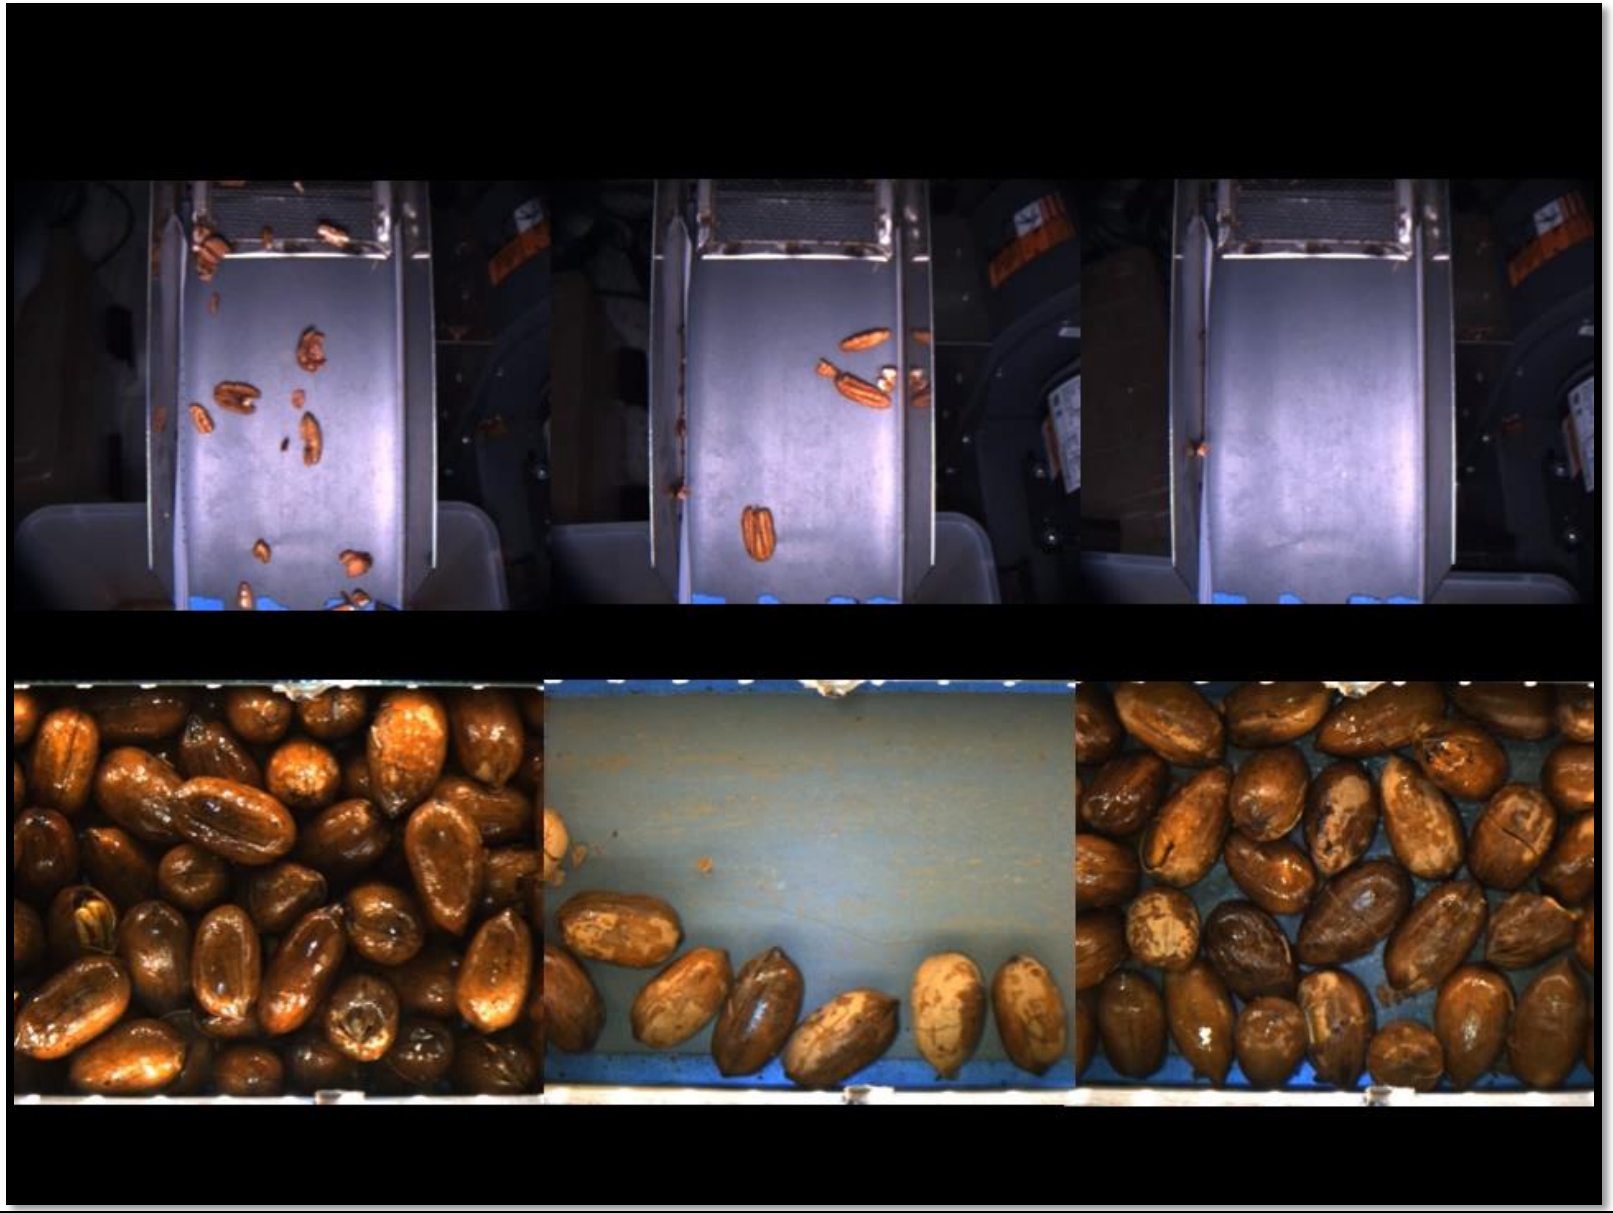

New Developments

UNIVERSITY OF GEORGIA

Half-Yield ratio Measurement using CV

The Model:
Trained on 2,400 labeled images with over 1,100 augmented samples.

Validation & Performance:
Strong class-level precision across four key classes (**Intact Half, Kernel Piece, Shell Piece, and Occluded**.)

Yield Calculation:
Instantly calculates the **Half-to-Full Yield Ratio**, providing a live metric for process optimization.

Crack Classification

Item	Value
Model / Framework	YOLO (Ultralytics 8.3.141)
Classes	Damaged, circumferential, cracked, longitudinal, open, uncracked
Data Split	1,527 Train (70%) / 435 Valid (20%) / 217 Test (10%)
Training Time	50 epochs (~2.3 hours)
Hardware	NVIDIA RTX 5000 Ada Generation
Model Size / Params	~87.7 MB / ~43.6M parameters
Speed	10.1 ms inference / image

Class	Precision (P)	Recall (R)	mAP50	mAP50-95
all (Overall)	**0.698**	**0.711**	**0.718**	**0.429**
Damaged	0.623	0.691	0.702	0.437
circumferential	0.425	0.300	0.266	0.096
cracked	0.915	0.945	0.971	0.719
longitudinal	0.618	0.610	0.602	0.235
open	0.868	0.845	0.899	0.439
uncracked	0.741	0.876	0.869	0.650

