

Label Foundry

From a prompt to a dataset you can review.


Independent portfolio demonstration. All inputs are newly created fictional documents or synthetic imagery. No customer assets or outcomes are shown.

Problem

Annotation work needs faster proposals while retaining control of labels, boundaries and quality.

Working application

Local SAM3 proposes six bottle masks in an original scene. The annotator reviews boundaries, corrects boxes and labels, and exports PNG masks, COCO with exact RLE or YOLO boxes with the image.



Actual app-only capture from fresh local model inference.

Measured evidence

Check	Passed / evaluated
Proposal count on the scene	1 / 1
Generated masks (structural check)	6 / 6

NVIDIA GeForce RTX 4070 Laptop GPU / 8 GB

Median / p95 complete HTTP case: 8.433 / 8.433 seconds. Includes model loading, OCR and any local queue wait.

Recorded: 2026-10-08 UTC. Exact fixture-level outputs are supplied in evaluation/results.json.

Limits

One original AI-generated scene. Six mask outputs and their exports are checked without claiming segmentation IoU. Manual boxes are identified as human annotations. The single-image archive is not a validated training dataset.

Possible pilot

An assisted annotation tool with human control, provenance and exports for your training format.

Before implementation: Taxonomy, annotation guidelines, output formats, quality review and independent training/validation splits.

Watch the English demonstration

<https://quarkeon-website.pages.dev/en/projects/label-foundry/>

Quarkeon Labs on Upwork