OVERVIEW

High-Volume, Medium-Mix Inspection, Visualized.

A visual walkthrough of how synthetic data is changing inspection in high-volume, medium-mix manufacturing. One full Contaminate Verification deployment on a real production line, plus the Style Transfer mode that scales it across product variants.

CONTENTS

What's inside.

Four manufacturing applications, one paper. A 15-minute walk-through on a food-safety belt, plus deep dives on connectors, wood flooring, and car seats. Headline numbers at the end.

INTRODUCTION

APPLICATION 01 · CONTAMINANT VERIFICATION

STYLE TRANSFER APPLICATIONS

RESULTS & NEXT STEPS

Real inspection lines, trained on synthetic defects.

In the next twenty pages we walk through four manufacturing applications running today on Overview AI cameras. We start with a single 15-minute Contaminate Verification deployment, narrated frame by frame from clean oatmeal pile to a model catching a real nail dropped live on the belt. Then we go wide: 300+ connector SKUs trained without breaking a single housing, wood flooring scratch inspection retargeted across finishes via Style Transfer, and a car seat trim line where one labeled wrinkle covers every fabric variant in the OEM catalog.

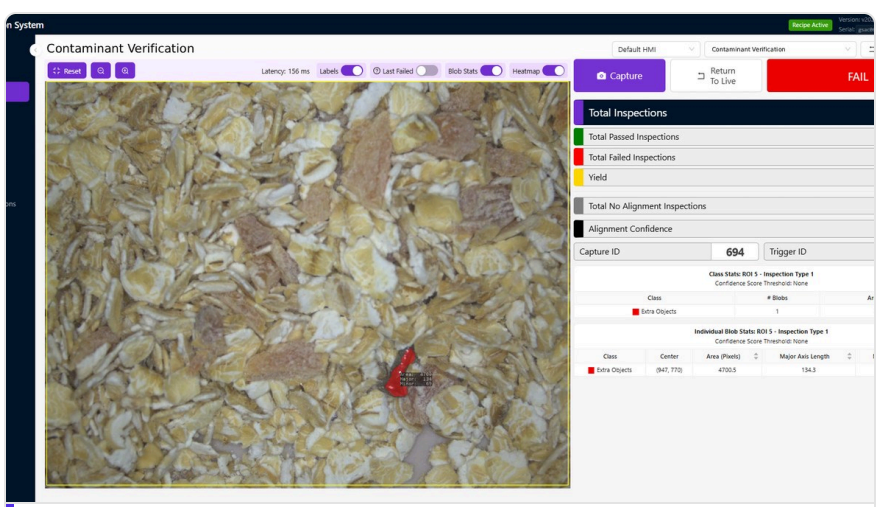

APPLICATION 01 · FOOD SAFETY

Contaminate Verification

Clean oatmeal reference plus synthetic contaminants. The model caught a real nail dropped live on the belt.

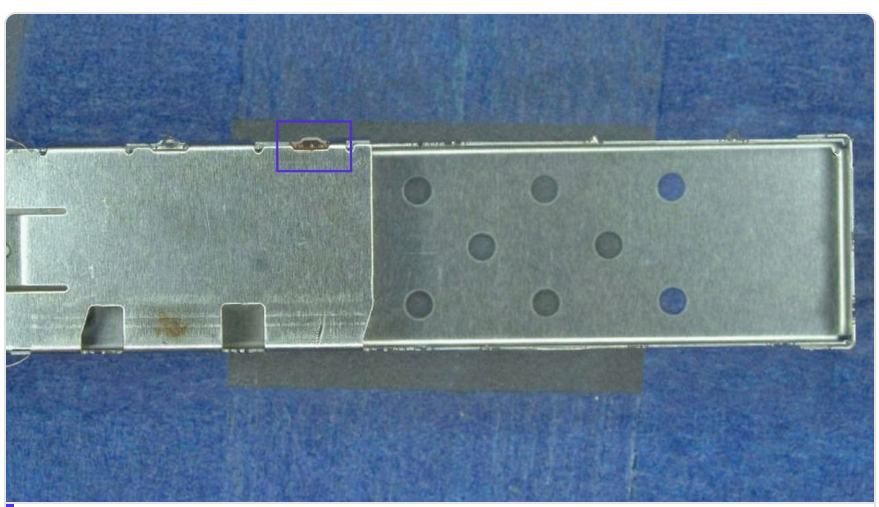

APPLICATION 02 · ELECTRONICS

300+ Connector SKUs

Every defect class authored in the studio. Trained without breaking a single housing.

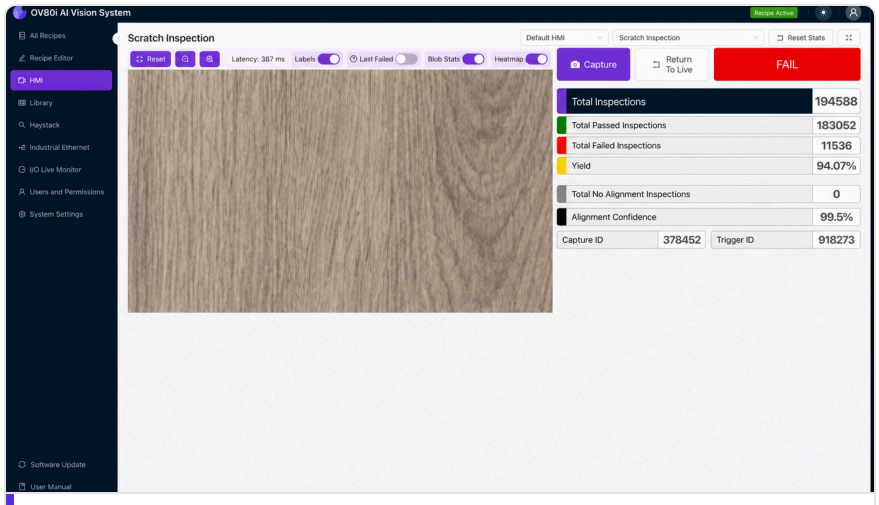

APPLICATION 03 · ENGINEERED FLOORING

Wood Flooring

One scratch retargeted across wood, laminate, and composite finishes via Style Transfer.

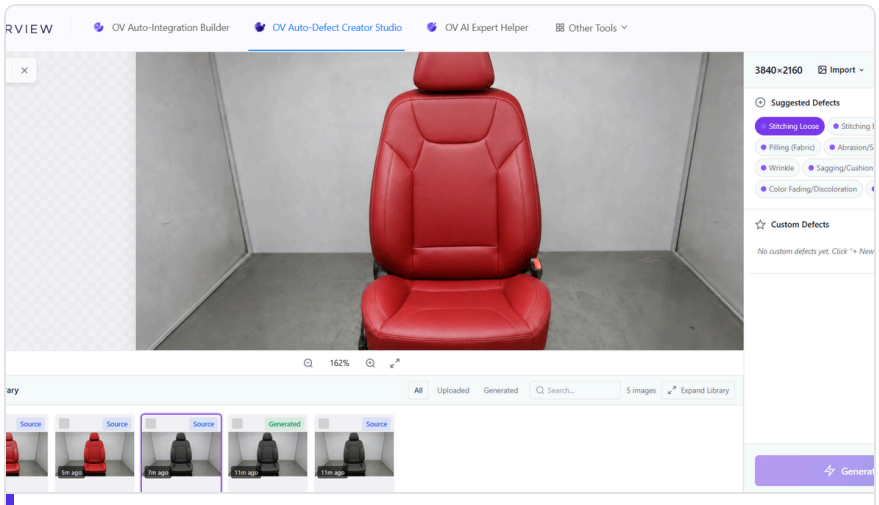

APPLICATION 04 · AUTOMOTIVE TRIM

Car Seat

One labeled wrinkle covers every fabric and color variant in the OEM catalog.

