# KINBOT-AI: Design, Implementation and Controlled Demonstration of a Resource-Constrained Edge-AI Mobile Robot for Mixed-Traffic Observation in Kinshasa

Kinkete Mfumabi Hervé[1]; Belo Belaris Kevin[2]; Longa Kapesa Emilie[2]; Tshibangu Lutumba Israêl[2]; Makengele Lukusa Gradi[2]; Mabasa Ndontoni John[2]; Mbiye Tshiama Sharonne[2]

[1]Principal Author and Academic Secretary of the Faculty
[2]Student Authors
[1;2]Faculty of Computer Science and Artificial Intelligence, William Booth University Kinshasa, Democratic Republic of the Congo

Publication Date: 2026/08/29

**Abstract: Mixed traffic in Kinshasa combines cars, motorcycles, and pedestrians on roads with limited fixed sensing and intermittent connectivity. This paper presents KINBOT-AI as a design, implementation, and controlled-demonstration artifact: a resource-constrained mobile Edge-AI platform that performs local perception, tracking, counting, queue estimation, safety control, event logging, and human-supervised alert handling. The evaluation comprised 40 navigation missions on a closed 85 m route, 25 counting sequences, 15 queue scenarios, 12 auxiliary environmental traces, 20 simulated alerts, and eight forced 30 s network outages. The prototype completed 37/40 nominal missions (92.5%; 95% Wilson interval: 80.1–97.4). Classwise F1 scores were 0.87 for vehicles, 0.82 for motorcycles, and 0.79 for pedestrians. Counting errors recalculated from the reported manual and automatic totals were 8.6%, 14.1%, and 16.5%, respectively. Queue-length MAE was 1.7 units, average autonomy to 20% battery was 58 min, and median alert latency was 1.4 s. During all eight outages, local obstacle avoidance continued and no logged event was lost or duplicated after reconnection. These measurements characterize the prototype only within the controlled study environment and must not be interpreted as city-wide operational performance.**

***Keywords:*** *Edge AI, Mobile Robot, Mixed Traffic, Object Detection, Human-in-the-Loop, Intermittent Connectivity, Kinshasa.*

**How to Cite:** Kinkete Mfumabi Hervé; Belo Belaris Kevin; Longa Kapesa Emilie; Tshibangu Lutumba Israêl; Makengele Lukusa Gradi; Mabasa Ndontoni John; Mbiye Tshiama Sharonne (2026) KINBOT-AI: Design, Implementation and Controlled Demonstration of a Resource-Constrained Edge-AI Mobile Robot for Mixed-Traffic Observation in Kinshasa. *International Journal of Innovative Science and Research Technology*,11(8), 1919-1926. https://doi.org/10.38124/ijisrt/26aug908

## I. INTRODUCTION

Urban traffic observation increasingly depends on camera networks, roadside sensors, and centralized processing. Such architectures can provide continuous coverage, but they also assume stable power, adequate backhaul, and sustained infrastructure investment [1], [2]. These assumptions do not always hold in rapidly growing cities where heterogeneous traffic, informal stopping patterns, and limited instrumentation coexist.

Kinshasa illustrates this mismatch. Private vehicles, taxis, motorcycles, pedestrians, and roadside activity share constrained road space, while continuous public traffic series remain scarce [3], [4]. Low-cost air-sensing studies also document the broader shortage of reference-grade urban monitoring infrastructure [5]. A mobile observer cannot replace a city sensor network, but it can temporarily carry sensing and computation to controlled sites where fixed infrastructure is unavailable.

Edge computing reduces bandwidth and remote-service dependence by moving inference and safety-relevant decisions close to the data source [6], [7]. Recent fixed-site traffic systems have demonstrated compressed on-site video analysis [8], deep-learning traffic infrastructure [9], YOLOv11 with BoT-SORT on Jetson Orin Nano [10], and synchronization strategies for low-end traffic devices [11]. However, these

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26aug908

works predominantly study fixed cameras or infrastructure nodes. The narrower gap addressed here is the integrated operation of traffic observation, independent local safety, event buffering, and human supervision on a lightweight mobile platform evaluated in an infrastructure-constrained African context.

This work answers four research questions: RQ1—How can a resource-constrained mobile robot be architected for local operation? RQ2—How can the architecture be implemented using accessible hardware classes? RQ3—Can the integrated prototype execute its intended functions in a controlled environment? RQ4—Do critical local functions and event integrity persist during a temporary communication outage?

The contributions are: C1, an Edge-AI mobile architecture separating vision analytics from an independent ultrasonic safety layer; C2, a local implementation integrating detection, tracking, counting, queue estimation, event storage, communication, and human supervision; C3, a continuity mechanism based on identifiers, buffering, replay, and deduplication; and C4, a transparent controlled demonstration reporting both successful and weak outcomes without claiming open-road validation.

## II. RELATED WORK

➢ *Traffic Monitoring and Edge AI*

Classical intelligent transportation systems use loops, radars, wireless sensor networks, and fixed cameras [1], [2]. Modern detectors from the YOLO family and lightweight backbones enable increasingly capable on-device inference [12]–[14]. Lee et al. [8] reported an edge smart-intersection deployment with model compression and near-real-time processing. Villa et al. [9] combined YOLOv5 and Deep SORT for fixed traffic infrastructure, while documenting degradation under poor weather and occlusion.

➢ *Recent 2025 Developments*

Gayas et al. [10] combined YOLOv11, BoT-SORT, directional counting, and a Jetson Orin Nano for intersection monitoring. Li and Hou [11] addressed synchronization of YOLO traffic detection on low-end devices. These systems raise the contemporary baseline for detector documentation, tracking, and resource characterization. KINBOT-AI does not claim superior detection performance; its contribution is mobility plus safety/communication separation and controlled outage continuity in Kinshasa.

➢ *Mobile Robotics and Human Supervision*

Reactive obstacle avoidance and low-speed waypoint navigation are established mobile-robot strategies [15]–[17]. In safety-sensitive observation, human-in-the-loop validation remains appropriate when false alerts and incomplete generalization prevent automatic physical action. Accordingly, KINBOT-AI treats alert confirmation as a supervisory decision rather than an autonomous traffic-control command.

Table 1 Comparative Contribution Matrix (✓ = Addressed; — = Not Reported as a Central Feature).

| Work | Mobile | Edge | Veh. | Moto | Ped. | Track | Count | Queue | Air | Offline | Human | Africa |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Zhang et al. [1] | — | Partial | ✓ | Partial | Partial | Partial | ✓ | Partial | — | — | — | — |
| