



INTERNATIONAL JOURNAL OF SCIENTIFIC RESEARCH AND INNOVATIVE STUDIES

Advancing Knowledge, Inspiring Innovation

www.ijsrisjournal.com



 **PEER REVIEWED**
Rigorous & Transparent

 **DOI: CROSSREF**
Member

 **OPEN ACCESS**
For Global Visibility

 **BI-MONTHLY**
Timely Publication

 **WIDE READERSHIP**
High Impact & Visibility

 **AUTHOR FRIENDLY**
Fast | Supportive | Reliable



RECEIVED
17 August 2026



ACCEPTED
09 September 2026



PUBLISHED
1 October 2026



October 2026



Volume 5



Number 5

319 to 331

Integration of an intelligent humanoid system into road traffic management at the Roads Agency: An approach based on DevOps knowledge-based systems and Big Data technologies

Blaise KAPALALA KAPENDA¹, Merveille BONTE KASSONGO², Thierva KANIKI MADILA³

¹Computer Science, Institut Supérieur Pédagogique de la Gombe, City-Province of Kinshasa, Democratic Republic of the Congo, blaisekapalala94@gmail.com , (+243) 811 459 397, <https://orcid.org/0009-0000-0815-6956>

²Science and Technology, Department: Mathematics, Statistics and Computer Science, National Pedagogical University, City-Province of Kinshasa, Democratic Republic of the Congo, merkasbonte@gmail.com , (+243) 812 180 450, <https://orcid.org/0009-0000-1668-2710>

³Computer Science, Gombe Higher Institute of Education, City and Province of Kinshasa, Democratic Republic of the Congo, (+243) 823532619

LICENSE



This article is licensed under the Creative Commons Attribution 4.0 International License (CC BY 4.0).

<https://creativecommons.org/licenses/by/4.0/>



PREFIX :

10.63883

IJSRIS JOURNAL

OPEN ACCESS



This is an open access article distributed under the terms of open access.

ABSTRACT

Road traffic management requires information that is rapid, reliable and actionable for operational managers. This article proposes an architecture to support road surveillance, combining a humanoid device, a knowledge base, artificial intelligence models, a Big Data platform and a DevOps/MLOps pipeline. The approach falls within the scope of Design Science Research and comprises a targeted literature review, the formulation of requirements, architectural design and the definition of an evaluation protocol. The proposed system integrates field perception, edge processing, event ingestion, hybrid inference and human supervision. The numerical evaluation is limited to illustrative values and a fictitious confusion matrix comprising 1,000 events. Recalculation of this matrix yields an accuracy of 91.00 per cent, a precision of 91.11 per cent, a recall of 89.13 per cent and an F1 score of 90.11 per cent. These values demonstrate the calculation of the indicators, without establishing the operational effectiveness of the system. The analysis highlights the benefits of gradual integration and the need to compare the humanoid's specific contribution with that of a fixed installation. The main contribution is a traceable design framework for road decision support. Its validation requires local data, reproducible trials, a controlled comparison and clarification of institutional responsibilities.

Keywords: Intelligent humanoid system; road traffic; knowledge-based system; artificial intelligence; Big Data; DevOps; MLOps; human supervision; Office des Routes.

1. Introduction

1.1. Background and rationale

The monitoring of road traffic flows is a key issue for mobility management and safety. The World Health Organisation's global report describes the scale of road risks and highlights the importance of prevention and monitoring measures. It provides an international context, without serving as a measure of the performance of the network studied here (World Health Organisation [WHO], 2023).

Digital transformation enables sensors, information systems and monitoring centres to be integrated. Cameras, GPS tracks and traffic measurements can feed into traffic flow analyses and alerts. Computer vision provides methods for recognising and estimating movement, but its reliability depends on the images and the conditions under which they are captured (Szeliski, 2022).

This study focuses on an assistance system intended for use by the Roads Authority in the Democratic Republic of the

Congo. It does not report on either an institutional roll-out or a survey of its departments. Responsibilities for monitoring, traffic policing and intervention will need to be clarified with the relevant authorities prior to a pilot scheme. The humanoid is envisaged as a cyber-physical node capable of perceiving, communicating and presenting structured information to operators.