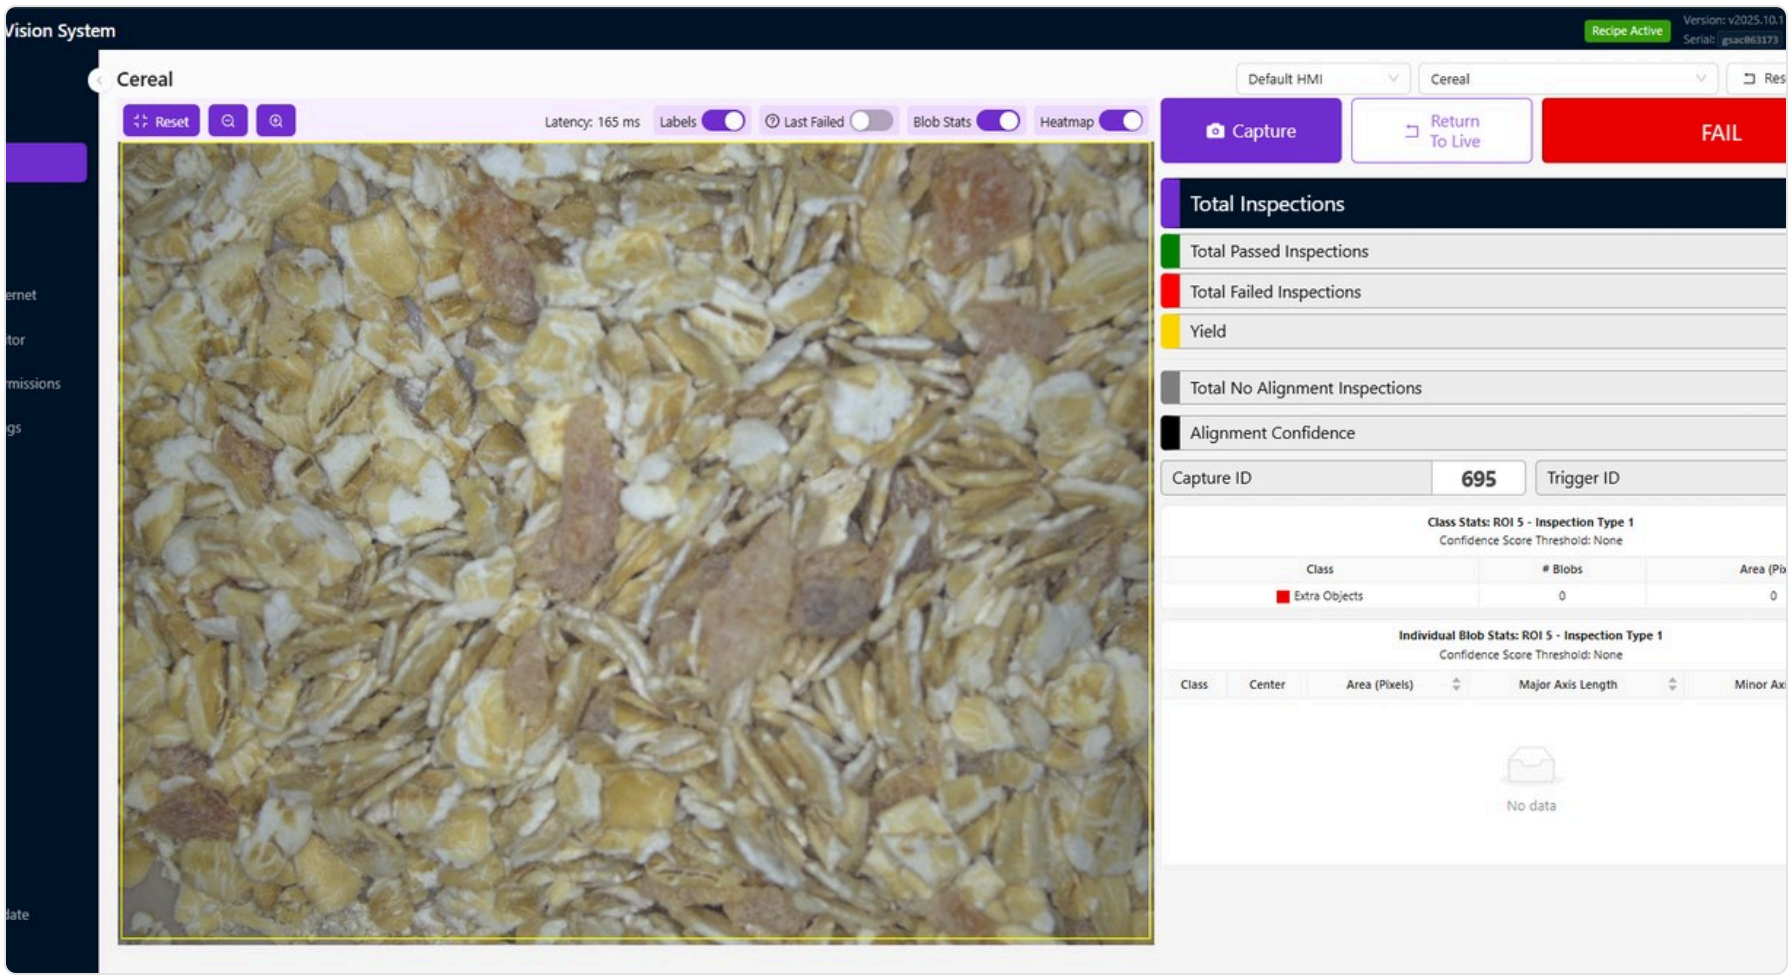

APPLICATION 01 · PURE DEFECT GENERATION

Contaminate Verification.

Real good images. Synthetic defects only. We trained the model on a clean oatmeal reference set plus contaminants authored entirely inside the studio, then walked up to the line and dropped in real metal, real plastic, real nails. The model caught every one.

The first five clicks.

Mount the camera, skip the aligner for pile inspection, set the ROI to the whole frame, capture a baseline reference set, and load it into the OV Auto-Defect Creator Studio.

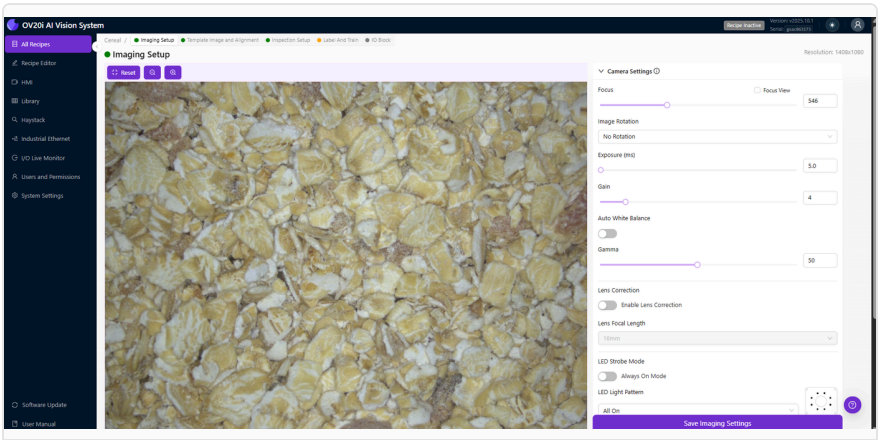

STEP 01 · CAMERA

Frame, expose, light.

OV80i mounted above the belt, imaging tuned for the pile texture.

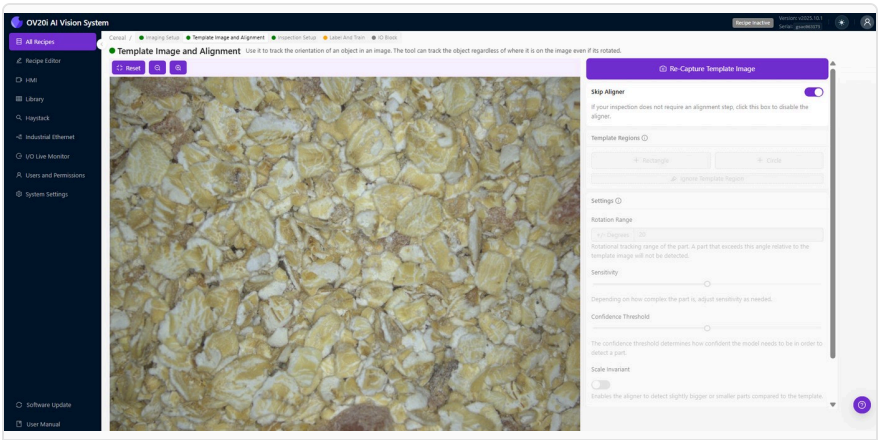

STEP 02 · ALIGNER

Skip for piles.

Alignment matters for discrete parts. A pile is always a pile in the frame.

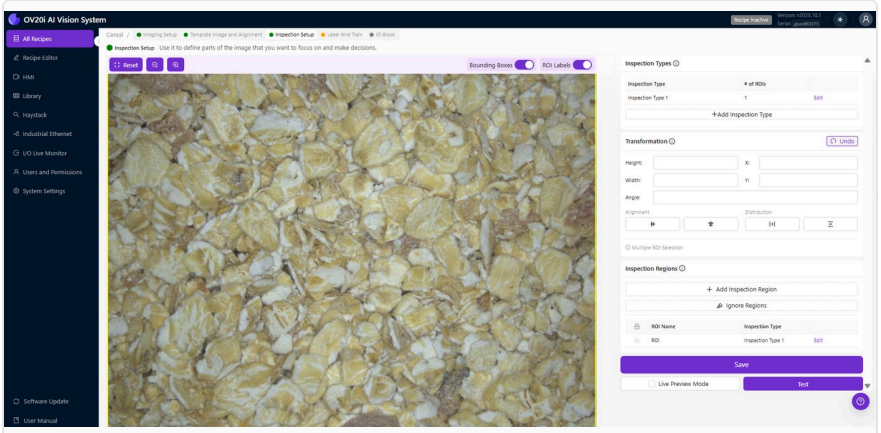

STEP 03 · ROI

Whole frame.

Smaller 512×512 ROIs work better on parts. Pile inspection uses the whole frame.

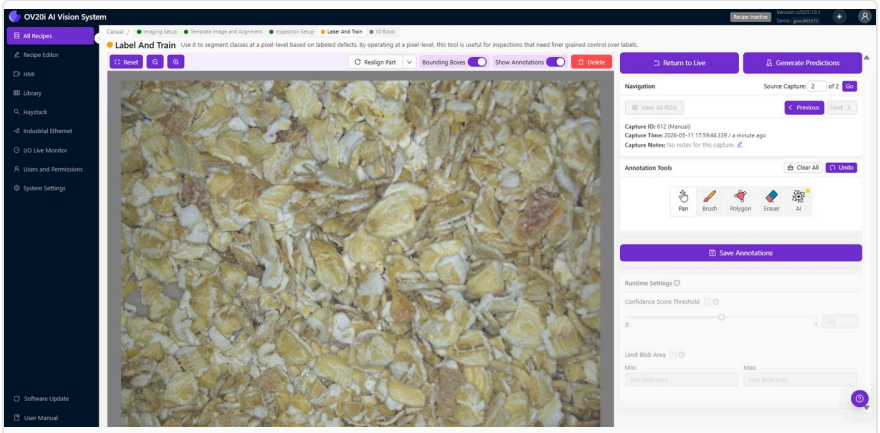

STEP 04 · CAPTURE SET

13 captured + 13 generated.

A 26-image baseline of clean pile variations, no defect samples needed.

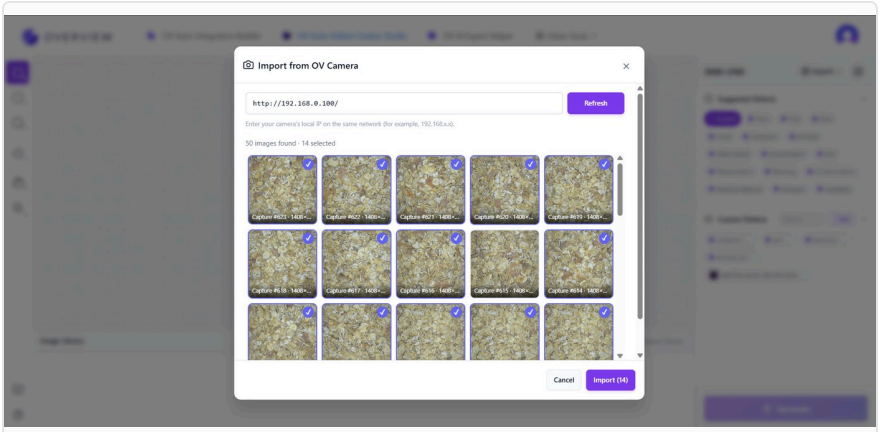

STEP 05 · LOAD STUDIO

Library imported.

Camera library flows straight into the Auto-Defect Creator. No exports.

UP NEXT

Author the first defect.

Region + class + prompt. The studio renders the synthetic defect photorealistically on the pile.

One session, dozens of defects.

Each defect is a plain-English prompt drawn on the area where it should appear (e.g., *"Add a quarter into the pile"*). The studio anchors to the grain texture and renders the synthetic defect surface-aware, photorealistic, and auto-annotated with a pixel-level mask. Three generations run in parallel while you author the next prompt. Renaming a class is a single click, so the catalog reaches a dozen labeled defects in the time it took to brew coffee.

