

# Multi-Object Tracker Decision Worksheet

Four questions, four trackers, three topologies — pick the right MOT stack in ninety seconds.

## Four-question decision tree

### 1. Is the camera moving?

- Yes (phone, drone, dashcam) → BoT-SORT (CMC).
- No (fixed CCTV) → continue to question 2.

### 2. Non-linear / acrobatic motion?

- Yes (dance, gymnastics, animals) → OC-SORT.
- No (pedestrians, vehicles) → continue.

### 3. Many streams / modest hardware?

- Yes (10+ streams / server) → ByteTrack.
- No → continue to question 4.

### 4. Long-occlusion identity matters?

- Yes (sports, broadcast, POI) → BoT-SORT-ReID.
- No (counting, dwell) → ByteTrack.

## Metric cheat sheet

- MOTA — detection-biased; legacy single number.
- IDF1 — association-focused; detection-blind.
- HOTA — balanced det + assoc; the 2026 default.

## Deployment topology — default tracker

Topology	Tracker	Cost owner	Latency
Browser	ByteTrack (ONNX Web)	User device	~16 ms
Edge (many)	ByteTrack on Jetson	Hardware	~30 ms
Edge (acc.)	BoT-SORT-ReID	Hardware	~50 ms
Server	BoT-SORT-ReID on CPU	Cloud / stream	~100 ms

## Four pitfalls — pre-launch checklist

- Detector recall on real data above 90 percent.
- DeepSORT is not the default tracker.
- track\_buffer / match\_thresh tuned on real footage.
- Cross-camera IDs handled by separate Re-ID model.

## Pipeline stages

- 1 · Detection — YOLO, RT-DETR (recall > 90 percent).
- 2 · Motion prediction — Kalman or observation-centric.
- 3 · Association — IoU or IoU + appearance.
- 4 · Lifecycle — birth / death / re-id thresholds.