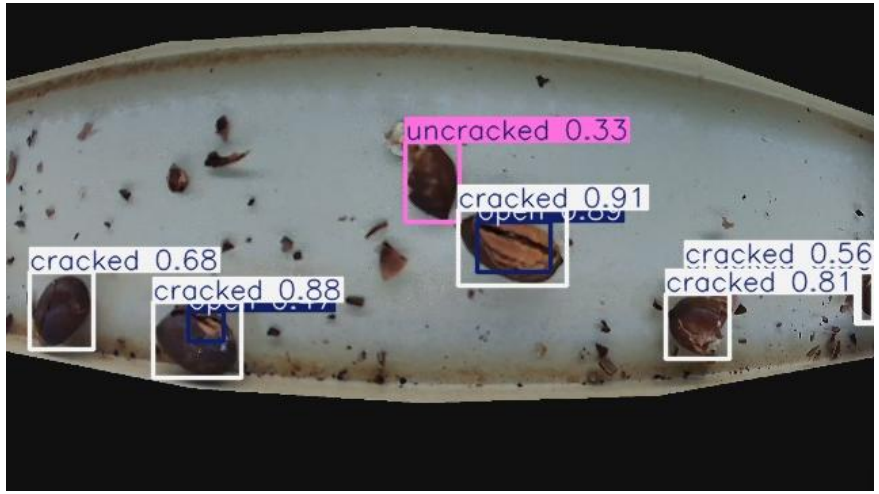

Crack Classification

Item	Value
Classes	Damaged, cracked, hairline, open, uncracked
Model	YOLOv8m fine-tuned on the Pecan v10 dataset
Parameters	25,842,655 (~25.8M)
Training GPU	NVIDIA GeForce RTX 5090, ~32 GB VRAM
Total training time	2.443 hours

Class	Precision	Recall	mAP50	mAP50-95
all (Overall)	**0.772**	**0.762**	**0.784**	**0.461**
Damaged	0.682	0.673	0.731	0.420
cracked	0.895	0.971	0.971	0.689
hairline	0.674	0.432	0.486	0.177
open	0.871	0.895	0.916	0.427
uncracked	0.737	0.840	0.814	0.590

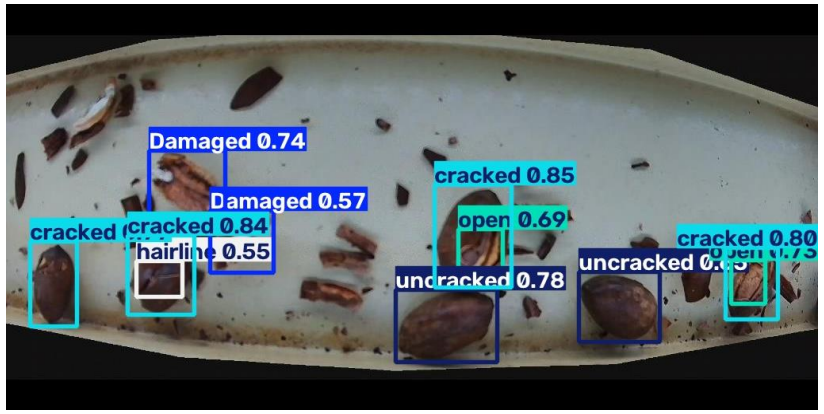

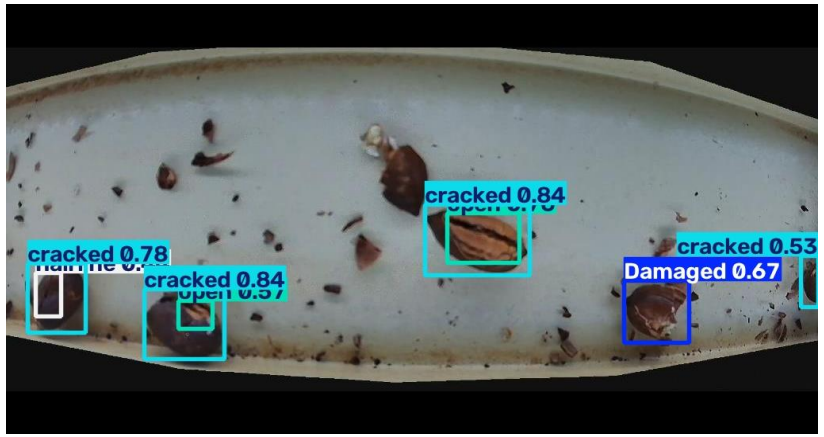

The bottleneck in supervised learning

Computer Vision Model Development

The Bottleneck: Manual, frame-by-frame data annotation.