RENAME · SINGLE CLICK

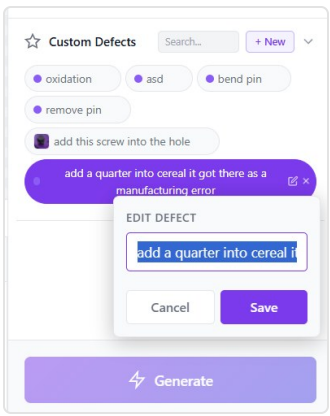

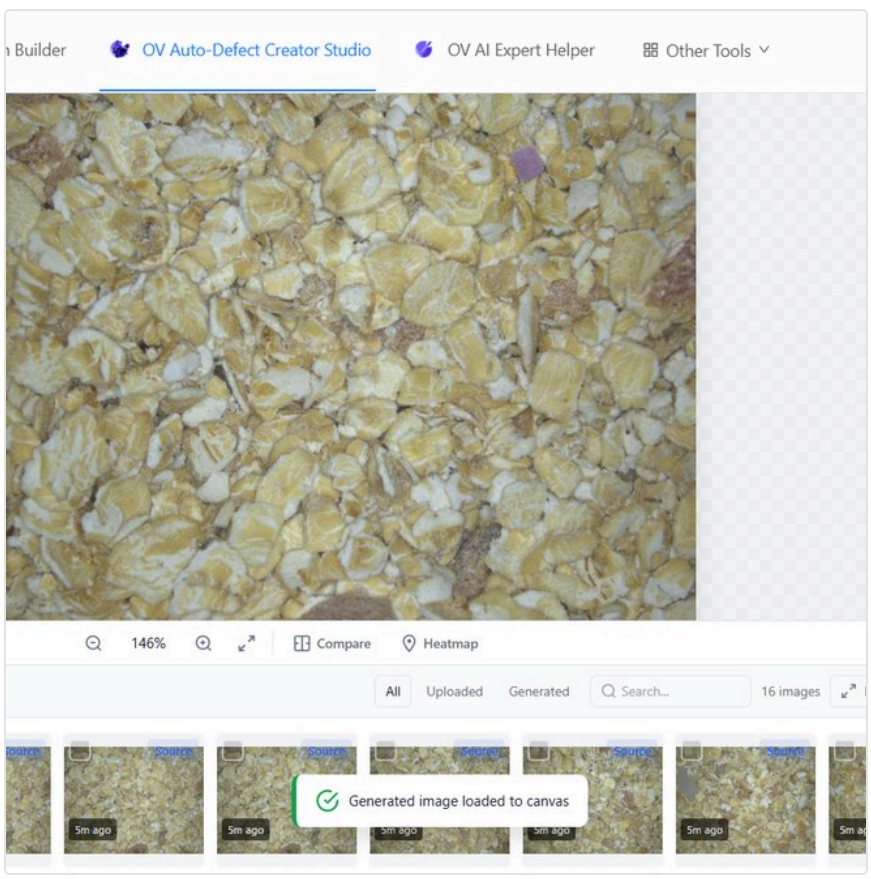

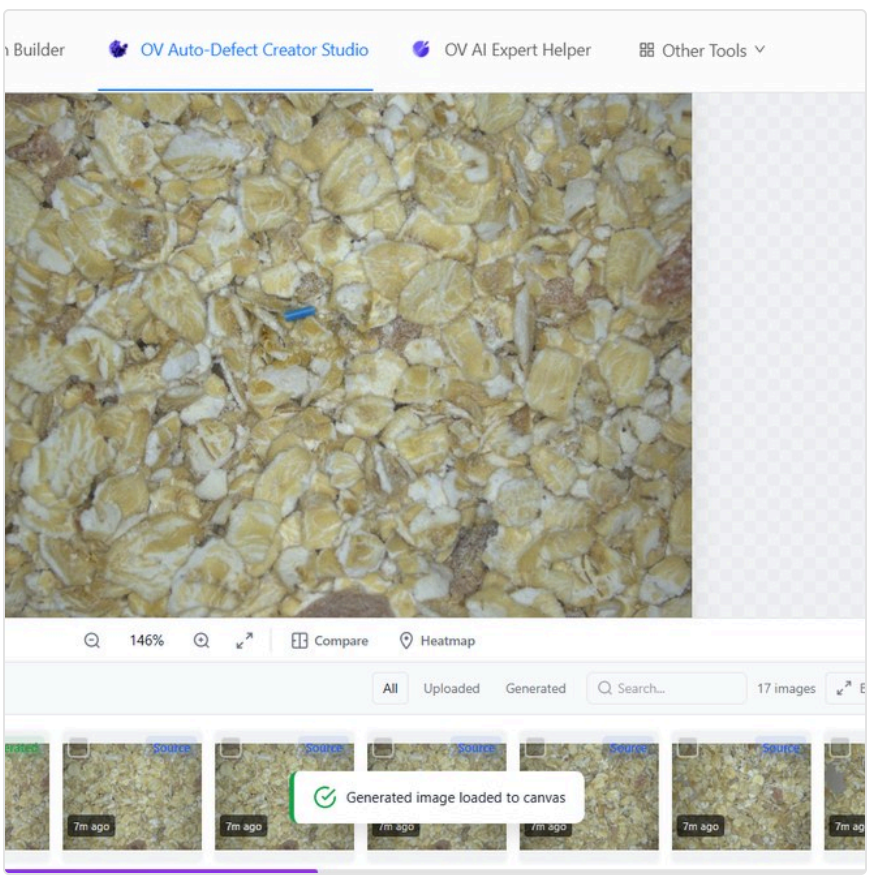

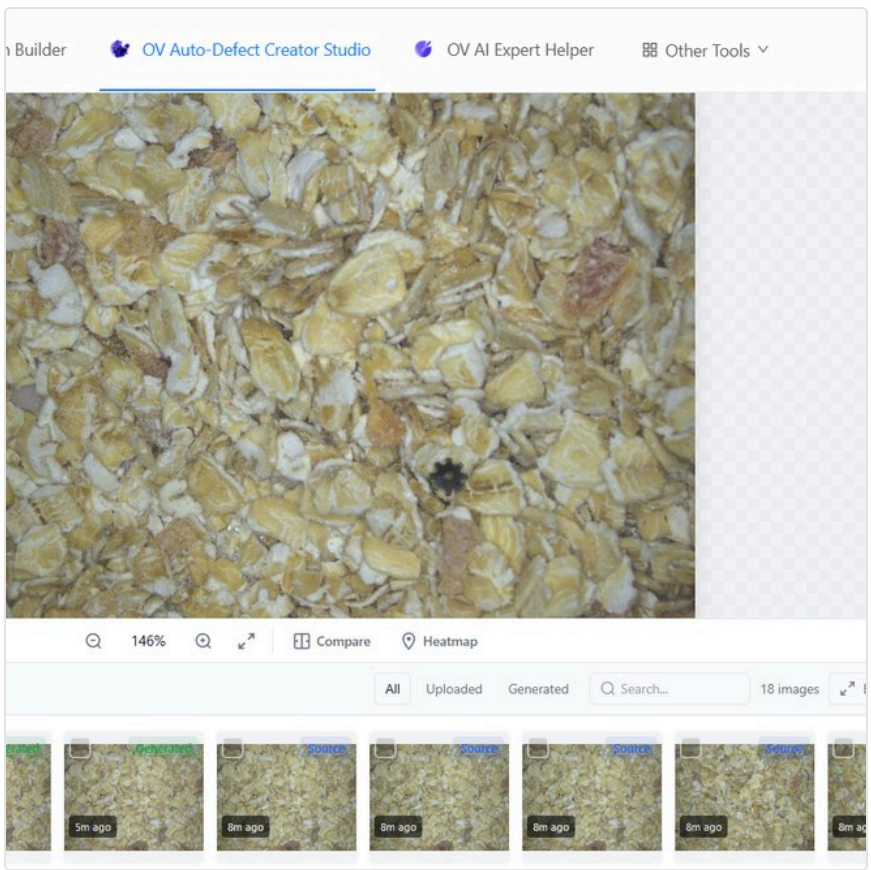

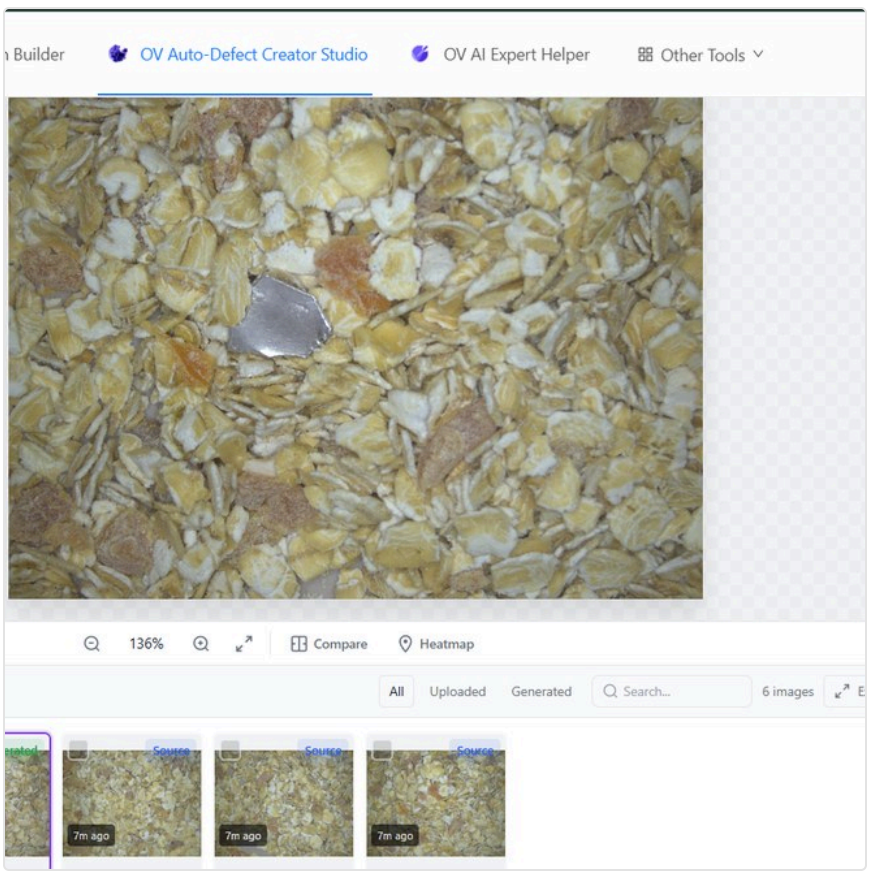

Pixel-level ground truth, marked by hand.

Once the images are on the camera, the built-in annotation tool lets you segment each defect, painting its exact pixel boundaries on top of the bounding box. Below: thirteen examples masked this way.

