

Computer Vision

Pixels, camera geometry, and learned models. You measure real objects from photographs, train a detector on data you labelled yourself, and run it live on cheap hardware.

Foundational

4 phases

20 sessions

30 contact hours

Syllabus

4 phases, 20 sessions. Each phase ends in something you have built.

1

Images Are Arrays

Read, clean, and measure an image as numbers before any model touches it.

5 sessions

\$72 alone

1.1 Pixels, Channels, dtypes

How an image becomes an array, why OpenCV hands you BGR, and what uint8 does at 256.

[FREE PREVIEW](#)

1.2 Loading Things That Fail

imread returning None instead of raising, headless imshow, and checking every frame you are handed.

1.3 Filters, Edges, Thresholds

Blur, gradients, adaptive thresholding, and the lighting change that makes all three collapse.

1.4 Morphology & Contours

Cleaning a binary mask, then reading area, centroid, and orientation off the shapes you kept.

1.5 A Pipeline You Can Debug

Composing the steps into one tool with tunable parameters and a visual dump of every stage.

By the end: **A counting tool that uses no trained model and shows its work at every stage.**

2

Cameras & Geometry

Turn pixels into real-world units, the part of vision no model hands you.

4 sessions

\$72 alone

2.1 How a Camera Makes an Image

Projection, focal length, exposure, and why your lens bends every straight line near the edge.

2.2 Calibration & Undistortion

Print a board, solve for intrinsics and distortion, and reject a calibration by its reprojection error.

2.3 Features, Matching, Homography

ORB keypoints, matching, RANSAC, and flattening a surface you photographed at an angle.

2.4 Measuring in Real Units

Scale from a known reference, pose from fiducial markers, and quoting your error instead of hiding it.

By the end: **A calibrated measurement tool reporting millimeters with a reprojection error.**

3

Learned Vision

Get a model finding your objects, and know which model you are allowed to ship.

6 sessions

\$72 alone

3.1 Transfer Learning

Fine-tune a pretrained backbone on a few hundred images; you will not train from scratch.

3.2 Prompt First, Train Later

Grounding DINO and SAM 3 finding objects nobody labeled, then the prompts where they quietly fail.

3.3 Detection & Segmentation Families

Boxes against masks, anchors against queries, NMS-free heads, and reading mAP without flattering yourself.

3.4 The License Is Part of the Model

AGPL and its network clause, Apache weights, non-commercial checkpoints, and picking what you can actually deploy.

3.5 Building a Dataset

Auto-label with a foundation model, correct by hand, and audit what your camera never saw.

3.6 Splits That Don't Lie

Near-duplicate video frames, perceptual hashing, and why your 0.95 mAP is fake.

By the end: **A fine-tuned detector on a dataset you auto-labeled, scored on a split that is not leaking.**

4

Live Video & The Edge

Run it on a real camera, on cheap hardware, and know when not to run it at all.

5 sessions

\$72 alone

4.1 Video, Not Images

Decoding, dropped frames, buffering, and processing at capture rate instead of falling further behind.

4.2 Tracking & Counting

IDs across frames, recovery after occlusion, and counting over a line without double-counting.

4.3 Export & Quantization

PyTorch to ONNX to INT8, choosing calibration data, and measuring the accuracy you traded away.

4.4 Hitting a Frame Budget

Profile on a Pi, a Jetson, or a laptop, and discover the bottleneck is your resize, not the model.

4.5 Failure, Bias & Consent

Domain shift, subgroup error rates, faces, and the deployment you decline to build.

By the end: **A live counting system on edge hardware with a measured frame budget and a deployment review.**

Tools you will use

OpenCV 5

NumPy

PyTorch

RF-DETR

Ultralytics YOLO26

SAM 3

Grounding DINO

ONNX Runtime

What you will build

1 Count It Without a Model

Thresholds and contours, a number you verify by hand

2 The Calibrated Tape Measure

Millimeters from a photo, with an error you state

3 Edge Counter

Prompted labels to fine-tuned model to a live frame budget

Registration

Phase 1: Images Are Arrays	\$72
Phase 2: Cameras & Geometry	\$72
Phase 3: Learned Vision	\$72
Phase 4: Live Video & The Edge	\$72
Whole track, all 4 phases	\$240

Buying every phase separately costs \$288, so the whole track saves \$48.

No payment is taken online. Registering creates an order and payment instructions are sent with it. Access opens once payment is confirmed.