1.2. Problem statement and research question

Traditional road management systems rely heavily on human observation, administrative reports, telephone calls, periodic inspections and information systems that are sometimes independent of one another. These methods can lead to delays in data collection, scattered data and limited predictive use of historical data.

The scientific problem is formulated as follows: How can an integrated architecture be designed to enable an intelligent humanoid system, combined with a knowledge-based system, Big Data technologies and a DevOps/MLOps approach, to contribute to the monitoring, analysis and management of road traffic at the Roads Agency?

Secondary issues relate to environmental perception, knowledge representation, big data processing, interoperability, system availability, performance indicators and human supervision.

1.3. Objectives and contribution

The overall objective is to design an integrated architecture combining humanoid robotics, knowledge-based reasoning, data processing and continuous operation. The specific objectives are to define monitoring requirements, organise multimodal traffic flows, formalise business rules, propose a DevOps/MLOps pipeline, identify safety constraints and establish a comparative evaluation protocol.

The expected contribution is a model for integration and governance. Simply combining existing technologies is not sufficient to demonstrate empirical originality: the scientific value must be established through comparisons that measure the contribution of each component and its suitability for real-world applications.

1.4. Research hypotheses

H1. The integration of an intelligent humanoid system with a road data processing platform could reduce the time between the occurrence of an incident and its detection.

H2. The use of a knowledge-based system combined with artificial intelligence models would improve the ability to classify and interpret road situations.

H3. The use of Big Data technologies would enable the effective integration and analysis of large volumes of heterogeneous road data.

H4. Adopting a DevOps/MLOps architecture would improve the maintainability, traceability, reproducibility and availability of services.

H5. A hybrid architecture combining artificial intelligence and human validation would be more suitable for a critical environment than fully autonomous automation.

1.5. Positioning within the literature

Traffic forecasting relies on temporal and spatial relationships. Graph neural networks enable the representation of links between road segments or junctions. Jiang et al. (2023) summarise these approaches and the challenges involved in applying them to forecasting. Their relevance to the present system forms the basis for a model selection hypothesis, to be tested against local data.

Symbolic AI enables the representation of explicit rules and the generation of fact-based explanations. Statistical AI estimates properties based on examples. Their combination forms the basis of a design proposal here: a detection score is contextualised by a business rule prior to an alert (Russell & Norvig, 2021).

Kreuzberger et al. (2023) describe MLOps as the organisation of the model lifecycle, combining automation, operations and monitoring. This perspective complements the delivery and reliability practices examined by DORA (2024). On its own, it does not demonstrate an improvement in the availability of the future road system.

2. Methodology

2.1. Type of research and status of the approach

The research adopts a Design Science Research approach centred on the construction of an artefact intended to address a practical problem. The framework proposed by Peffers et al. (2007) distinguishes between problem identification, objectives, design, demonstration, evaluation and communication. In this article, the design and conceptual demonstration are presented; the empirical evaluation remains to be carried out.

The structure follows the IMARaD framework: the introduction sets out the problem; the methodology describes the approach; the analyses examine the artefact and the indicators; the discussion interprets its implications; the results present the deliverables and calculations. The interpretative conclusions refer to the detailed data in the Results section.



Fig 1: Methodology and approach based on the IMRaD framework.

2.2. Literature review and knowledge acquisition

The literature review is a targeted narrative review. It draws on scientific publications on traffic forecasting and MLOps, works on vision and data architectures, as well as institutional risk management frameworks. The recent references prioritised cover the period 2021 to 2024; a few seminal works are included for the method and concepts. This is not an exhaustive systematic review, and no claim is made regarding the number of items selected without documentation.

Future acquisition of domain knowledge will be based on interviews with operators and road experts, consultation of procedures, and validation of rules by the relevant departments. These interviews are proposed and are not presented as having already taken place. Each rule must include an author, a justification, a date, a version and conditions of application.



Fig 2: Literature review methodology.

2.3. Units of analysis and required data

The detection unit will be a time-stamped event located on a road segment. The congestion assessment unit will be a time window at a junction. The required data include video sequences, speeds, counts, GPS positions, sensor statuses and human annotations. No empirical dataset specific to the Office des Routes is available for this study.