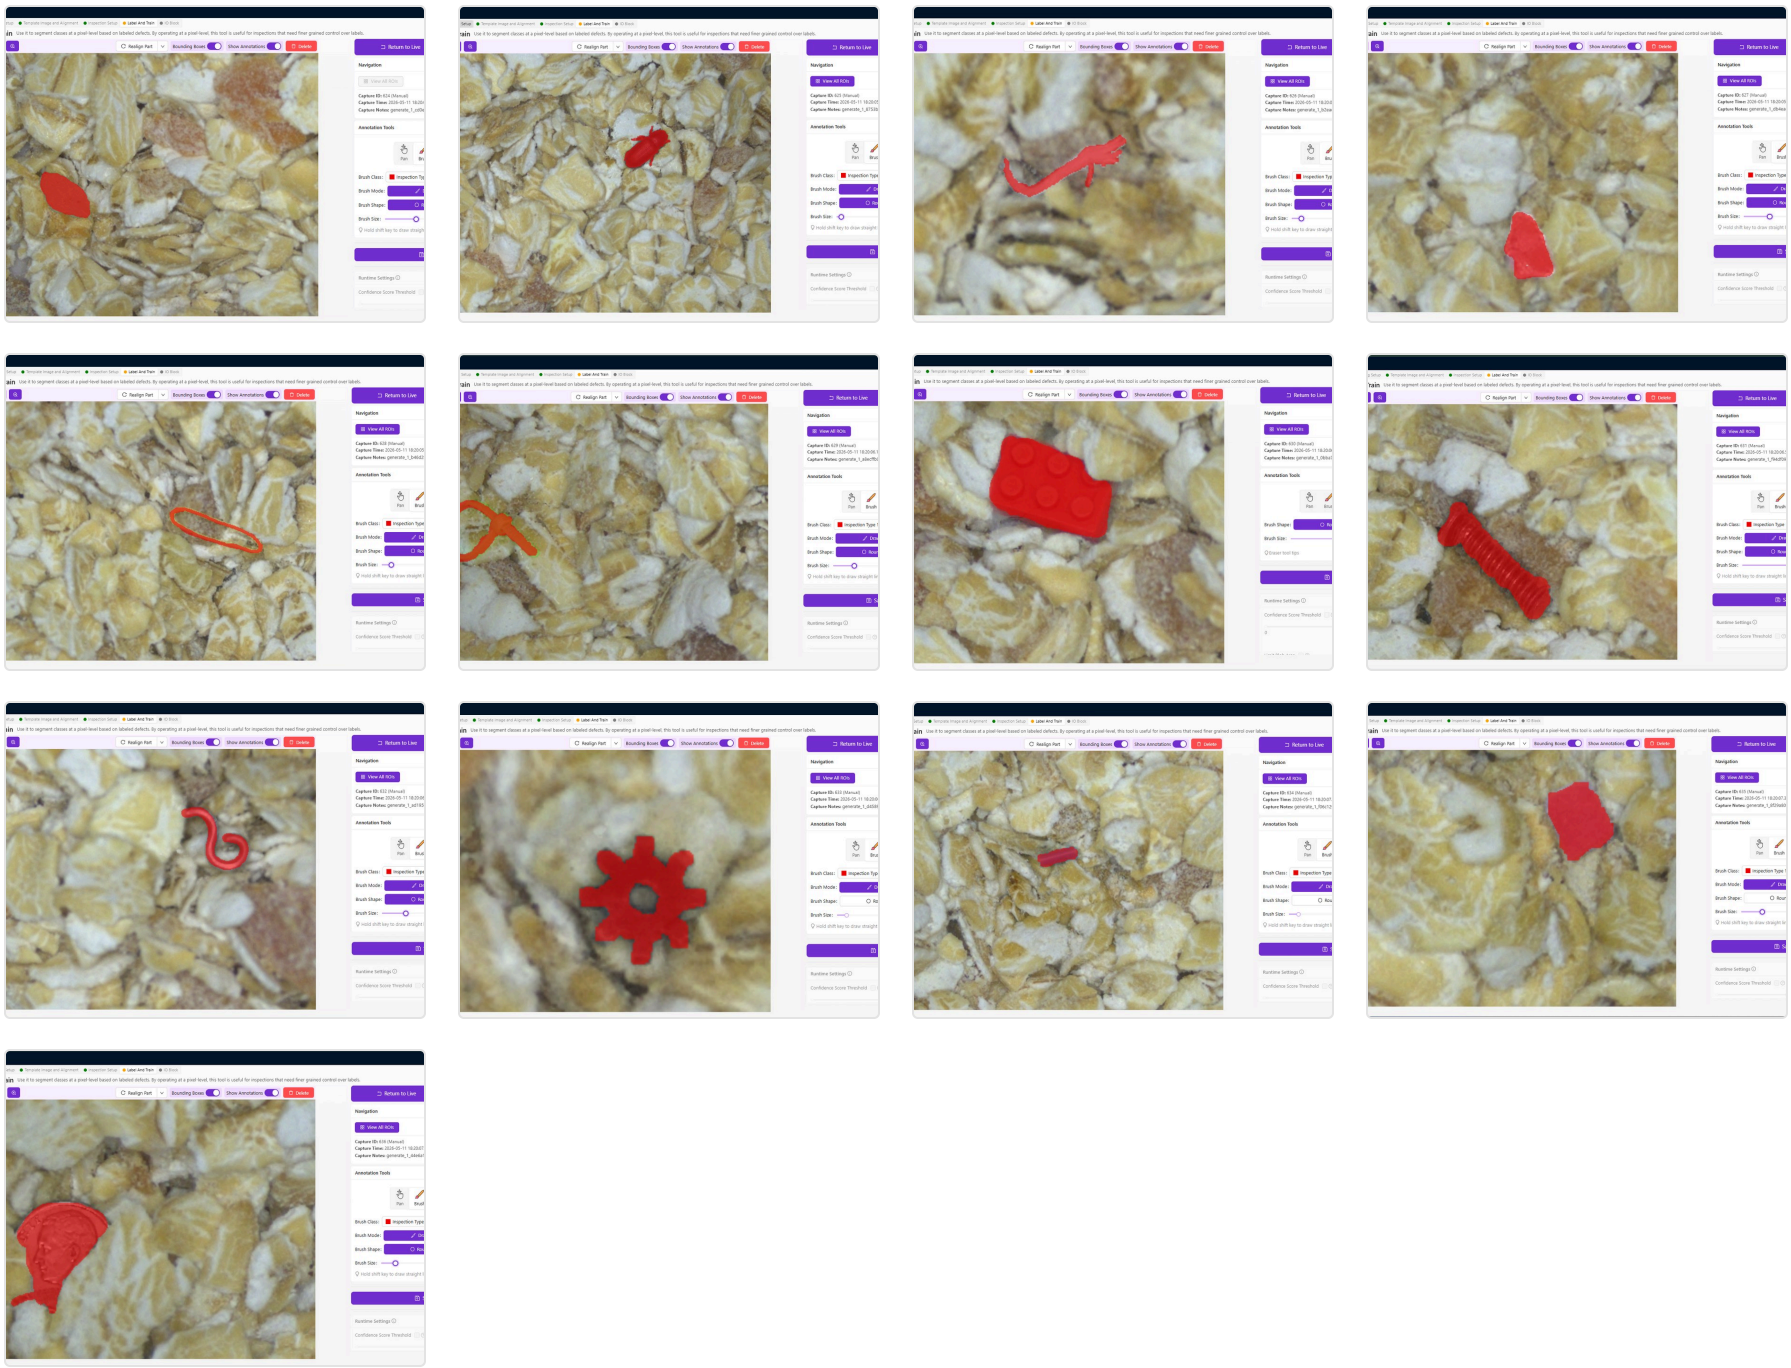

From studio library to running recipe.

The studio library moves to the camera as a single labeled bundle. No file shuffling, no manifest editing. Training runs locally on the OV80i. When validation passes, the recipe deploys straight to the camera.

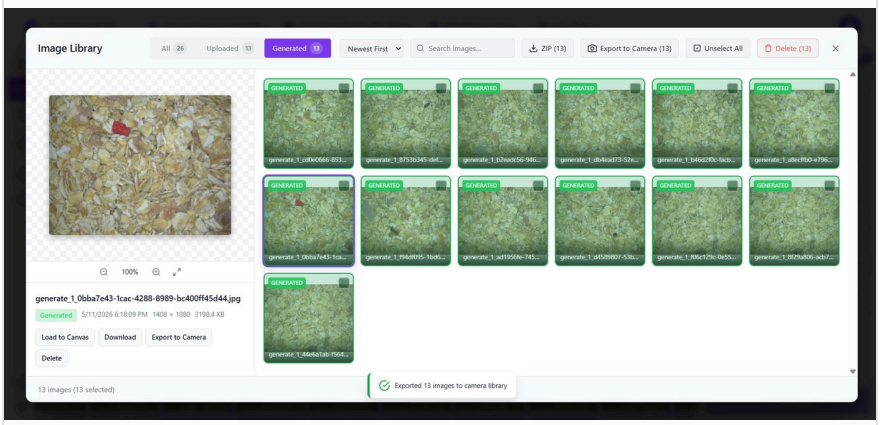

STEP 01 · EXPORT

Studio library to camera.

Dataset and labels move as one bundle. No conversion.

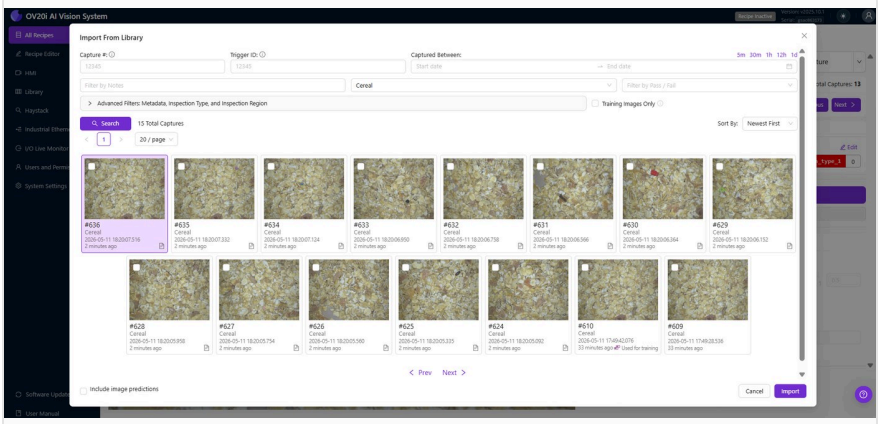

STEP 02 · ARRIVED

Camera library = the dataset.

The OV80i sees the same labeled images the studio generated.

