Roadwatch: Leveraging AI and Configurable Rules for Real-Time Traffic Anomaly Detection on Existing Camera Infrastructures

Roberto Saia, **Alessandro S. Podda, Livio Pompianu, Mirko Marras, Nicola Floris, Marco M. Manca, and Gianni Fenu**

Department of Mathematics and Computer Science
University of Cagliari, Italy

ICCSIT 2026, 19th International Conference on Computer Science and Information Technology
25÷27 September 2026 – Paris, France

Outline

- Context
- System Architecture
- Validation Methodology
- Results
- Conclusion

Context

The **growing density** of vehicles and pedestrians, together with the **increasing complexity** of modern road infrastructures limits the effectiveness of traditional human-centered monitoring approaches, which **hardly scale** to large urban areas

Context

Existing traffic monitoring systems face **open challenges**:

- Many **require dedicated hardware**, such as carefully positioned sensors or specific camera models
- Most rely on **low-level detections** rather than semantic traffic events

Roadwatch addresses both by combining **deep learning** with configurable, **rule-based** event interpretation

System Architecture

Roadwatch is built on **NVIDIA DeepStream** and **GStreamer**, organized as a real-time **modular pipeline** structured as follows:

- the **video preprocessing** stage: working-area definition, perspective correction, anonymization
- **AI** perception combines **detection** (**YOLOv4-CSP)** with **tracking** (**NvDCF)**
- a **microservice** layer applies **rule-based interpretation**

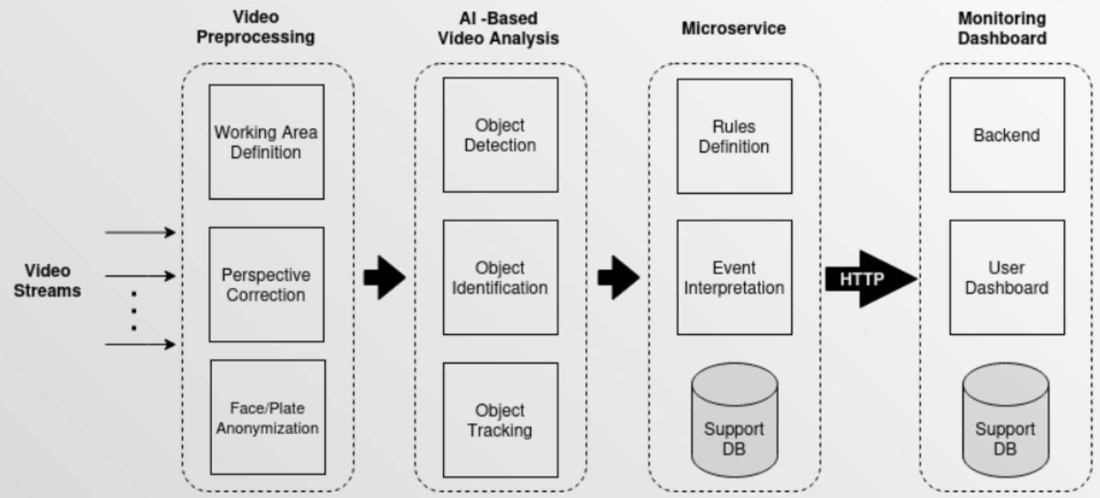

Validation Methodology

- Roadwatch adopts an **event-based validation** paradigm: **events**, not frames, are the unit of evaluation
- each event is **characterized**, then reviewed to confirm whether it is a true or **false positive**:

$$TP = \{ e \in E \mid e \text{ is semantically and spatially correct} \}$$

$$FP = \{ e \in E \mid e \text{ is semantically or spatially incorrect} \}$$

Validation Methodology

- False negatives are estimated via a **temporal sampling strategy** applied to a subset of the **observation window**
- The **sample recall** is used to **extrapolate** the **estimated false negatives** over the full **24h window**:

$$Recall_{sample} = \frac{TP_S}{TP_S + FN_S}$$

$$N_{real} \approx \frac{N_E}{Recall_{sample}}$$

$$FN_{est} = N_{real} - N_E$$

Results

- The experimental evaluation was carried out over a continuous 24-hour **observation window**, recorded by the Municipality of Cagliari:
- the system automatically detected **899 anomalous events**, fully logged by the system
- a **temporal sample of 82 real events** was reviewed: 74 correct, 8 missed
- Manual review confirmed **zero false positives**

Indicator	Value
Events detected by the system (TP)	**899**
Real events in the sample	**82**
Correctly detected / missed in sample	**74 / 8**
Estimated recall	**0.902**
Estimated total real events	**≈997**
Estimated total false negatives	**≈98**

Results

We analyzed the ***distribution of detected events*** by class, with false negatives **proportionally projected** across event types; manual review confirmed **zero false positives** across every category

Event	TP	FP	FN (est.)	TN *	Share (%)
ILLEGAL_WAY_CROSSING	616	0	67	n/a	68.52
ENTITY_ON_ROAD	273	0	30	n/a	30.37
AREA_INVASION	4	0	0.4	n/a	0.44
NO_PARK	5	0	0.5	n/a	0.56
RECKLESS_DRIVING	1	0	0.1	n/a	0.11
Total	**899**	**0**	**98**	n/a	**100.00**

* **Note**: True Negatives are not defined in event-based evaluation, as there is no finite set of "non-events" to count.

Conclusions

- Roadwatch combines **deep learning** perception with **rule-based** reasoning, **operating on existing camera** infrastructures without additional sensors

- An **event-based validation paradigm** with temporal sampling enables recall estimation on continuous real-world footage, where frame-level annotation is impractical

- Over 24h of real urban traffic, the system detected **899 events** with **zero false positives** and an estimated **recall of 0.902**

Roadwatch: Leveraging AI and Configurable Rules for Real-Time Traffic Anomaly Detection on Existing Camera Infrastructures

THANK YOU FOR YOUR ATTENTION