A future dataset will need to cover day and night, adverse weather conditions, varying levels of traffic density, occlusions and network outages. Events will be annotated according to a common guideline, with any disagreements subsequently resolved. The training, validation and test datasets must be separated by time period or location to minimise information leakage between closely related sequences.

2.4. Artifact design and planned comparison

The design follows these steps: formulation of requirements, modelling of entities, definition of interfaces, formalisation of rules, selection of models and processing methods, definition of operational controls, followed by preparation of tests. The technological choices constitute implementation options; no operational prototype has been verified at this stage.



Fig 3: Artifact design and planned comparison.

Four scenarios have been selected: free-flowing traffic, heavy traffic, severe congestion and accident/obstacle. Two configurations are compared: the conventional system and the proposed intelligent architecture.

The conventional reference system must be described in detail: procedures, equipment, staffing levels, timeframes and observation frequency. Two additional variants are required: a platform with fixed sensors and an identical platform equipped with the humanoid. Tests with and without a rule engine will also enable the contribution of symbolic inference to be isolated. The configurations must be subjected to the same events or comparable traffic conditions.



Fig 4: Traffic scenarios and comparative configuration.

2.5. Metrics and analysis plan

Latencies will be measured using synchronised logs. The detection time corresponds to the duration between the annotated start of the event and its recognition; the transmission time spans the sending and receipt of the alert; the overall response time covers the event through to a defined action. Stages that may overlap must not be added together without checking their timestamps.

For binary detection, TP refers to correctly detected incidents, TN to non-incidents correctly excluded, FP to false alarms and FN to missed incidents. The formulas are:

- Accuracy = $(TP + TN) / (TP + TN + FP + FN)$
- Precision = $TP / (TP + FP)$
- Recall = $TP / (TP + FN)$
- F1 = $2 \times TP / (2 \times TP + FP + FN)$
- False positive rate = $FP / (FP + TN)$

Accuracy and recall measure different properties. Recall should be prioritised to track missed incidents, whilst monitoring the burden of false alarms. Availability should be defined as a proportion of the expected uptime; data availability and robot availability should be measured separately.




Fig 5: Metrics and analysis plan.

For future tests, means will be supplemented by medians, quantiles and confidence intervals. Proportions must specify their denominators. A paired comparison will be preferred when the same events are processed by several configurations. The choice of tests will depend on the distribution and dependence of the measurements; no statistical significance can be calculated from the illustrative values presented here alone.

2.6. Reproducibility and risk management

A publishable simulation must specify the simulator, its version, the network, the injected traffic, the incident distributions, the random seeds, the durations, the number of repetitions and the computing resources. The data, configurations, models, rules and logs must be versioned. The figures in this article do not have such associated artefacts and are therefore treated as fictitious numerical examples.

AI risk management must define responsibilities, characterise the context, measure risks and plan for their mitigation. The AI RMF framework provides a voluntary reference for this approach (National Institute of Standards and Technology [NIST], 2023). The present architecture provides for human validation of critical alerts and a fallback procedure.



Fig 6: Simulation reproducibility and AI risk management.

3. Analyses

3.1. Functional requirements and the role of the humanoid

A humanoid robot is a robotic system whose physical or functional characteristics are modelled on the human body and human interactive capabilities. In this research, it refers to a

device with capabilities for perception, analysis, reasoning, interaction, communication and decision support.

Visual sensors can be used to detect vehicles, people and obstacles, and to estimate density. Audio sensors facilitate voice interaction. LiDAR or depth sensors provide three-dimensional perception, whilst GPS/GNSS provides location data.

The specific role of the humanoid relates to local mobility, the presentation of information and interaction with users. Detection can also be carried out by fixed equipment. The analysis must therefore distinguish between the utility of the digital platform and the added value of the humanoid form, which entails additional costs and mechanical constraints.



Fig 7: Requirements and role of a humanoid.

3.2. Architecture and data flow

The architecture comprises eight layers: road environment; perception; communication; Big Data; intelligence; knowledge; services; human supervision. The humanoid acts as a cyber-physical node connected to the whole system.