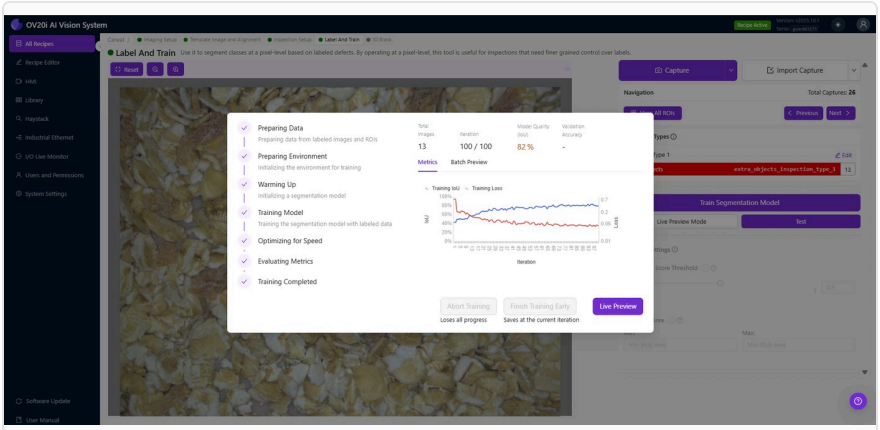

STEP 03 · TRAIN

Local training on the edge.

Runs on the device. No cloud round-trip, no data leaves the plant.

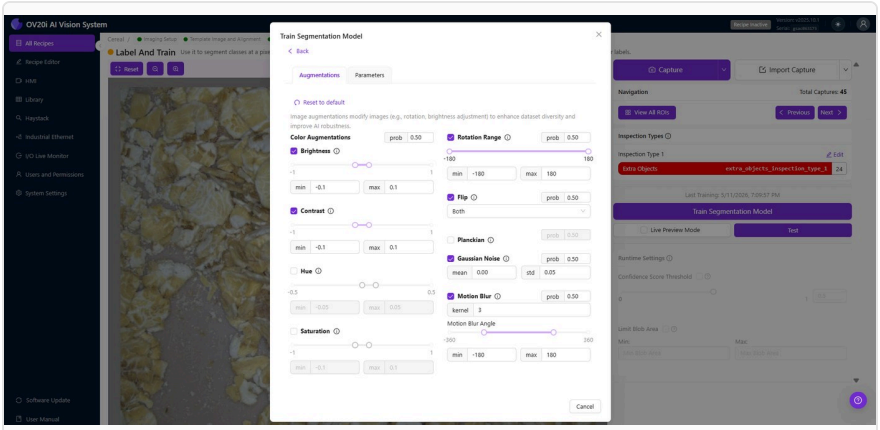

STEP 04 · TUNE

Sensible defaults, full control.

Production teams ship on defaults nine times out of ten. Every parameter is exposed when needed.

One-click deploy to the camera

A passing validation run sends the recipe to the OV80i directly from the studio. Redeploying a new model takes minutes, not a maintenance window.

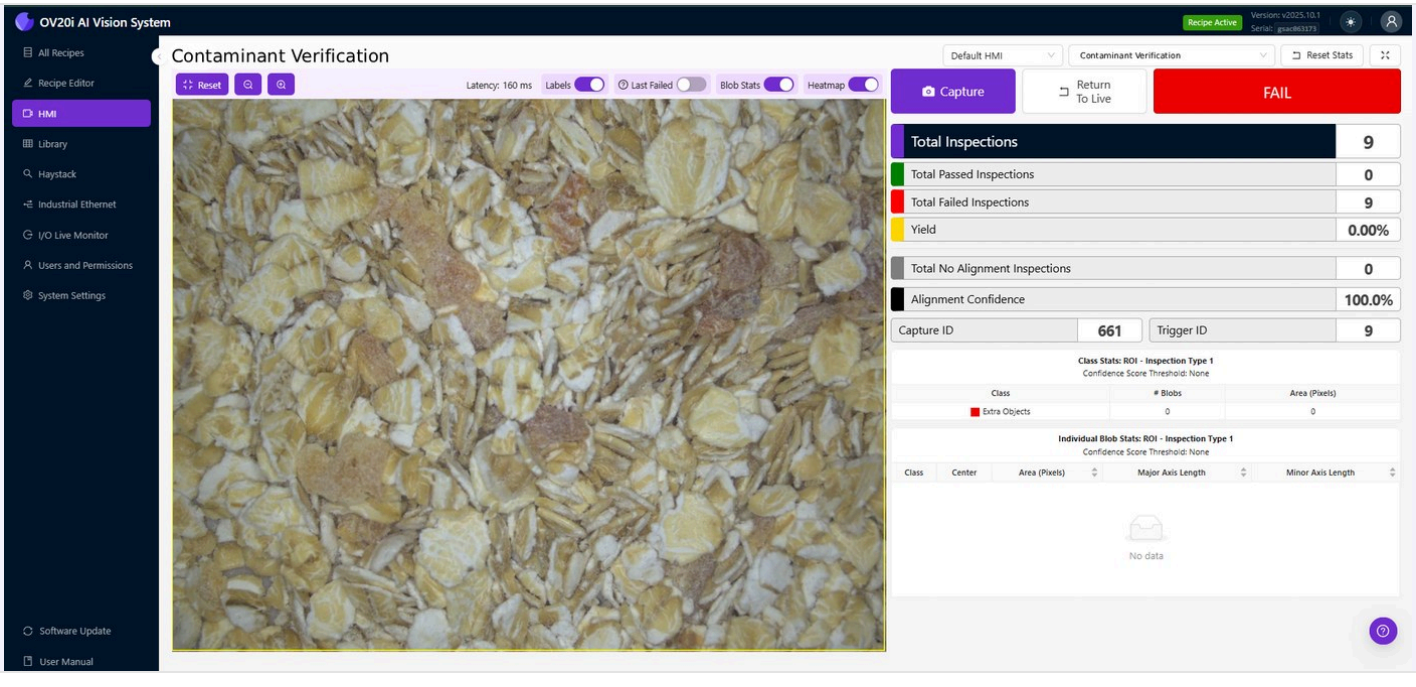

FIRST LIVE FRAME

Clean pile. 0 defect pixels. PASS.

The recipe is live on the OV80i. The first scrambled pile reads clean, exactly as expected. The model has learned what "good" oatmeal looks like, including the texture variability of normal product flow. Operator confidence in the system starts here.

0

Defect pixels on the first clean frame. PASS verdict, no operator intervention required.

Real objects. Trained on synthetic. Caught on the first frame.

Each row below is a real foreign object placed by hand on the belt under production lighting. The model never saw a real defect during training, only synthetic ones authored in the studio. It catches every one on the first frame.

LIVE #1

Green Plastic

Foreign material placed on the belt by hand. Detected, classified, called out.

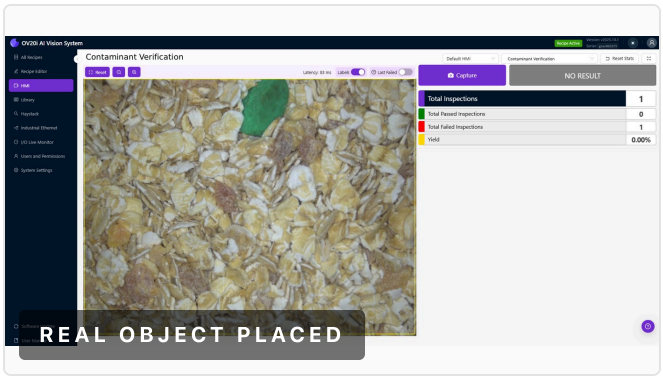

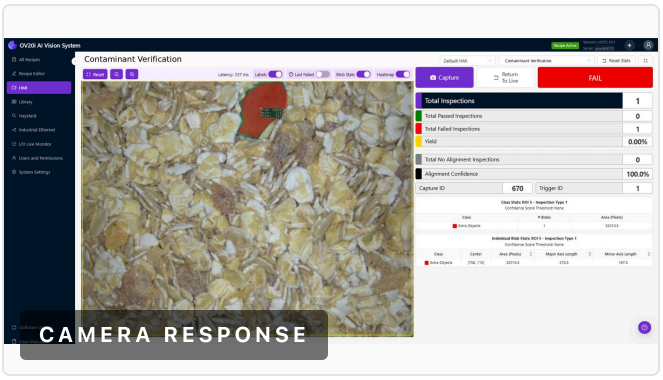

LIVE #2

Metal Piece

Reflective intruder under production lighting. Different reflectance, same outcome.

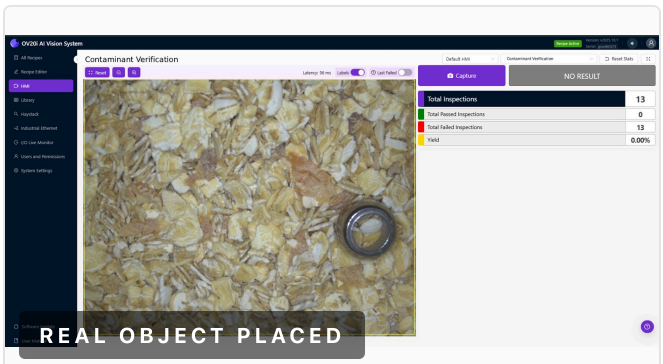

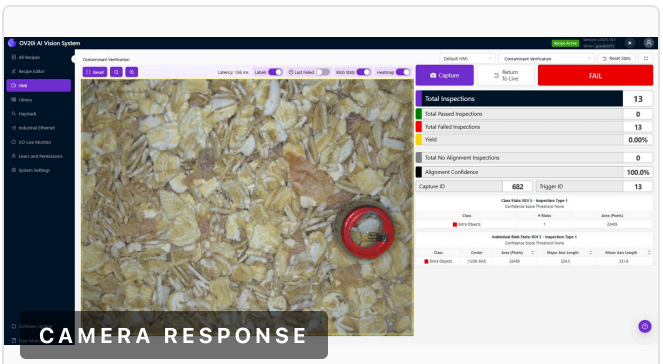

LIVE #3

Real Nail

The food-safety hard case. Sharp metal, textured belt. Caught and annotated red.

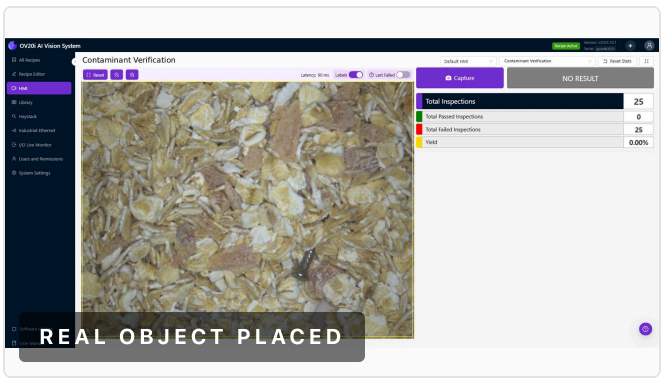

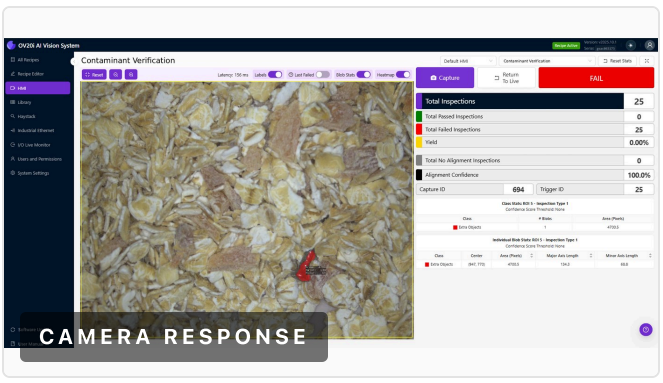

From product change to deployed inspection.

