Motivation and Feasibility of an AI-Enabled Road Safety Planning and Monitoring Framework for African Contexts

Kyandoghere Kyamakya[1], Mahmoud Mohamed[1], Asha Ashokkumar Choudhary[1], Claude Takenga[2], Kamala Mitume Boni [2], Héritier Kipimo[2], Nsenge Mpia Héritier[3], Wido Hamel[4], Tomasz Mackun[5], Wojciech Kustra[5]

[1]Alpen–Adria Klagenfurt University (AAU), Austria

[2]Université Officielle de Ruwenzori (UOR), Democratic Republic of Congo

[3]Université de l'Assomption au Congo (UAC), Democratic Republic of Congo

[4]Bauhaus Universitaet Weimar, Germany

[5]Gdansk University of Technology (GUT), Poland

***Corresponding author:** Claude Takenga: takenga@yahoo.fr

Keywords: Road safety; network screening; Africa; DRC; vulnerable road users; road geometry; speed management; multi-modal AI; large language models; human-in-the-loop.

Background

Road safety work requires decisions under uncertainty: where is risk concentrated, which mechanisms dominate (speed, visibility, conflict points, roadside hazards), and which interventions are feasible within budgets and institutional capacity. Common failure modes include inconsistent geometric parameters (curve radii, crossfall/superelevation), insufficient sight distance, unsafe access density near markets and schools, drainage-related edge breaks, inadequate delineation/markings, and unprotected roadside objects. In large territories such as the DRC, these issues recur across paved and unpaved networks and are exacerbated by seasonal disruptions and limited maintenance.

AI can help by extracting comparable safety-relevant features from remote sensing and opportunistic field data, enforcing safety checks systematically, and generating reproducible risk rankings. The feasibility question is whether useful, auditable outputs can be produced with the data that are realistically available—and how uncertainty is communicated to engineers and stakeholders.

Feasibility in data-constrained African contexts

Feasibility depends on a "progressive enhancement" approach: deliver value with minimal data, then refine as measurements and institutional processes mature. Baseline inputs often available include public satellite imagery, global elevation models, rainfall/flood proxies, road centerlines, and administrative boundaries. Targeted field inputs can be collected at low cost via smartphones (GNSS tracks, accelerometer roughness proxies), vehicle video, and short safety-audit surveys on critical segments.

Governance is central: a human-in-the-loop workflow is assumed where AI proposes risk diagnostics and compliance warnings, while licensed road safety and design engineers approve decisions. Outputs must flag low-confidence segments, document assumptions, and support traceable review (engineering sign-off, stakeholder consultation, and donor reporting).

Table 1 — Illustrative coverage of safety-relevant evidence and outputs

Safety mechanism	Minimal data sources	AI-derived outputs	Engineering actions
Speed & operating-speed mismatch	Geometry + slope from DEM; video-based speed samples	Speed-consistency flags; credible-speed map	Speed management; design-speed review; traffic calming
Visibility & sight distance	DEM; curve/grade reconstruction; roadside imagery	Sight-distance warnings; crest/sag risk markers	Reprofile; clear zones; signing/markings
Conflict points & access density	OSM/imagery; POIs (markets/schools); traffic counts (spot)	Hotspot clustering; VRU exposure index	Access management; crossings; protected shoulders
Roadside hazards	Imagery; video; inventory of poles/trees/ditches	Hazard proximity score; barrier need flags	Remove/relocate; guardrails; forgiving roadsides
Weather/drainage-related crash risk	Rainfall/flood proxies; drainage paths; field notes	Wet-risk segments; shoulder/edge failure cues	Drainage fixes; skid resistance treatments

Proposed multi-modal, LLM-assisted road safety framework

The framework is organised into four operational layers:

- Multi-modal evidence layer: extract slope/curvature, visibility proxies, roadside hazard indicators, access density, and delineation/marking cues from satellite/UAV imagery, elevation models, and vehicle/smartphone sensing.
- Safety rules and compliance checks: implement geometry and safety checks (design-speed consistency, stopping-sight-distance warnings, curve-radius plausibility, crossfall/superelevation plausibility) and screen for common VRU hazards (schools/markets, missing shoulders, unsafe crossings).
- Risk modelling and prioritisation: produce corridor and site risk scores with uncertainty bands, explain dominant mechanisms, and rank candidate interventions by expected risk reduction and feasibility (quick-build vs capital works).
- LLM governance co-pilot: retrieval-augmented assistance over approved standards and agency templates to generate auditable checklists, road-safety-audit drafts, risk registers, and monitoring plans; all outputs include explicit assumptions and engineer sign-off fields.

The core contribution is the coupling of uncertainty-aware safety diagnostics with an LLM layer that improves traceability, documentation consistency, and stakeholder communication—without replacing engineering accountability.

Priority use cases

The initial deployment should concentrate on high-leverage workflows that do not require complete crash datasets:

- Network screening: combine sparse crash records with imagery/geometry-derived risk indicators to identify candidate high-risk corridors for safety audits and enforcement targeting.
- Design and rehabilitation safety checks: automatically flag segments where geometry and expected operating speeds imply elevated crash severity (curves, crests, junction approaches).
- VRU protection planning: map exposure around schools, markets, and settlements and prioritise low-cost treatments (shoulders, crossings, signage, lighting) paired with speed management.
- Before–after learning: standardise pre- and post-intervention monitoring using repeatable sensing (video, smartphone tracks) to validate risk reduction and update local safety parameters.

Evaluation and DRC-focused pilot roadmap

Evaluation should balance technical validity and decision usefulness: (i) accuracy of geometry and safety flags on audited sites; (ii) stability and interpretability of corridor/site risk rankings under data uncertainty; (iii) time saved in producing audit-ready documentation; and (iv) observed changes in surrogate safety indicators and, where available, crash outcomes.

A pragmatic pilot can proceed in three phases: Phase 1 builds a baseline geospatial safety layer for 2–3 corridors with known risk; Phase 2 deploys automated safety checks and network screening with targeted field sensing; Phase 3 applies prioritised countermeasures and establishes post-implementation monitoring and reporting with a reproducible audit trail.

Conclusion

An AI-enabled road safety planning and monitoring framework is feasible for African contexts when designed for low-data realities and strong governance: multi-modal sensing provides scalable evidence, rule-based checks institutionalise safety-by-design, and an LLM governance layer standardises auditable safety artefacts while preserving human accountability. The resulting pipeline supports faster, more transparent prioritisation of interventions aimed at reducing crash risk and severity across mixed road types and user groups.