Capture → Pre-processing → AI → Classification → Knowledge base → Inference → Risk level → Recommendation → Validation → Action → Logging.




Fig 8: Eight-layer architecture and data flow.

An example of a congestion scenario might combine the number of vehicles, speed and queue length, then generate a structured alert including location, priority and recommendation.

The eight layers describe logical functions and not necessarily eight distinct applications. Perception generates observations; communication transmits them; the Big Data platform stores them; models estimate states; rules interpret them; services structure the alerts; supervision reviews the recommendations. The humanoid device interacts with several of these functions.

3.3. Big Data Processing and Service Continuity

Road data illustrates the dimensions of Big Data: volume, velocity, variety, veracity and value. It may include images, videos, GPS coordinates, IoT data, events, text and historical records.

The architecture follows the chain: Sources → Ingestion → Streaming → Storage → Processing → AI → Decision → Visualisation. Sources include the humanoid robot, cameras, road sensors, GPS and applications. Apache Kafka can serve as an event bus; Spark or Flink can handle aggregations and processing; PostgreSQL, MongoDB and a data lake can be combined for storage.

Critical processing must be carried out close to the source. The Edge can handle object detection, filtering, compression and urgent alerts. The Cloud or Data Centre handles historical analysis, machine learning, consolidation, dashboards and archiving.



Fig 9: Big data processing and service continuity.

Event durability, duplicate handling and concurrent write management are key challenges in a distributed architecture (Kleppmann, 2017). In the proposed system, each event must include an identifier, a timestamp, a source, a location and a quality indicator. A local buffer will enable observations to be retransmitted following an outage, whilst checking for duplicates.

3.4. Perception and predictive models

CNNs are well-suited to object detection and image classification. LSTMs can be used for time series analysis and traffic flow forecasting. Transformers enable the analysis of complex sequences. Graph Neural Networks can model the road network as a graph $G=(V,E)$, where V represents junctions and E represents road segments.

Visual detection, traffic flow forecasting and incident classification must be evaluated separately. A good vehicle detector does not necessarily demonstrate the ability to recognise an accident. Perception tests should include challenging scenarios and errors by class; comparisons of temporal models will use an identical forecast horizon and a single benchmark (Szeliski, 2022).




Fig 10: Predictive models and perception.

3.5. Knowledge base and hybrid inference

A knowledge-based system uses formalised knowledge to produce conclusions or recommendations. Its general architecture comprises a knowledge base, a fact base, an inference engine, an interface and an explanation module. Example rule: IF density = high AND average_speed < 10 km/h AND queue_length > threshold THEN traffic_status = severe_congestion.

The hybrid approach combines statistical AI and symbolic AI: a model can estimate the probability of an accident, and the rule engine can contextualise this estimate before generating an alert. The main entities are Road, Segment, Sensor, Humanoid, Vehicle, Incident and Alert. Each entity contains operational attributes necessary for monitoring and decision-making.

Thresholds for speed, traffic density and queue length will be calibrated locally. The rule cited above is an example of formalisation, not a universal standard for congestion. In the event of a discrepancy between the model's score and the rule, the alert must flag this discrepancy to the operator rather than masking the uncertainty. The explanation module will indicate the facts used, the rules triggered and the corresponding versions.



Fig 11: Knowledge base and hybrid inference.

3.6. Use cases and alert management

The stakeholders are the system administrator, the centre operator, the traffic expert, the data scientist, the DevOps/MLOps engineer, the humanoid and the user. Responsibilities are separated to ensure system governance.

The robot monitors the road and transmits images; the model detects an abnormal situation and calculates a confidence score; the inference engine consults the rules; the event is classified; an alert is transmitted; the operator verifies it; and the event is logged.

The system receives sensor feeds, pre-processes the data, runs the detection models, constructs the facts, applies the inference engine and triggers an alert or notification depending on the risk level. The events are then forwarded to the Big Data platform. The dashboard can display real-time traffic, incidents, the location and status of robots, AI model metrics and IT infrastructure indicators.



Fig 12: Use cases and alert management.