The Contaminate Verification walkthrough on the previous pages took 15 minutes from the first clean image to a model catching real foreign objects on the belt. The classical computer vision path on a food-safety line is a multi-month data-collection and retraining cycle. Across Overview AI customer programs the average speedup is 12.4×. On the food-safety line in this paper, it is far more.

TIME FROM PRODUCT CHANGE TO A RUNNING INSPECTION MODEL

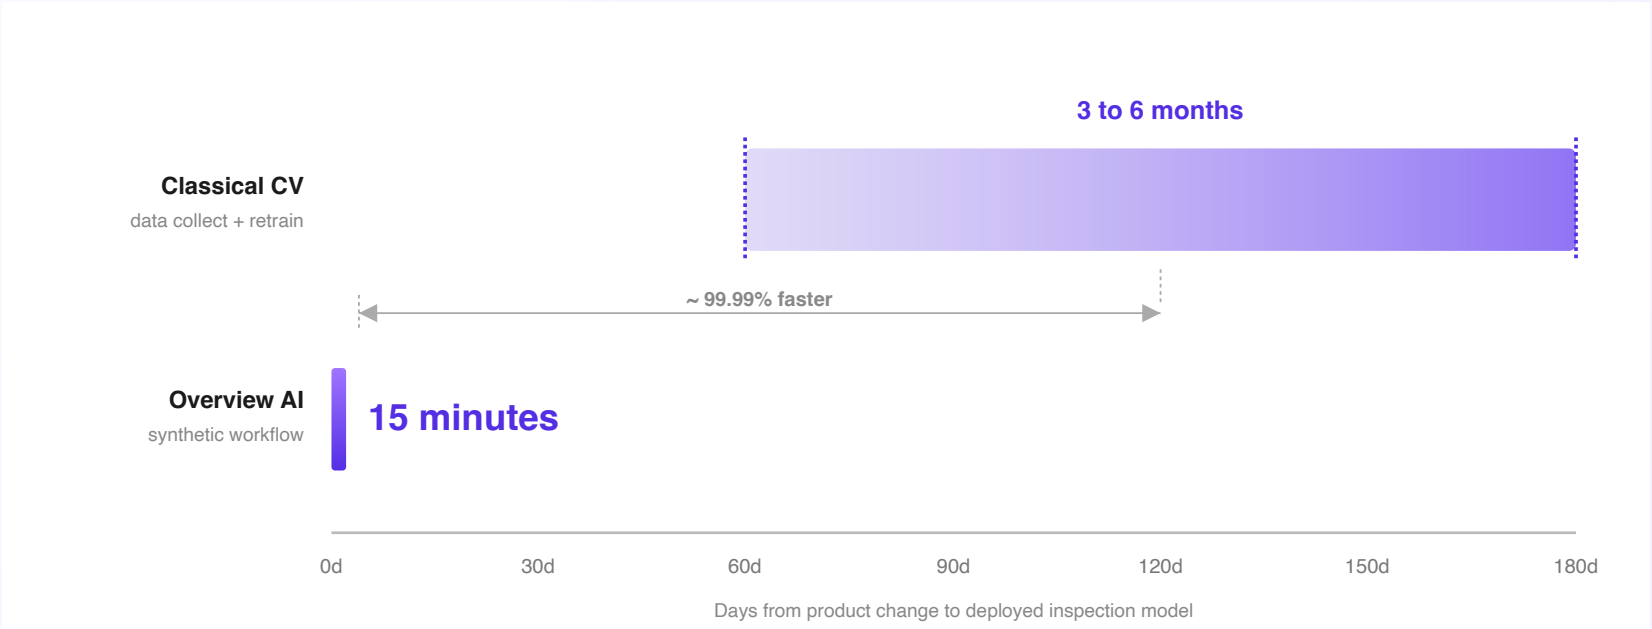

Manual range covers typical food-line retraining cycles where defect samples are collected from production runs and the model is retrained. Overview AI time is the end-to-end walk shown earlier in this paper.

12.4× faster on average. 15 minutes on the line in this paper.

12.4× is the across-the-board average across Overview AI customer programs. The Contaminate Verification number above is the elapsed time on the walk in this whitepaper.

One real housing, eight synthetic defect classes.

A connector inspection line needs coverage across many failure modes: scratch, dent, tool mark, oxidation, corrosion, chip, stain, porosity. Collecting or producing real samples of each class on real housings is slow and destroys parts. Here, one clean OV80i capture seeds the entire defect catalog. Each frame below adds one synthetic defect to the previous frame, with a pixel-level mask attached. Nine frames, eight defect classes, one real housing.

STARTED CLEAN

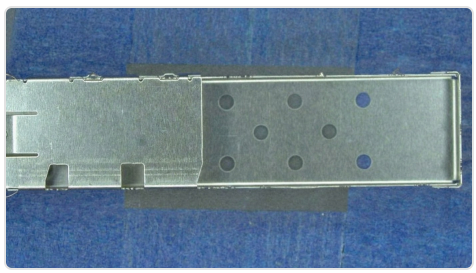

01

Good

REAL OV80I CAPTURE

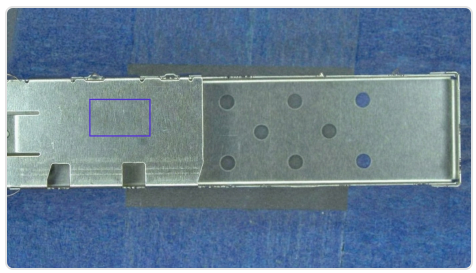

02

Scratch

+1 DEFECT CLASS

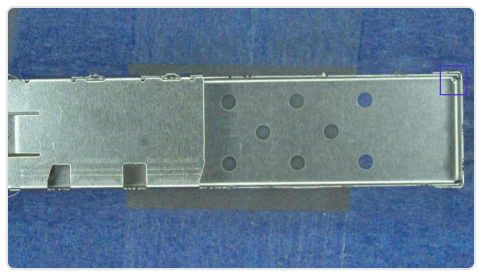

03

Dent

+2 DEFECT CLASSES

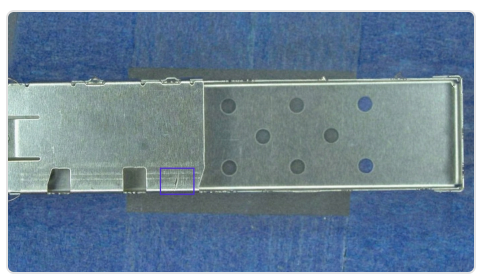

04

Tool Mark

+3 DEFECT CLASSES

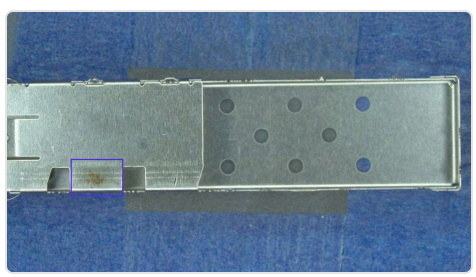

05

Oxidation

+4 DEFECT CLASSES

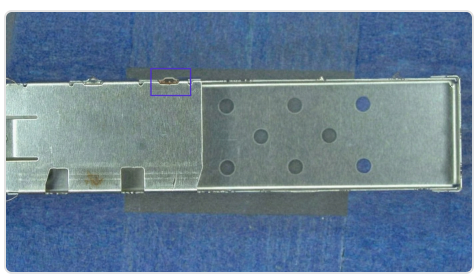

06

Corrosion

+5 DEFECT CLASSES

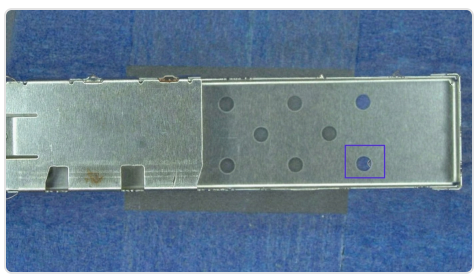

07

Chip

+6 DEFECT CLASSES

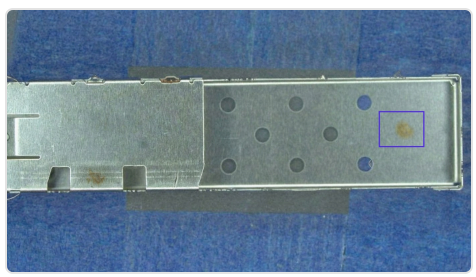

08

Stain

+7 DEFECT CLASSES

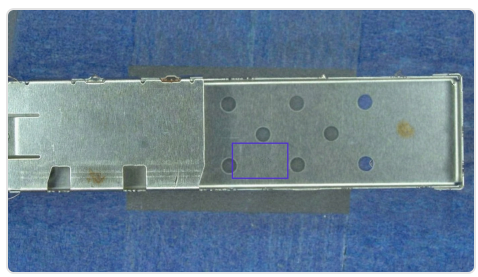

09

Porosity

+8 DEFECT CLASSES

ONE STUDIO SESSION · EIGHT CLASSES FROM ONE CAPTURE

8 defect classes from 1 real capture. 0 housings damaged.

Every newly added defect carries its own pixel-level segmentation mask. The cumulative training image (frame 09) feeds the multi-class inspection recipe directly — no per-class data collection, no manual labeling.

Wood Flooring: one good plank, every finish, every defect.

Start with one real OV80i image of a clean plank. Style Transfer renders that plank in cherry and oak. Defect Transfer adds scratches, scuffs, and peel delamination, all in the same surface location, propagated across every finish in the catalog. Twelve labeled training images from a single real capture.