- *2 minutes for 1 frame*
- *277 frames in a video = 9.3 hours*
- *~1500 videos in total !!!!*

PecanPy – Automated Mask Generation

Interactive video tools for pecan inspection & segmentation

A *napari* plugin that helps generating pseudo-masks for training a robust ML model

Area	Tools
Color prep	Adjustments · Color Thresholding (RGB, HSV, LAB)
Masks	Mask Retouching · Mask Operations
Automation	Pipeline Recorder · Batch Pipeline
ML	Segmentation (YOLO/U-Net++) · SAM 2 · Tracking · Contrastive Coding

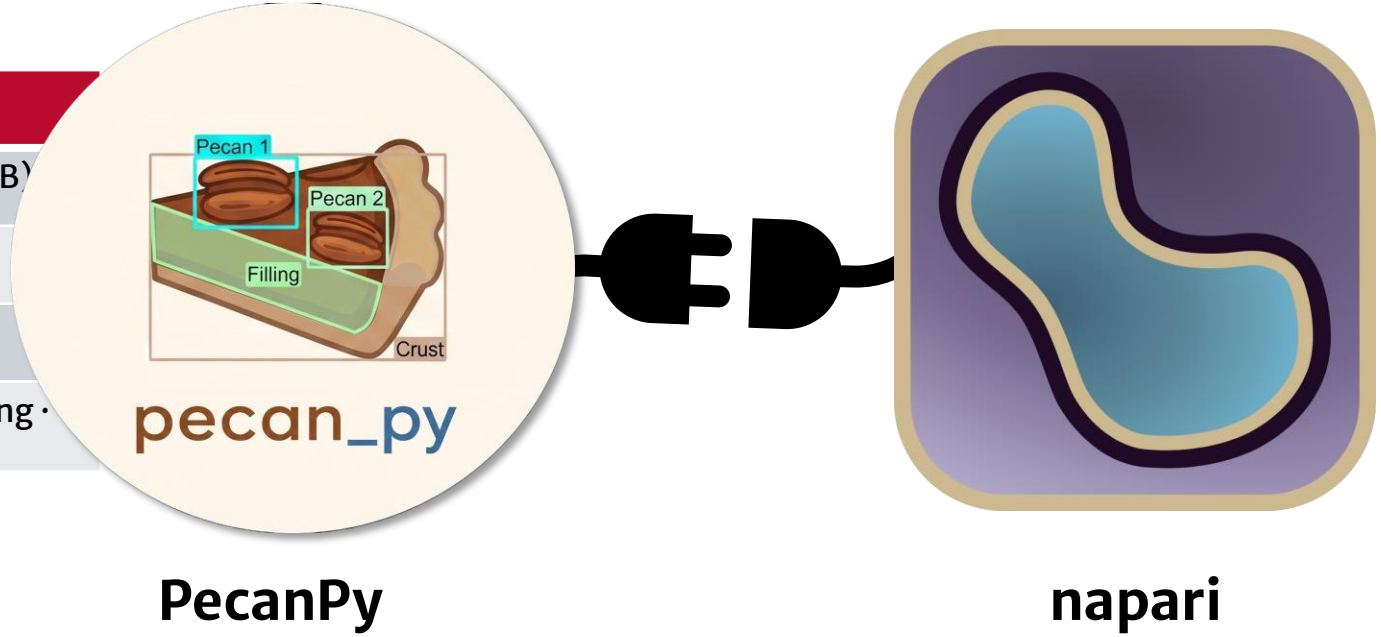

Typical workflow

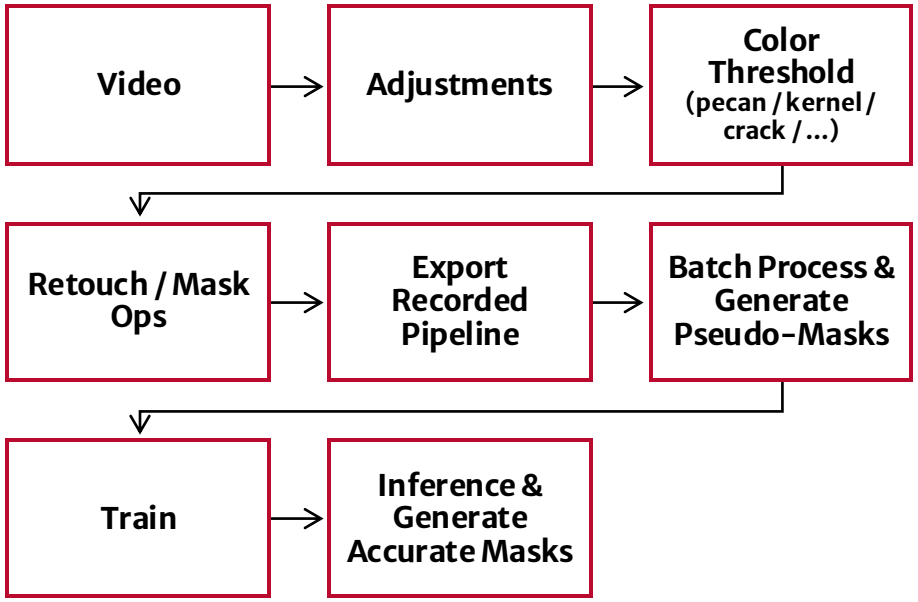

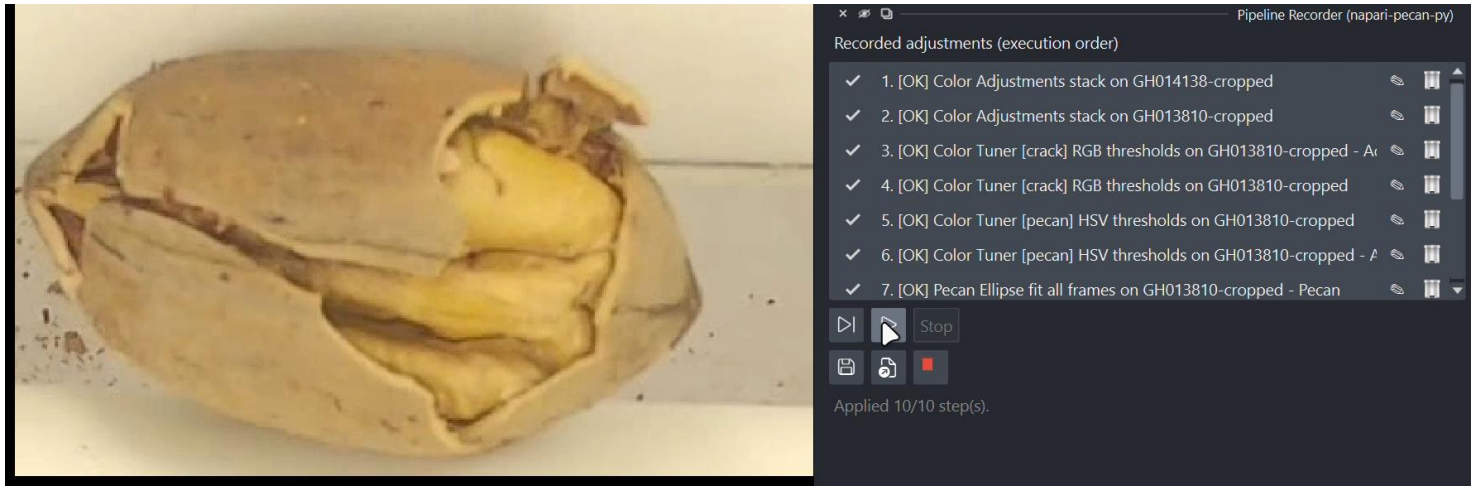

PecanPy Features

- **Tone / contrast**
- **Brightness / Contrast**
- **Levels** (input min/max, gamma, output min/max)
- **Curves (RGB)**
- **Normalization methods**
  - Min-Max (per channel)
  - Percentile stretch
  - Z-score
  - Robust (median/IQR)
  - Unit L2 (per pixel)
  - Luminance Min-Max
  - Histogram Equalization
  - CLAHE (adaptive)

- **Blur / denoise**
- **Surface Blur** (edge-preserving)
- **Denoising methods**
  - Gaussian blur
  - Median
  - Bilateral (edge-preserving)
  - Non-local means
  - Total variation (TV)
  - Wavelet

PecanPy Features

Glare / illumination

Glare / Specular

- HSV V soft roll-off
- Specular detect + inpaint
- Retinex
- Homomorphic
- Dichromatic

PecanPy Features

Method	What it does
Median background	Still pixels look like the clip's median; a moving nut does not. |frame − median| is the motion score.
Frame difference	Subtract a recent frame. Highlights change between consecutive frames, not a static background.
Instance tracking	Match YOLO IDs frame-to-frame so each pecan keeps one label while it crosses left → right.
SAM 2 propagate	Follow a prompted object through the video. Labeling aid, not a motion score.