3.7 DevOps and MLOps pipeline

DevOps brings software development and IT operations closer together. The intended pipeline is:

Code → Build → Test → Security → Deployment → Monitoring → Feedback.

It improves test automation, version management, reproducible deployment, microservices monitoring and logging.

MLOps extends DevOps to the model lifecycle:

Data → Training → Validation → Versioning → Deployment → Monitoring → Retraining.

Monitoring data drift and performance enables a re-evaluation to be triggered when performance falls below a minimum threshold. Promotion to production must depend on testing, data validation, comparison with the previous model and the option to roll back. MLOps monitoring will distinguish between input drift, performance degradation measured against annotations, and operational incidents; drift will not automatically necessitate retraining (Kreuzberger et al., 2023).

Operational metrics will include deployment failures, recovery time and delivery frequency. Monitoring these will enable the assessment of H4. DORA research provides a framework for analysing delivery practices, but its general findings will need to be considered within the context of the platform (DORA, 2024).



Fig 13: DevOps vs MLOps.

3.8. Cybersecurity and data governance

The security of such a device goes beyond that of an ordinary information system. The humanoid is a cyber-physical node: it perceives its surroundings, moves about and interacts with users in public spaces. This has two consequences. A breach of the system not only threatens data, but can also distort a traffic management decision, with very real consequences. Furthermore, e of the threats do not target the software, but the sensors themselves, prior to any data processing. These two distinctive features – the convergence of safety and security, on the one hand, and the vulnerability of perception, on the other – guide the measures presented here.

Rather than simply listing a series of controls, it is preferable to categorise the threats. A STRIDE-type framework applied to the eight-layer architecture reveals risks that a mere list of best practices fails to highlight: the compromise of a sensor that would inject false readings; the tampering with an event stream between the periphery and the platform; the repudiation of an alert due to a lack of reliable logging; the disclosure of images containing personal data; a denial-of-service attack saturating the ingestion bus; or the elevation of privileges in human supervision. Each of these risks calls for a countermeasure designed in from the outset, not a patch applied after the event.

Modèle de menace STRIDE appliqué à l'architecture cyber-physique

Catégorie STRIDE	Mesure concrète sur le système	Contrôle(s) concerné(s)	Contre-mesure de conception
Interception (Eavesdropping)	Capture du GPS, écoute passive de données véhiculées	Intégrité, Confidentialité	Sécurité cryptée des données, masquage des données
Altération (Tampering)	Faux événements envoyés, manipulation des données d'événement	Confidentialité, Intégrité	Catégorisation et vérification des données, contrôle d'intégrité de l'envoi
Repudiation	Origine d'un événement ou d'un alerte contestée, fausses données	Sécurité, Intégrité	Journalisation horodatée, traçabilité des règles et actions
Disclosure (Data disclosure)	Fuite d'images, données personnelles et images	Intégrité, Confidentialité	Pseudonymisation / Masquage, destruction contrôlée des données
Denial of service (DoS)	Saturation du bus d'événements, saturation des ressources	Confidentialité, Disponibilité	L'ajout de ressources, gestion des ressources
Elevation of privilege	Perte de contrôle de la supervision de l'opérateur	Intégrité, Confidentialité	Processus d'urgence, détection des anomalies, validation humaine

Fig 14: STRIDE threat model applied to the eight-layer cyber-physical architecture.

The most specific issue concerns perception. Vision models tasked with detecting vehicles and incidents are susceptible to adversarial examples: an imperceptible disturbance, or a simple sticker on a sign, may be enough to cause a misclassification (Goodfellow et al., 2015). Location sensors are not immune either: GPS can be spoofed to shift the apparent position of an event, and LiDAR can be deceived by reflected signals (Petit & Shladover, 2015). A system that triggers alerts based on such inputs must therefore operate on the assumption that its sensors may be unreliable. In practice,

this means cross-referencing multiple sources before issuing any critical alert, flagging inconsistencies between sensors rather than smoothing them over, and treating an isolated detection as a hypothesis to be confirmed.