STYLE AND DEFECT TRANSFER

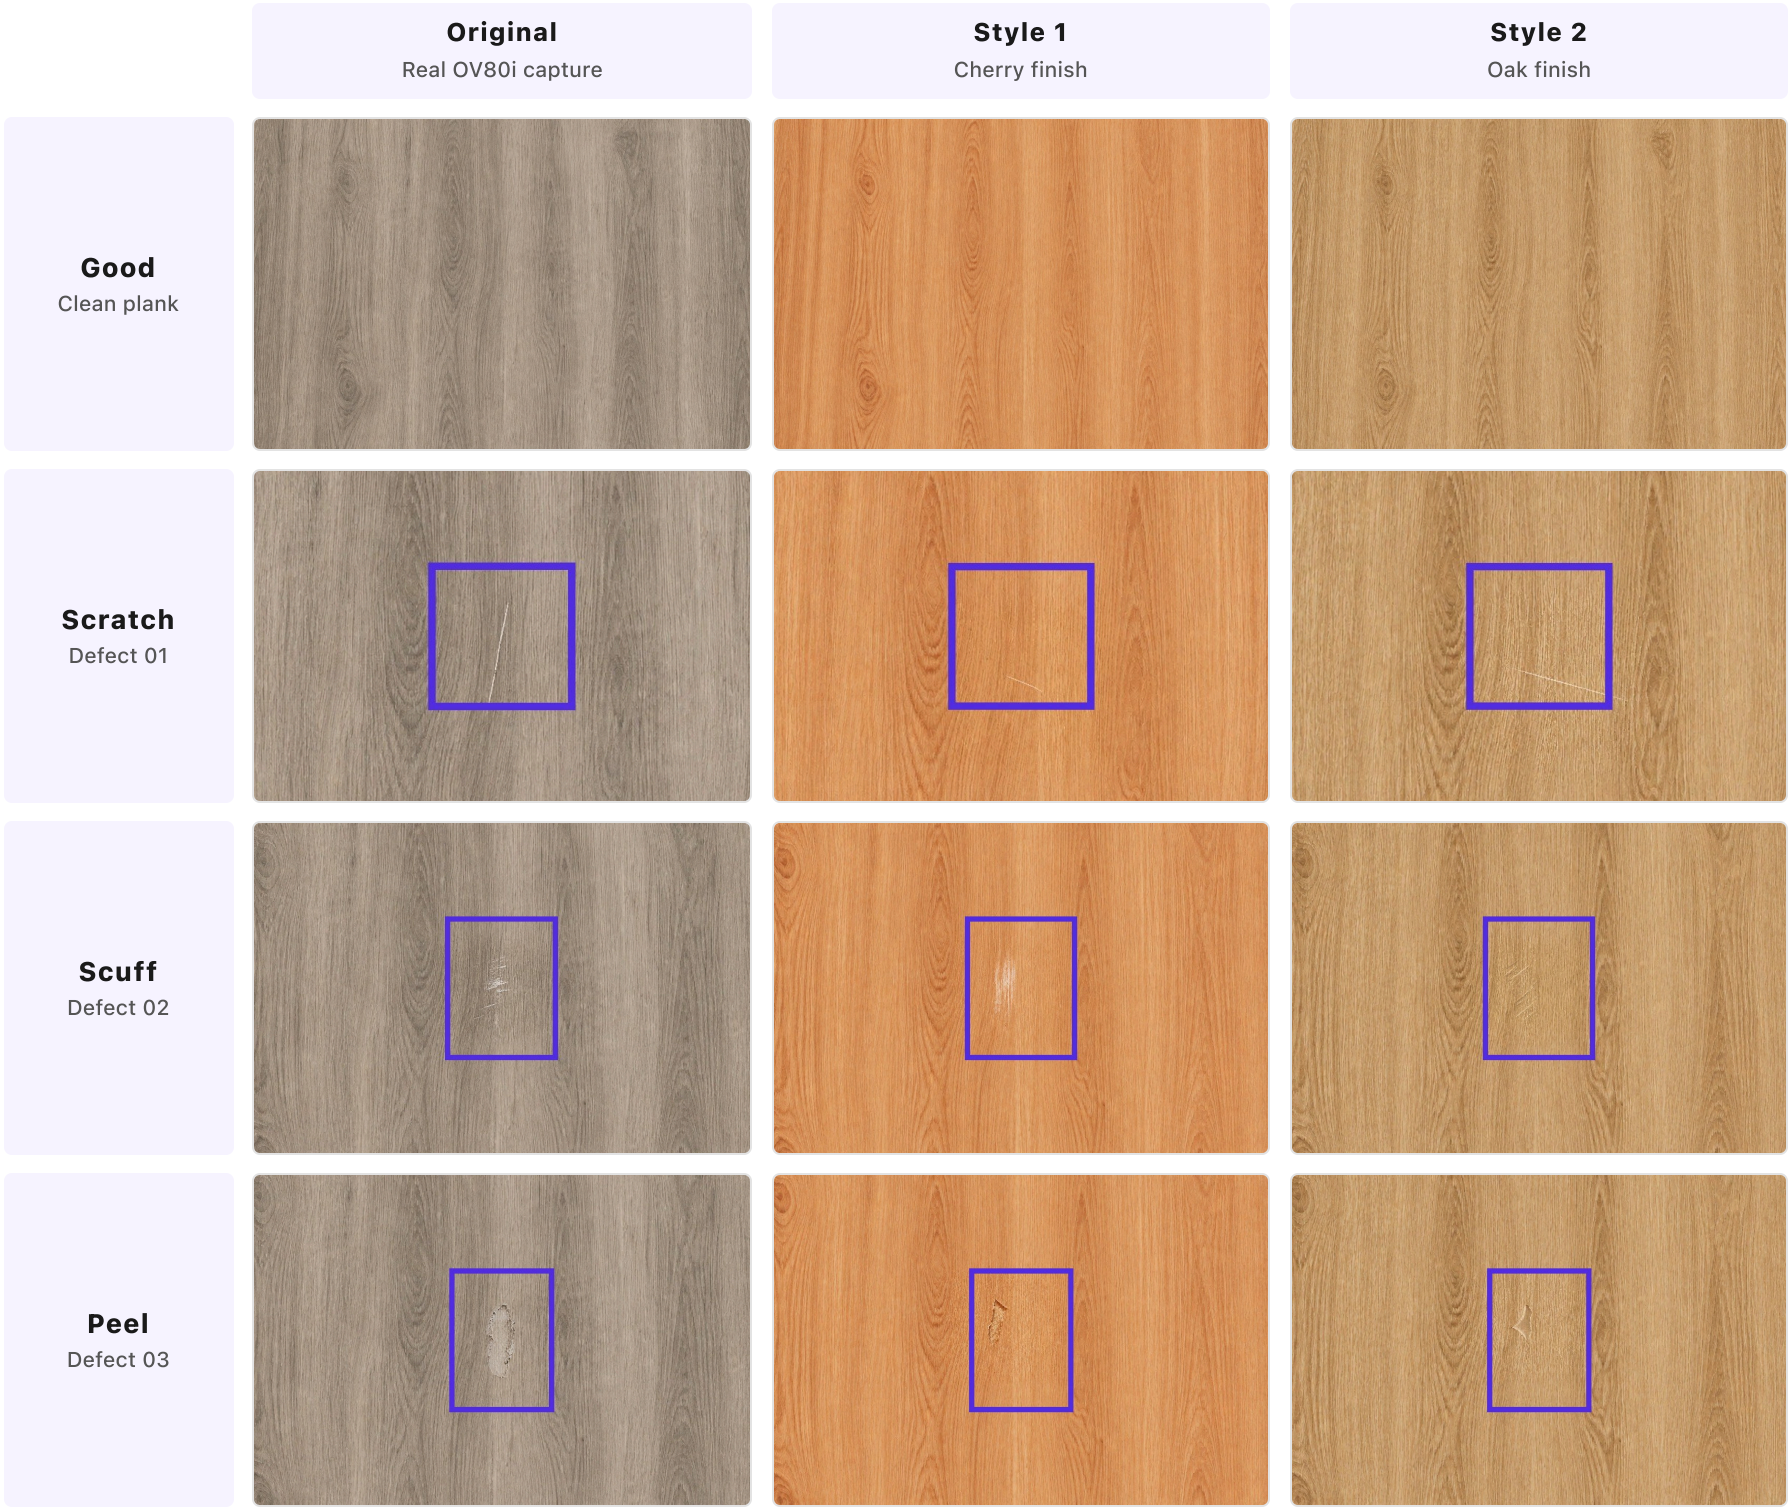

Defect region marked by the studio → Style Transfer across finishes ↓ Defect Transfer across classes

3 finishes × 4 defect classes = 12 labeled training images from 1 real capture.

Add a new finish to the catalog? Apply Style Transfer once more. Every defect class propagates automatically. ~12.4× faster than per-finish data collection, zero planks intentionally damaged.

Car Seats: one good seat, every fabric, every defect.

Start with one real OV80i capture of a clean leather seat. Style Transfer renders the same seat in blue fabric and vinyl. Defect Transfer adds scratches, stains, and wrinkles in the same surface location, propagated across every fabric variant. Twelve labeled training images from a single real capture.

