

From Fragmented Traffic-Violation Records to Intelligent Monitoring in Kinshasa: A Mixed-Methods Prototype Integrating Big Data, Computer-Vision Readiness and Machine-Learning–Supported Severity Assessment

Bebe Bebe Ngwaba^{1*}; Augustin Pambi Tadiamba²

Teaching Assistant¹; Research²

^{1,2} Department of Computing, Management and English for Businesses, University of Kinshasa, Kinshasa, Democratic Republic of the Congo

¹Corresponding Author: Bebe Bebe Ngwaba^{1*}

Publication Date: 2026/07/21

Abstract: Traffic-violation monitoring in Kinshasa is constrained by discontinuous controls, fragmented case records, and limited traceability. This study develops a socio-technical framework for a traceable traffic-violation intelligence system and tests the feasibility of an analytical prototype in Gombe and Barumbu. A convergent mixed-methods design combined 200 questionnaires, 18 semi-structured interviews, 24 observation sessions, and 22 documentary sources. These materials were triangulated to identify operational needs, institutional constraints, and functional requirements. An anonymised experimental dataset of 500 cases was then used to test a supervised classification workflow for predicting violation-severity levels (low, medium, and high). The prototype used one-hot encoding, a stratified 75/25 train-test split, and a Random Forest classifier with 150 trees and balanced class weights. The field evidence shows distinct local profiles: Gombe records more dynamic violations, particularly speeding and red-light running, whereas Barumbu presents more illegal stopping, road-use conflicts, and overloading. The most pronounced barriers are insufficient digital tools, weak archiving, limited staff training, poor coordination, and unreliable energy or connectivity. On 125 test observations, the prototype reached 52.8% accuracy; high-severity cases achieved the strongest F1-score (0.65). The result supports the feasibility of a structured analytical core for case triage, but it does not validate automated sanctions or a deployed computer-vision system. A phased implementation with human review, data-governance rules, and locally annotated evidence is recommended.

Keywords: Traffic-Violation Monitoring; Data Governance; Machine Learning; Computer Vision; Random Forest; Kinshasa.

How to Cite: Bebe Bebe Ngwaba; Augustin Pambi Tadiamba (2026) From Fragmented Traffic-Violation Records to Intelligent Monitoring in Kinshasa: A Mixed-Methods Prototype Integrating Big Data, Computer-Vision Readiness and Machine-Learning–Supported Severity Assessment. *International Journal of Innovative Science and Research Technology*, 11(7), 462-468. <https://doi.org/10.38124/ijisrt/26jul347>

I. INTRODUCTION

Road safety is both a public-health concern and an urban-governance challenge. The World Health Organization (2023) estimates that road traffic crashes cause approximately 1.19 million deaths each year, with a particularly heavy burden in low- and middle-income countries. Beyond crash prevention, credible monitoring of traffic violations is essential because it

allows institutions to identify recurrent risk patterns, preserve evidence, and direct prevention or enforcement resources toward priority locations.

In Kinshasa, dense road interactions coexist with discontinuous observation, largely paper-based or fragmented records, and uneven capacity to preserve and analyse evidence. The presence of enforcement officers and the use of

sanctions remain important, but they cannot by themselves guarantee continuous observation, comparable records, or timely analytical reporting. A practical response therefore requires more than cameras: it requires an information architecture that connects capture, validation, storage, analysis, reporting, and human review.

This paper addresses two questions: (1) which technical, organisational, and governance factors currently limit credible traffic-violation monitoring in Gombe and Barumbu; and (2) how can these findings be translated into a traceable, scalable design that supports operational prioritisation without overstating what the available data can prove? The objective is to propose and test an analytical framework that links empirical diagnosis, structured data, and supervised classification. The computer-vision component is positioned as a future extension for locally annotated imagery; it is not presented as a deployed or validated visual-detection system.

II. RELATED WORK AND CONCEPTUAL POSITIONING

Intelligent transportation systems increasingly combine administrative records, sensor streams, video, and spatial data to improve safety management. In this setting, Big Data is valuable not merely because of data volume, but because it enables the integration, interoperability, and analysis of heterogeneous information (De Mauro et al., 2016; Kaffash et al., 2021; Montoya-Torres et al., 2021). Computer vision can transform visual scenes into detectable traffic events, while machine-learning models can support classification and prioritisation when the resulting data are sufficiently structured and representative (Coifman et al., 1998; Buch et al., 2011; Szeliski, 2022).