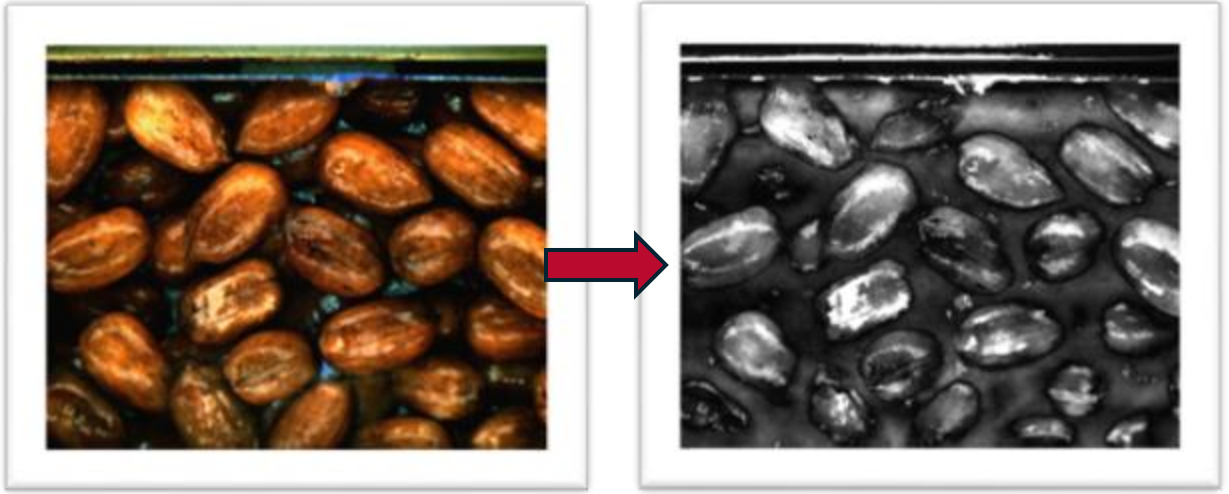

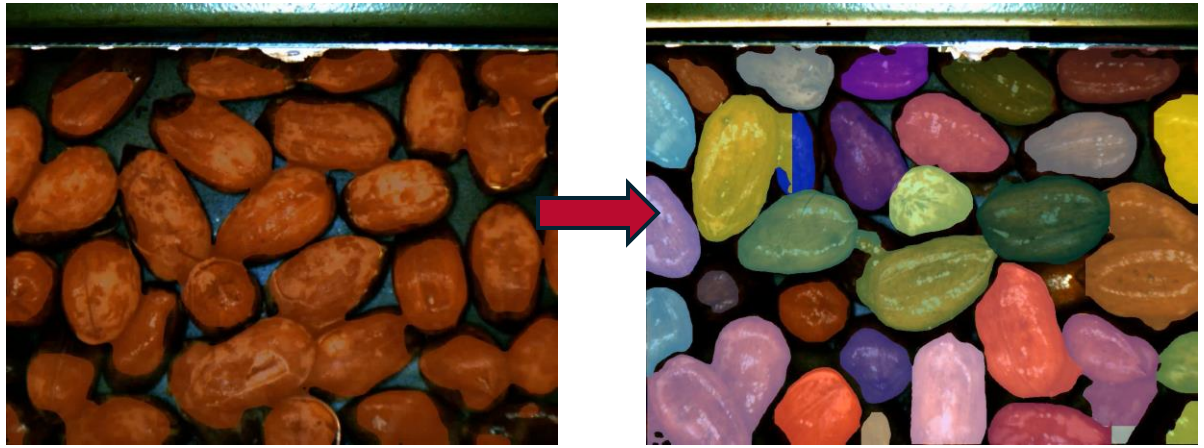

Pecan_Py Features

ML/DL Training & Inference

Model	Architecture	What it does
YOLO-seg	yolov8n-seg	Instance segmentation. Each nut/crack/kernel is a separate polygon; touching objects stay distinct.
U-Net / U-Net++	One SMP network, one output channel per class	Pixel-wise semantic masks. Simpler than cascade.
Cascade U-Net	One U-Net per class, later stages see parent masks	**Pecan** first, then **Crack** and **Kernel** as siblings inside the nut.