STYLE AND DEFECT TRANSFER

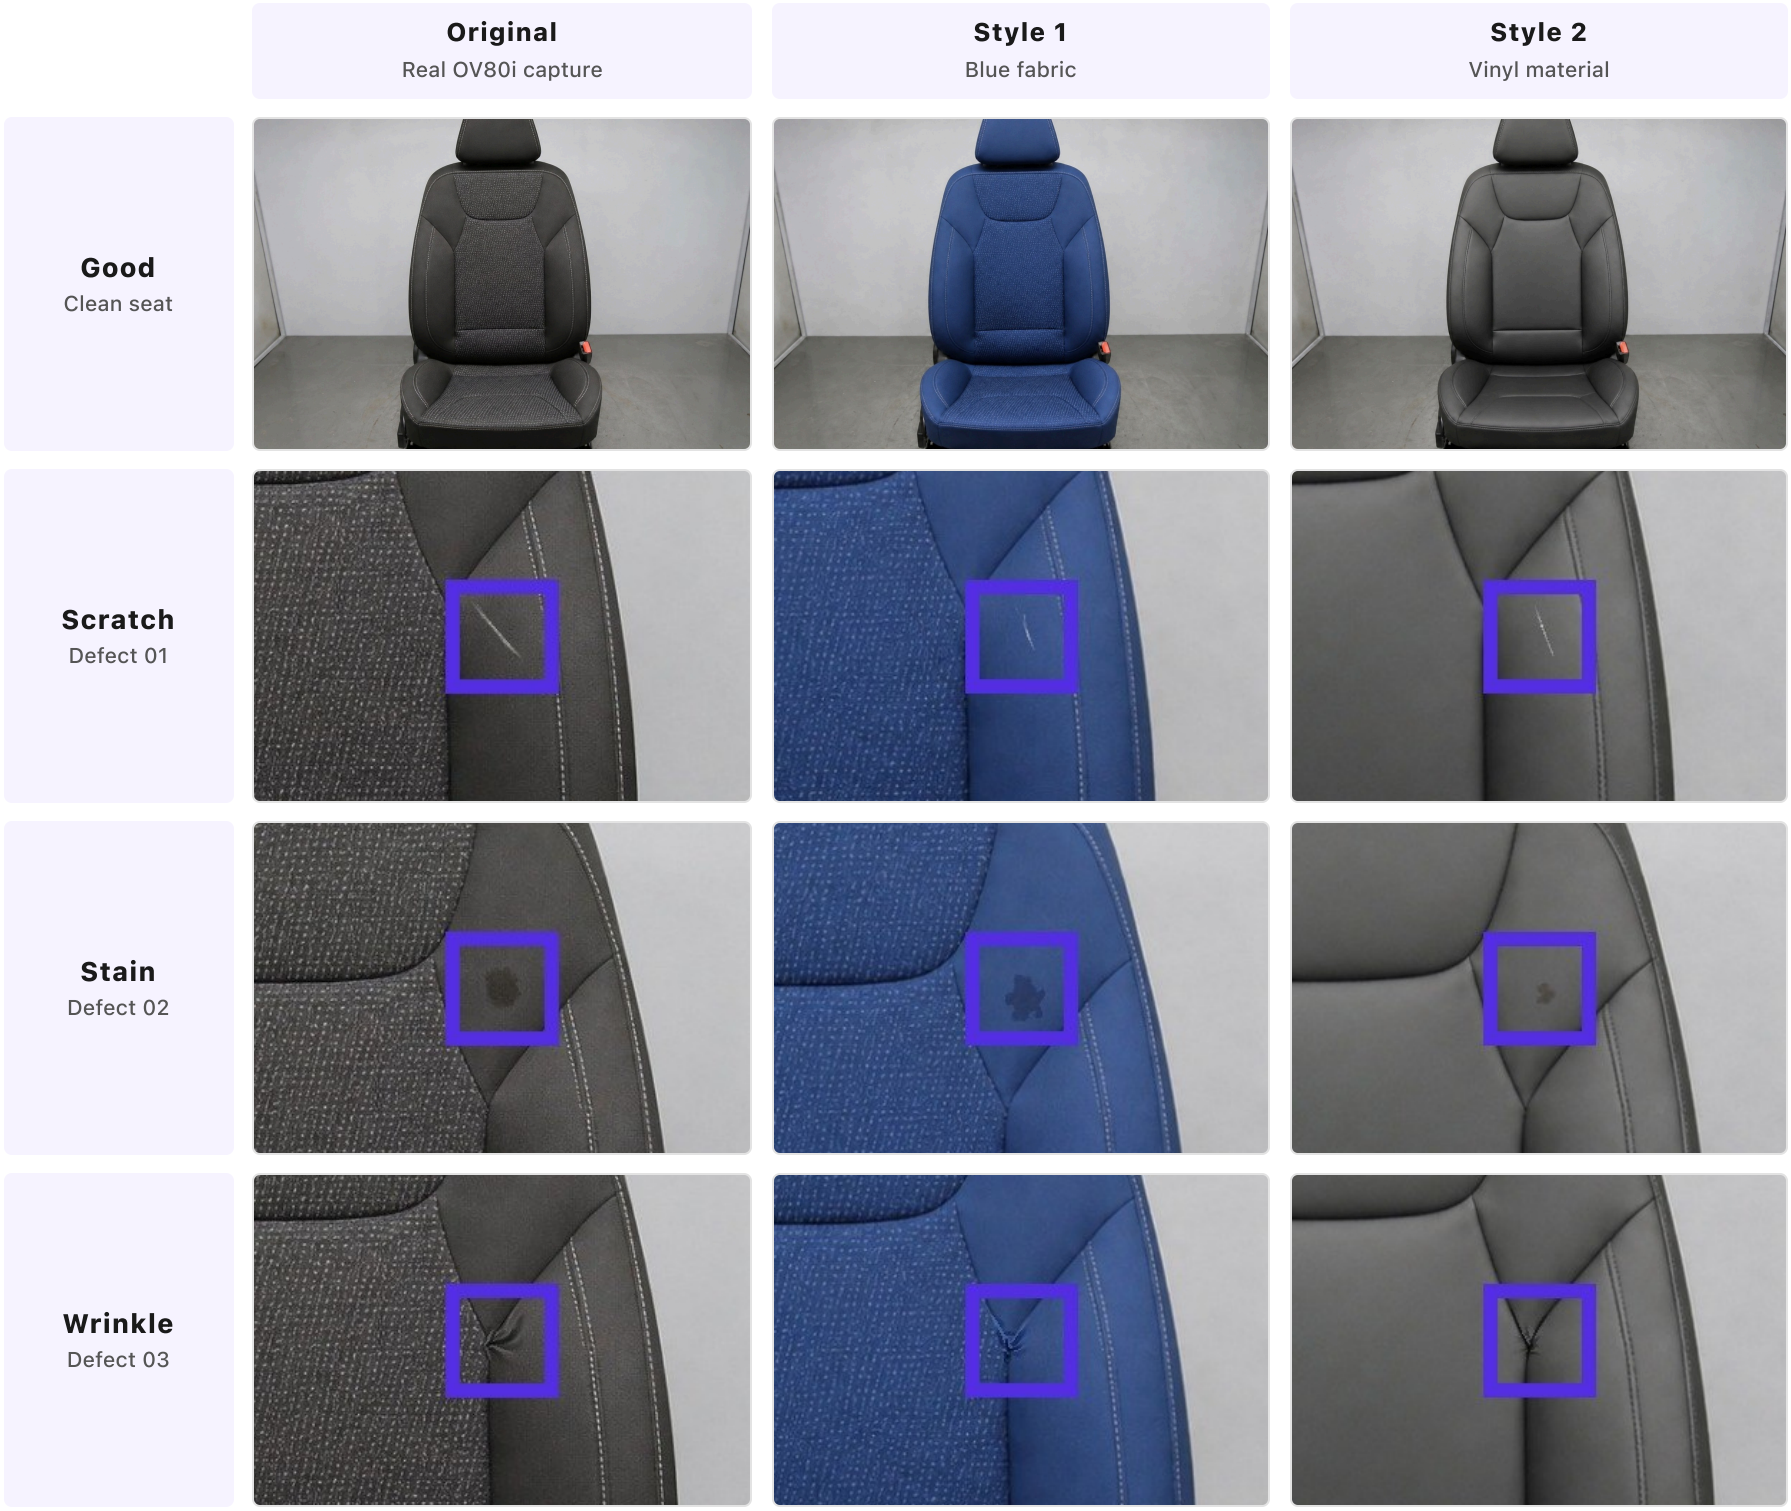

Defect region marked by the studio → Style Transfer across fabrics ↓ Defect Transfer across classes

3 fabrics × 4 defect classes = 12 labeled training images from 1 real capture.

New OEM trim spec? Style Transfer covers it before the changeover. Every defect class propagates automatically. False reject rate stays under 1% across the entire variant catalog.

What this actually saves you.

The 15-minute Contaminate Verification walk is one of many. Across the deployments in this paper, three numbers carry the program: how much faster the model gets to the line, what does not get destroyed to train it, and how much operator effort disappears.

12.4×

faster deployment

Average across Overview AI customer programs vs the classical CV path of collecting and labeling real defect samples per variant.

>$1M

in destroyed training parts avoided

Connector program alone. 300+ SKUs at roughly 200 deliberately damaged housings per SKU and ~$15-25 per housing in a classical pipeline. Wood and food lines add to this number.

12.4×

reduction in data-collection labor

Weeks of operator time per variant for image capture and hand-labeling collapse into a single shift of authoring in the OV Auto-Defect Creator Studio.

TIME TO ONBOARD A NEW VARIANT · PER APPLICATION

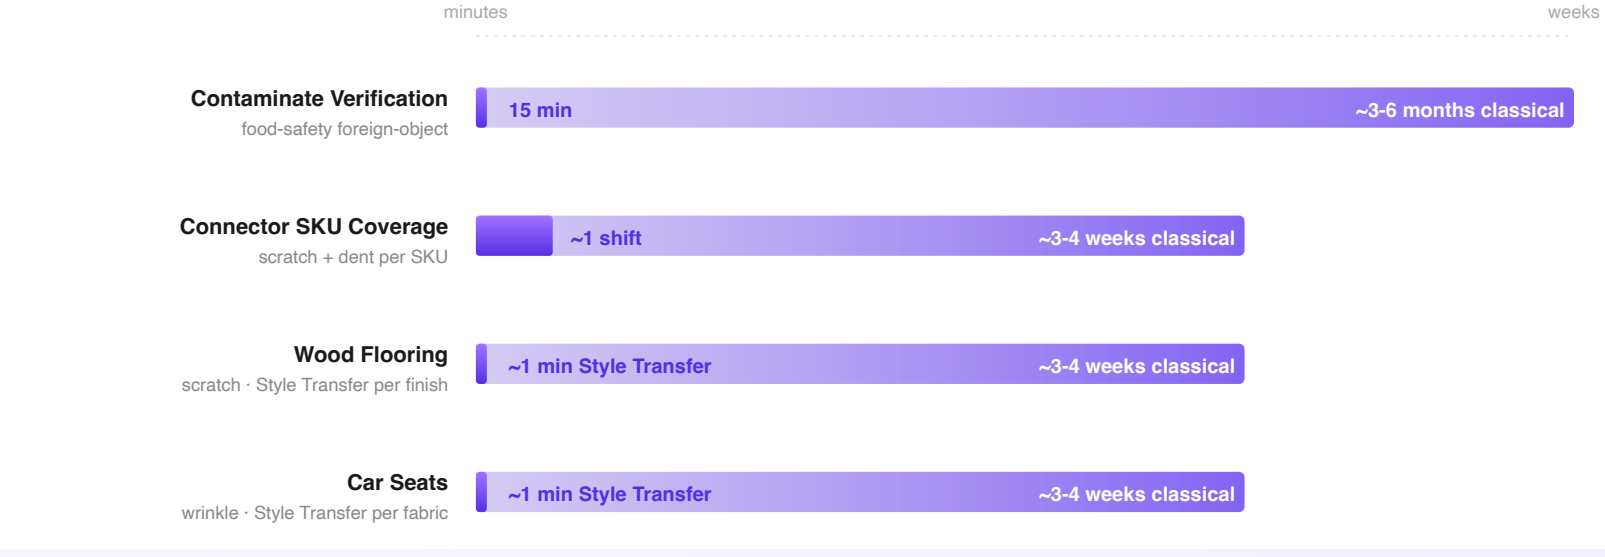

Per-variant onboarding time. Style Transfer applications push the speedup past 1000× per variant. The 12.4× headline is the weighted average across all customer programs.

See it on your line.

Reach out by email or through our website to book a demo, and see how we can solve your high-volume, medium-mix, or even high-mix vision inspection problems.

contact@overview.ai