This scepticism towards inputs ties in with the issue of safety. Because the system guides traffic decisions, a successful attack does not merely cause a traffic jam: it can trigger a false congestion alert or, more seriously, mask a real incident. Security and safety cannot therefore be treated separately. Human validation of critical alerts — already provided for in the architecture — plays a dual role here as a safeguard, both against model error and against its manipulation, provided that the operator also receives the evidence justifying the alert, and not just its conclusion.

Downstream of the sensors, the processing chain presents its own points of vulnerability. Peripheral devices, the event bus and distributed processing must each have their own unique identity, encrypted communications and secret management separate from the application configuration; access to any stolen or removed equipment must be revoked without interrupting the service. The training pipeline warrants the same level of vigilance: corrupted training data can permanently bias a model, and a deployed model can be tampered with or copied. This is why putting a model into production should be based on an integrity check and a comparison with the previous version, and why MLOps monitoring must distinguish between natural data drift and degradation caused by malicious activity: the two appear similar in the metrics, but require opposite responses.

Then there is the matter of data governance, which is inextricably linked to security in this context. Cameras installed on public roads capture number plates and faces: this constitutes personal data, and its processing is regulated in the Democratic Republic of the Congo by the Digital Code (Ordinance-Law No. 23/010 of 13 March 2023). The principles of purpose limitation, data minimisation and limited retention are not mere formalities; they determine what the system is permitted to collect, retain and reuse. In practical terms, the images should be analysed at the level of the event rather than the individual; identifying details should be pseudonymised or blurred as soon as the purpose permits; and the whole process should be subject to an impact assessment prior to any pilot scheme. Actual deployment will also require the clear designation of the data controller – an issue made all the more sensitive by the fact that several organisations are involved in road management.

These measures do not amount to a mere accumulation of tools but rather constitute a systematic approach. The CSF 2.0

framework (Pascoe et al., 2024), already cited, provides the backbone: govern, identify, protect, detect, respond, recover; and the AI RMF framework (NIST, 2023) extends this to risks specific to models. Neither of these constitutes certification; they offer a structured way of deciding what to protect, against what, and in what order of priority. For a cyber-physical system operating in the public domain, this discipline determines both user confidence and the validity of the decisions it supports.

4. Discussion

4.1. Scientific scope and links to existing work

The design ties in with work on spatial traffic forecasting and the continuous operation of models. However, these references do not validate the proposed complete architecture. Graphs can represent interactions between segments, whilst MLOps organises updates; their combination with a humanoid must be the subject of specific testing (Jiang et al., 2023; Kreuzberger et al., 2023).

The expected benefit of hybrid inference is improved traceability of the recommendation. An explicit rule may be understood by operators, but it may also be incomplete or inappropriate. Explaining a decision therefore does not guarantee its validity. The comparison will need to simultaneously measure accuracy, recall, critical errors and the usefulness of the explanations.

4.2. Interpretation of numerical examples

The values detailed in the Results section illustrate possible discrepancies between configurations. They do not stem from a documented experiment and do not demonstrate any actual reduction in response times. The fictitious matrix allows the calculation of the indicators to be verified and highlights the necessary distinction between overall accuracy, recall and false positive rate.

The superiority of the humanoid cannot be inferred from the illustrative differences: the proposed system simultaneously modifies the sensors, software, inference and information exchange. A comparison with an equivalent stationary platform is essential to attribute any improvement to the humanoid's mobility or interaction.

4.3. Status of assumptions

H1 remains untested empirically; the illustrative latency figures do not constitute validation results. H2 requires a comparison between standalone models and models combined with rules. H3 requires measurements of throughput, latency,

data loss and costs under load. H4 requires observations across multiple versions and operational incidents. H5 is retained as a design choice to be evaluated with operators, and not as statistically proven superiority.

Hypothesis H5, which posits that a hybrid architecture combining artificial intelligence and human validation is better suited to a critical environment than full automation, finds further support in the security analysis (section 3.8). Given that the sensors themselves can be deceived and that the system directs decisions with physical consequences, human supervision ceases to be a mere functional safeguard: it becomes a line of defence against both error and manipulation. This convergence between safety requirements and security requirements reinforces the relevance of an assistance mechanism, rather than one of autonomous decision-making.