However, transferring such technologies into a local monitoring context requires a socio-technical perspective. Information-system performance depends on the fit between the technical tool, data quality, organisational practices, skills, energy and connectivity, legal safeguards, and user acceptance (Bostrom & Heinen, 1977; Rogers, 2003; Venkatesh et al., 2003). Recent feasibility work on smart mobility in Kinshasa also stresses the importance of context-sensitive deployment rather than direct replication of solutions designed for better-equipped settings (Kayisu et al., 2024). This study therefore treats the prototype as a proof of concept for an analytical workflow, not as evidence that an automated surveillance system is already operational.

III. MATERIALS AND METHODS

➤ Study Setting and Research Design

The fieldwork focused on Gombe and Barumbu, two complementary urban settings in Kinshasa. Gombe is characterised by dense administrative, institutional, and commercial mobility, whereas Barumbu combines commercial and residential activities with more heterogeneous circulation and frequent conflicts over road use. The comparative design was selected to distinguish common system constraints from those related to local mobility patterns.

The study used a convergent mixed-methods and applied-design approach. Quantitative and qualitative materials were triangulated to identify operational needs and translate them into data and functional requirements. This orientation combines the logic of mixed-methods research with the design-science aim of producing an artefact grounded in observed problems (Creswell & Plano Clark, 2018; Hevner et al., 2004).

Table 1 Empirical Corpus and Prototype Data

Source	Volume	Role in the Study
Questionnaires	200	Measure perceptions and compare the two municipalities.
Semi-structured interviews	18	Identify institutional, technical, and human constraints.
Observation sessions	24	Describe behaviours and violations observed in the field.
Documents analysed	22	Contextualise governance standards, practices, and challenges.
Anonymised experimental dataset	500 cases; 25 variables	Test the processing workflow and severity classification.

Source: Authors' Empirical Corpus and Experimental Dataset.

➤ Prototype Data, Preprocessing, and Evaluation

The experimental prototype used an anonymised dataset with a three-class target variable: low, medium, and high violation severity. The data dictionary contained territorial,

event, mobility, and evidence-related fields. The analytical workflow was deliberately limited to structured tabular data, allowing the study to test the full sequence from variable preparation to classification and result interpretation.

Table 2 Variables Used in the Analytical Core

Family	Representative variables	Purpose in the prototype
Territorial context	Municipality, location, time of day	Situate events and compare local patterns.
Road event	Violation type, detection source, visual-evidence status	Define cases and document traceability.
Mobility context	Vehicle type, traffic condition, weather condition	Qualify contextual risk conditions.
Predicted outcome	Severity level: low, medium, high	Prioritise review and intervention.

Source: Prototype Variable Dictionary.

Categorical inputs were encoded through one-hot encoding. A stratified 75/25 train-test split with random_state = 42 was used, producing 375 training observations and 125 test observations. The estimator was a Random Forest classifier with 150 trees and balanced class weights. The encoder and classifier were fitted within the training pipeline to reduce the risk of information leakage. Model reporting used precision, recall, F1-score, accuracy, and a confusion matrix. Predictions were interpreted as a triage signal for review and indicator production, never as an autonomous sanction decision.

IV. RESULTS

➤ Local Profiles of Observed Violations

The 24 observation sessions show that the two municipalities do not display the same traffic-risk profile. Gombe shows more dynamic violations, especially speeding and red-light running. Barumbu records more illegal stopping, overloading, and road-use conflicts. This difference supports the need for configurable alerts and locality-sensitive priorities rather than one undifferentiated monitoring policy.

Table 3 Observed Violations During Field Sessions

Observed violation	Gombe	Barumbu
Speeding	31	18
Running a red light	22	13
Illegal stopping or parking	19	27
Using a phone while driving	14	9
Failure to yield / wrong-way driving	11	18
Overloading / irregular loading	7	21
No helmet or seat belt	9	14

Source: Results of the 24 Field-Observation Sessions.




Fig 1 Aggregated Profiles of Violations by Municipality

Source: Authors' Aggregation of the Detailed Observations Reported in Table 3.

➤ Constraints that Limit Monitoring

The questionnaire and interview results indicate that the problem is not limited to the absence of cameras. Respondents report major gaps in digital equipment, case archiving,

technical training, coordination, and the practical conditions needed to maintain a reliable service. In general, the perceived constraints are more pronounced in Barumbu, notably for equipment, archiving, training, and energy or connectivity.

