Anomaly Detection in Surveillance Video - One-Page Reference

What it is, how it differs from rule-based analytics, the model families, why false alarms are the real metric, where it runs, the ONVIF catch, and the legal line. Representative 2025-2026 figures - all move with scene, lighting, angle, and tuning. Never 100%.

A. What it is - and what it is not

- Anomaly detection learns what is NORMAL for a camera view and flags whatever deviates - no rule authored. It answers 'tell me when something I did not anticipate happens', the opposite of a rule (line, zone, timer), where you specify the event in advance.
- It is not a separate camera or a 'detect everything' button: it runs on object detection + tracking, scoring the scene description. Anomalous is NOT the same as dangerous - it flags the statistically unusual (a wheelchair, an odd outfit), which is its value and its false-alarm problem.

B. The model families (internals: AI for Video Engineering)

- Reconstruction / prediction (autoencoders): trained on normal video only; high rebuild or next-frame error = anomaly. Unsupervised, no anomaly labels. One-class / feature-distance: model the normal distribution, outliers score high.
- Weakly-supervised (UCF-Crime / deep MIL): video-level labels only; better on messy real footage. 2025-2026 vision-language methods (LAVAD, VERA): training-free and explainable - return a sentence saying WHY. The model internals are owned by the AI section; this is the application.

C. Accuracy is a range - and false alarms are the real metric

- Benchmark frame-level AUC: ~97% UCSD Ped2, ~96% CUHK Avenue, ~98% ShanghaiTech (controlled), but ~81-88% on real-world UCF-Crime and ~86-94% AP on XD-Violence. A live camera is harder - 'normal' drifts with day/night, weather, crowds. AUC is a ranking quality, NOT a false-alarm rate.
- The base-rate math: one camera, ~86,400 scored moments/day, ~10 truly unusual. At a 1% false-positive rate you get ~864 false alarms vs ~8 true catches - precision ~0.9%. Even at 0.1% it is ~86 false vs 8 true. Rare events swamp the list; the floor is never zero.
- So deploy it as a TRIAGE layer that ranks footage for a human or a forensic search - never an autonomous alarm. Alert fatigue (operators ignoring or disabling noisy analytics) is the true failure mode. And it scores EVERY frame continuously - a real, ongoing compute cost - so plan for drift and retraining.

D. Where it runs + how it surfaces - the ONVIF catch

- On the camera (edge): under 0.1 s, only a score/clip leaves, video stays local. On a local server: under 1 s, heavier model, on-prem. In the cloud: 0.5-12 s, every frame up, continuous egress + compute bill. Continuous scoring makes the edge or a local server the usual default.
- ONVIF standardizes named rule types (Line, Field, Loitering, counting) - but 'anomaly' is NOT one of them. A vendor's anomaly analytic travels over ONVIF Profile M as vendor-defined metadata: portable plumbing, vendor-defined meaning. Expect more SDK lock-in; conformance is interoperability, not accuracy.

E. The legal line: personal data, usually not biometric

- Scoring identifiable people is PERSONAL DATA under GDPR (Art. 4(1)) - lawful basis, notice, and a DPIA for systematic public monitoring (Art. 35; EDPB Guidelines 3/2019). But scoring how a scene DEVIATES, not WHO is in it, is generally NOT Art. 9 biometric data.
- It tips into the heavier regime the moment it IDENTIFIES the unusual person (face watchlist) or infers emotion/intent (EU AI Act Art. 5 restricts emotion recognition and real-time remote biometric ID; prohibitions since Feb 2025, high-risk biometric duties from 2 Dec 2027). Keep it anonymous by design: cue a human, never judge a person.

Accuracy is a range tied to scene, lighting, angle, and tuning, never a single number and never 100%. The benchmark AUC/AP figures are 2025-2026 academic numbers (representative, condition-dependent); the false-alarm arithmetic (86,400 moments, ~10 anomalies, 1%/0.1% FPR) is illustrative base-rate math, not a measured constant. The detection/anomaly MODELS are engineered in the AI for Video Engineering section; this is the surveillance APPLICATION - the false-alarm reality, the deployment, the standard, and the law. Sources: ONVIF Analytics Service Specification + Profile M; GDPR Reg. (EU) 2016/679 Art. 4(1), 35; EDPB Guidelines 3/2019; EU AI Act Reg. (EU) 2024/1689 Art. 5; Sultani et al. (UCF-Crime, CVPR 2018); 2025 benchmark surveys (IET); edge-anomaly studies (MDPI Sensors).