4.4. Limitations and conditions for transfer

The study is based primarily on a conceptual architecture and experimental scenarios. The quantified results are simulated. Actual effectiveness will depend on conditions such as rain, dust, light levels, connectivity, road surface condition, pedestrian density and energy autonomy. Cost and institutional acceptance will also need to be investigated.

The main threats to validity are the lack of local data, the conventional uninstrumented reference, the lack of repetitions, and the lack of evidence regarding the simulation. An overall confusion matrix also obscures the differences between accidents, obstacles and congestion. Finally, reaction time depends on human and logistical resources that are not solely determined by the software.

The economic analysis will need to compare the costs of acquisition, energy, maintenance, spare parts, connectivity, training and equipment replacement. The acceptability of the interaction will be assessed amongst users and operators. The collection of images must be limited to surveillance requirements and governed by the applicable regulations, which must be verified prior to the pilot scheme.

4.5. Phased roll-out and institutional implications

Future work should focus on developing a prototype. An initial stage could utilise a camera + edge AI + big data + expert system. A second phase would add a mobile robot, a third would involve several junctions, and a fourth would incorporate a digital twin of the road network.

The recommended deployment is phased: needs assessment; proof of concept at a single junction; humanoid prototype and

sensors; evaluation; extension to multiple areas; and industrialisation to ensure high availability.

Such a platform could contribute to monitoring, planning, maintenance and decision support. The system can thus evolve into a Road Intelligence platform.

The first phase should prioritise a baseline measurement and a fixed data-collection installation. The humanoid will only be introduced once its usefulness, physical stability and the safety of its interactions have been verified. Shutdown criteria will include loss of supervision, a deterioration in detection or a safety anomaly. Human supervision and responsibilities fall within the remit of the system’s governance (NIST, 2023). Similarly, the processing of images captured on public roads is governed by the Digital Code, which requires the appointment of a data controller and the conduct of an impact assessment prior to any pilot scheme (section 3.8).



Fig 15: Phased roll-out and institutional implications.

5. Results

5.1. Conceptual deliverables

The results established at the design stage comprise a distributed architecture, an alert processing workflow, an entity model, a set of proposed rules and an evaluation protocol. These constitute conceptual artefacts; no deployed prototype or measured operational benefit is claimed.

Table 1: Comparison of alert processing workflows.

Deliverable	Function	Validation required
Distributed architecture	Linking sensing, data, AI and monitoring	Interface and fault testing
Knowledge base	Putting observations into context	Expert review and rule testing
DevOps/MLOps pipeline	Tracking versions and updates	Testing, rollback and operational monitoring
Comparative protocol	Isolate component contributions	Annotated corpus and repeated trials

5.2. Illustrative numerical comparison

Table 1 sets out illustrative values for the comparative scenario. These are retained as working numerical assumptions. The sample sizes, distributions and logs used to produce these figures are not documented; no confidence intervals can be associated with the differences.

Table 2: Fictitious comparison of surveillance indicators.

Indicator	Conventional	Proposed system	Illustrative difference
Average detection time	8.5 mins	1.9 mins	−77.65%
Average transmission	5.2 mins	0.7 mins	−86.54%
Incidents correctly detected	72 per cent	91 per cent	+19 percentage points
Information availability	78 per cent	96%	+18 points
False positives without a specified denominator	14%	7%	−7 points
Overall reaction	15.4 mins	5.8 mins	−62.34%

Source: fictitious values from the study; differences recalculated. For timeframes, the relative variation is calculated as (proposed value – conventional value) / conventional value. Differences in proportions are expressed in percentage points. Information availability remains an indicator to be defined operationally.

5.3. Confusion matrix and recalculated indicators

Table 2 relates to a separate example of 1,000 fictitious events. It comprises 460 incidents and 540 non-incidents; it is not linked to a temporal comparison with the conventional system.

Table 3: Fictitious confusion matrix.

Actual status	Predicted incident	Predicted non-incident	Total
Incident	410 TP	50 FN	460
Non-incident	40 FP	500 TN	540
Total	450	550	1,000

The calculations give: accuracy = $910 / 1,000 = 91.00$ per cent; precision = $410 / 450 = 91.11$ per cent; recall = $410 / 460 = 89.13$ per cent; F1 = $820 / 910 = 90.11$ per cent. The false positive rate is $40 / 540 = 7.41$ per cent. The proportion of false alarms amongst the alarms issued is $40 / 450 = 8.89$ per cent. These last two proportions have different denominators.

The value of 91 per cent in the matrix refers to accuracy, not the proportion of incidents correctly detected, which corresponds to the recall of 89.13 per cent. The rate of 7 per cent in Table 2 cannot be equated with the false positive rate in the matrix without further information. The two examples are therefore treated separately and must not be aggregated.

5.4. Results yet to be established

No results are available regarding the actual length of queues, energy consumption, acceptability, costs, latency under load or mechanical stability. No conclusions can be drawn regarding the impact on road safety or accident reduction. These aspects will need to be included in the evaluation of the future pilot scheme.

6. Conclusion

This study proposes a design framework for integrating a humanoid system into a road traffic management support platform. The architecture links perception, data processing, statistical intelligence, business rules and human supervision. DevOps and MLOps provide a framework for versioning, testing and monitoring, to be put to the test in operational use.

The contribution is conceptual and methodological. The numerical examples help to explain the indicators and correct their interpretations; they do not validate the hypotheses on the real-world network. Proof of effectiveness requires a locally annotated dataset, a conventional instrumented reference system, comparisons isolating individual components, and reproducible tests.

Further work should begin with the collection and evaluation of data from a fixed platform, followed by an examination of the contribution of the humanoid system. The decision to deploy the system will depend on observed performance, safety, total cost, acceptability and the clarification of institutional responsibilities. The system will remain an assistance tool whose critical recommendations will be subject to human validation.

References

- DORA. (2024). *Accelerate State of DevOps report 2024*. Google Cloud. <https://dora.dev/research/2024/>
- Jiang, W., Luo, J., He, M., & Gu, W. (2023). Graph neural network for traffic forecasting: The research progress. *ISPRS International Journal of Geo-Information*, 12(3), Article 100. <https://doi.org/10.3390/ijgi12030100>
- Kleppmann, M. (2017). *Designing data-intensive applications: The big ideas behind reliable, scalable, and maintainable systems*. O'Reilly Media.
- Kreuzberger, D., Kühl, N., & Hirschl, S. (2023). Machine learning operations (MLOps): Overview, definition, and architecture. *IEEE Access*, 11, 31866–31879. <https://doi.org/10.1109/ACCESS.2023.3262138>
- National Institute of Standards and Technology. (2023). *Artificial intelligence risk management framework (AI RMF 1.0) (NIST AI 100-1)*. <https://doi.org/10.6028/NIST.AI.100-1>

- Pascoe, C., Quinn, S., & Scarfone, K. (2024). *The NIST Cybersecurity Framework (CSF) 2.0 (NIST CSWP 29)*. National Institute of Standards and Technology.
<https://doi.org/10.6028/NIST.CSWP.29>
- Peffers, K., Tuunanen, T., Rothenberger, M. A., & Chatterjee, S. (2007). A design science research methodology for information systems research. *Journal of Management Information Systems*, 24(3), 45–77. <https://doi.org/10.2753/MIS0742-1222240302>
- Russell, S., & Norvig, P. (2021). *Artificial intelligence: A modern approach (4th ed.)*. Pearson.
- Szeliski, R. (2022). *Computer vision: Algorithms and applications (2nd ed.)*. Springer.
<https://doi.org/10.1007/978-3-030-34372-9>
- World Health Organisation. (2023). *Global status report on road safety 2023*.
<https://www.who.int/publications/i/item/9789240086517>
- Goodfellow, I. J., Shlens, J., & Szegedy, C. (2015). *Explaining and Harnessing Adversarial Examples*. International Conference on Learning Representations (ICLR).
- Petit, J., & Shladover, S. E. (2015). Potential Cyberattacks on Automated Vehicles. *IEEE Transactions on Intelligent Transportation Systems*, 16(2), 546–556.
- Democratic Republic of the Congo. (2023). *Ordinance-Law No. 23/010 of 13 March 2023 on the Digital Code*. Official Gazette, 11 April 2023.