Table 4 Perceived Intensity of Barriers to Monitoring (Scale: 1 to 5)

Barrier	Gombe	Barumbu
Lack of digital tools	4.23	4.41
Insufficient data archiving	4.11	4.28
Insufficient technical training	3.96	4.12
High implementation cost	3.85	3.93
Weak inter-institutional coordination	3.74	3.91
Incomplete legal framework	3.69	3.84
Energy instability and weak connectivity	3.48	3.87

Source: Questionnaire Results and Descriptive Analysis.

Interview evidence converges with these descriptive results. The recurring themes were the need for digital records, staff training, preserved case histories, continuous or regular monitoring, coordination among services, and dependable energy or connectivity. These factors show that modernisation is simultaneously a technical, organisational, and governance task.

➤ *Translation of the Diagnosis into Design Requirements*

The empirical diagnosis was converted into functional requirements so that each observed weakness corresponds to a verifiable information or operational response. This step avoids treating digitisation as an end in itself and anchors the proposed system in traceability, comparability, and evidence quality.

Table 5 Correspondence Between Observed Problems and Proposed Functional Responses

Observed problem	Empirical manifestation	Proposed functional response
Insufficient traceability	Cases are difficult to follow and archives are fragmented.	Digital record with timestamp and case identifier.
Scattered information	Multiple actors produce information without a consolidated view.	Relational database and centralisation rules.
Intermittent control	Observation depends on the punctual presence of agents.	Structured event capture and configurable alerts.
Limited local comparison	Different violation profiles are not consistently described.	Territorial dashboards and indicators by area.
Evidence difficult to use	Events are poorly documented and visual archiving is limited.	Controlled evidence repository and progressive computer-vision module.

Source: Authors' Feasibility Matrix and Prototype Architecture.




Fig 2 Functional Architecture of the Proposed Prototype
Source: Authors' Design Based on Field Requirements.

➤ *Timing and Classification Results*

The experimental dataset highlights two sensitive periods: morning (38.4%) and evening (41.6%). Together, these periods represent 80.0% of the recorded cases. Such temporal indicators can support the prioritisation of pilot locations, deployment of field teams, and the definition of future visual-analysis scenarios.




Fig 3 Distribution of Cases in the Experimental Dataset by Time of Day
Source: Descriptive Output from the Experimental Prototype.

Table 6 reports the performance of the Random Forest model on the 125-case test set. Overall accuracy was 52.8%. The high-severity class produced the strongest F1-score (0.65), followed by medium severity (0.50) and low severity

(0.36). The results show that the workflow can help organise case review, but the modest overall performance means that predictions must remain subject to human verification.

Table 6 Random Forest Performance on the Test Set

Severity class	Precision	Recall	F1-score	Support
Low	0.44	0.31	0.36	26
Medium	0.46	0.54	0.50	54
High	0.66	0.64	0.65	45
Overall accuracy	0.528	-	-	125

Source: Outputs of the Random Forest Prototype.




Fig 4 Confusion Matrix for the Random Forest Severity Classification

Source: Outputs of the Random Forest Prototype.

The confusion matrix shows that the medium-severity class is the most difficult to distinguish. It contains cases classified as low or high, which suggests that the boundaries between severity levels are not fully captured by the available variables. In operational use, any medium- or high-severity signal should therefore be reviewed against the timestamp, location, visual evidence where available, and other information recorded in the digital case file.

V. DISCUSSION

The main contribution of this study is to show that intelligent traffic monitoring cannot be reduced to the acquisition of cameras or software. The proposed framework places data architecture and governance at the centre of the system: events must first be captured consistently, assigned a case identifier, stored securely, and made comparable before a classification model can produce a useful signal. This interpretation is consistent with work on Big Data and intelligent transportation systems, which emphasises integration, interoperability, and information quality rather than data volume alone (De Mauro et al., 2016; Montoya-Torres et al., 2021).

The comparison of Gombe and Barumbu also demonstrates why local configuration matters. Dynamic violations in Gombe may require different alert conditions from the illegal stopping, overloading, and road-use conflicts that are more visible in Barumbu. A scalable system should consequently maintain common core data fields while allowing dashboards, violation catalogues, and alert thresholds to be adapted to each priority area.

