Intelligent Video Analytics — Project Scoping Checklist

Size any retail, industrial, robotics, or surveillance vision project before the next vendor call.

1. Define the rule — what is the one question the camera must answer?

Retail

Did they pay?
Shrink, scan-avoidance, shelf gaps, shopper heatmaps.

Industrial

Is the part defective?
Anomaly detection on good-only data. 95–99% accuracy target.

Robotics

Where am I / what is that?
Detection + pose + SLAM. Camera + LiDAR fusion.

Autonomous vehicle

Is the path safe?
Fusion or vision-only. Hard real-time, recall in bad weather.

2. The four blocks — confirm a choice for each

#	Block	What it does	Production choices
1	Detection	Box + label + score per object	YOLO family · Grounding DINO (open-vocab)
2	Tracking	Stable identity across frames	ByteTrack · DeepSORT · OC-SORT
3	Rules & analytics	Tracks become business meaning	Line crossing · dwell · scan-mismatch · anomaly
4	Deployment topology	Where blocks 1–3 run	Edge chip · central GPU · hybrid (most common)

3. Edge or cloud — tick what your deployment needs

- ☐ Lowest latency — alert in milliseconds, no network round-trip → edge.
- ☐ Privacy first — raw video must never leave the device → edge.
- ☐ Tight bandwidth — ship alerts, not full video streams → edge.
- ☐ Largest model — need a heavier or VLM model for accuracy → cloud.
- ☐ Fleet updates — push one model to every camera instantly → cloud.
- ☐ Pooled GPUs — share expensive hardware across many cameras → cloud.

4. The four traps that kill the budget

Trusting per-frame accuracy

A 95%-accurate detector over a 900-frame session fires false alerts constantly.
Fix: aggregate evidence across the whole track before alerting.

Watching only the till

~80% of stolen items are concealed in the aisle, before checkout.
Fix: aisle-to-checkout tracking, alert at the exit.

Classifying defects you can't collect

Good factories produce too few defects to train a classifier per type.
Fix: anomaly detection trained on good-only images.

Ignoring the regulation

Cameras that identify people trigger GDPR Art. 9 + EU AI Act high-risk rules.
Fix: anonymised-by-design tracking; strict retention windows.