The classification result must be read cautiously. An accuracy of 52.8% on a limited experimental dataset is not adequate for direct enforcement, automatic sanctioning, or generalisation to all of Kinshasa. It only demonstrates that a reproducible data-preparation and classification chain can operate on structured cases. Research on machine learning for road safety similarly shows that performance depends heavily on variable quality, class balance, data representativeness, and the reliability of ground-truth labels (Silva et al., 2020; Chai et al., 2024).

Computer vision should likewise be positioned accurately. This prototype did not train a convolutional neural network or evaluate annotated video footage from Kinshasa.

The visual component remains an architectural extension that could later generate structured event records for the analytical core. A responsible next step is therefore to collect and annotate local images under a documented protocol, then

evaluate visual models separately before connecting them to operational workflows (Razi et al., 2022; Thompson et al., 2024).

Table 7 Conditions for a Phased Deployment

Condition	Operational justification
Energy and connectivity	Maintain continuity of data collection, storage, and processing.
Agent training	Support adoption, alert review, and data-entry quality.
Legal clarity	Regulate access, retention, use, and preservation of digital evidence.
Inter-institutional governance	Coordinate responsibilities, escalation rules, and decisions.
Data quality management	Control duplicates, inconsistencies, missing values, and unnecessary variables.
Maintenance and cybersecurity	Preserve availability, integrity, confidentiality, and auditability.

Source: Deployment Conditions Identified from the Field Diagnosis and Prototype Requirements.

A phased strategy is recommended. The first phase should focus on a small number of pilot intersections, a minimal case record, clear review rules, and basic dashboards. The second phase can integrate controlled image streams and a shared database. A later phase can pursue interoperability with relevant public services, provided that evidence retention, access rights, audit trails, and maintenance responsibilities are clearly defined. Across all phases, a human decision-maker should remain responsible for contextual review and final action.

VI. CONCLUSION

This study developed a traceable framework for intelligent traffic-violation monitoring in Gombe and Barumbu, Kinshasa. The field evidence confirms that the main constraints combine weak digital tools, fragmented archives, limited training, coordination gaps, and unstable operational infrastructure. These findings support a design in which data capture, secure case management, analysis, and human review are treated as a single information chain.

The 500-case experimental prototype shows that a Random Forest classifier can assist severity triage once cases are structured and documented. However, its 52.8% test accuracy does not justify automatic sanctions or broad operational claims. The next research and implementation priorities are the creation of an anonymised institutional database, multi-period data collection, local annotation of visual evidence, evaluation of visual models, and field testing of whether alerts improve officers' review practices. The final decision should remain human, contextualised, and verifiable.

VII. LIMITATIONS AND FUTURE RESEARCH

This study should be interpreted as a field-grounded design and proof-of-concept exercise rather than as a city-wide validation of an operational surveillance platform. The empirical investigation was limited to Gombe and Barumbu, two municipalities selected for their contrasting mobility patterns. Their results are informative for the proposed framework, but they cannot represent every road environment, enforcement practice, or infrastructure condition found across Kinshasa. In addition, perceptions reported through questionnaires and interviews may be influenced by the

respondents' direct experience, institutional role, and expectations regarding digital tools.

The analytical component also has clear boundaries. The prototype was built from 500 anonymised experimental records and evaluated through a single stratified train-test split. It was designed to verify the coherence of the data-preparation, encoding, and severity-classification workflow, not to establish a production-level predictive accuracy. The severity labels depend on the available variable dictionary and coding rules; relevant factors such as road-surface condition, pedestrian density, repeat-offence history, local lighting, and the quality of visual material were not represented in sufficient detail. Consequently, the reported accuracy must not be treated as a basis for automated enforcement or sanctions.

A further limitation concerns the visual dimension of the proposed architecture. No locally annotated image or video corpus was used to train or benchmark a computer-vision model. Computer vision is therefore presented as a future data-acquisition layer rather than as an empirically validated capability of this prototype. The study also did not test interoperability with police, municipal, judicial, or telecommunications information systems, nor did it conduct a formal privacy-impact assessment. These omissions are significant because reliable deployment depends as much on institutional safeguards and operating procedures as on algorithmic performance.

Future research should extend the study through a larger multi-period institutional dataset collected under a shared recording protocol. Repeated validation across municipalities, intersections, seasons, and peak periods would allow researchers to assess temporal stability, class imbalance, calibration, and potential bias. A next phase could build an ethically governed local image dataset, compare visual detection models with structured-record classifiers, and test whether alerts improve review speed and decision quality for officers. Participatory work with road users, enforcement services, legal authorities, and data-protection stakeholders would also be needed to define retention periods, access controls, appeal procedures, and the appropriate limits of automated support in public road governance.

REFERENCES

- [1]. Bostrom, R. P., & Heinen, J. S. (1977). MIS problems and failures: A socio-technical perspective. Part I: The causes. *MIS Quarterly*, 1(3), 17-32. <https://doi.org/10.2307/248710>
- [2]. Buch, N., Velastin, S. A., & Orwell, J. (2011). A review of computer vision techniques for the analysis of urban traffic. *IEEE Transactions on Intelligent Transportation Systems*, 12(3), 920-939. <https://doi.org/10.1109/TITS.2011.2119372>
- [3]. Chai, A. B. Z., Lau, B. T., Tee, M. K. T., & McCarthy, C. (2024). Enhancing road safety with machine learning: Current advances and future directions in accident prediction using non-visual data. *Engineering Applications of Artificial Intelligence*, 133, 109086. <https://doi.org/10.1016/j.engappai.2024.109086>
- [4]. Coifman, B., Beymer, D., McLauchlan, P., & Malik, J. (1998). A real-time computer vision system for vehicle tracking and traffic surveillance. *Transportation Research Part C: Emerging Technologies*, 6(4), 271-288. [https://doi.org/10.1016/S0968-090X\(98\)00019-9](https://doi.org/10.1016/S0968-090X(98)00019-9)
- [5]. Creswell, J. W., & Plano Clark, V. L. (2018). *Designing and conducting mixed methods research* (3rd ed.). SAGE.
- [6]. De Mauro, A., Greco, M., & Grimaldi, M. (2016). A formal definition of Big Data based on its essential features. *Library Review*, 65(3), 122-135. <https://doi.org/10.1108/LR-06-2015-0061>
- [7]. Hevner, A. R., March, S. T., Park, J., & Ram, S. (2004). Design science in information systems research. *MIS Quarterly*, 28(1), 75-105. <https://doi.org/10.2307/25148625>
- [8]. Kaffash, S., Nguyen, A. T., & Zhu, J. (2021). Big data algorithms and applications in intelligent transportation systems: A survey and bibliometric analysis. *Sustainable Cities and Society*, 75, 103299. <https://doi.org/10.1016/j.scs.2021.103299>
- [9]. Kayisu, A. K., Mikušová, M., Bokoro, P. N., & Kyamakya, K. (2024). Exploring the potential of smart mobility in Kinshasa (DR Congo) as a means of addressing traffic congestion and improving road safety: A comprehensive feasibility assessment. *Sustainability*, 16(21), 9371. <https://doi.org/10.3390/su16219371>
- [10]. Montoya-Torres, J. R., Moreno, S., Guerrero, W. J., & Mejía, G. (2021). Big data analytics and intelligent transportation systems. *IFAC-PapersOnLine*, 54(2), 216-220. <https://doi.org/10.1016/j.ifacol.2021.06.025>
- [11]. Razi, A., Chen, Y., Yu, H., & Zhang, Y. (2022). Deep learning for traffic safety analysis: A forward-looking review. *IET Intelligent Transport Systems*. <https://doi.org/10.1049/itr2.12257>
- [12]. Rogers, E. M. (2003). *Diffusion of innovations* (5th ed.). Free Press.
- [13]. Silva, P. B., Andrade, M., & Ferreira, S. (2020). Machine learning applied to road safety modeling: A systematic literature review. *Journal of Traffic and Transportation Engineering (English Edition)*, 7(6), 775-790. <https://doi.org/10.1016/j.jtte.2020.07.004>
- [14]. Szeliski, R. (2022). *Computer vision: Algorithms and applications* (2nd ed.). Springer. <https://doi.org/10.1007/978-3-030-34372-9>
- [15]. Thompson, J., Lowry, M., & Abdel-Rahim, A. (2024). *Computer vision for traffic monitoring*. Washington State Transportation Center, Washington State University.
- [16]. Venkatesh, V., Morris, M. G., Davis, G. B., & Davis, F. D. (2003). User acceptance of information technology: Toward a unified view. *MIS Quarterly*, 27(3), 425-478. <https://doi.org/10.2307/30036540>
- [17]. World Health Organization. (2023). *Global status report on road safety 2023*. WHO.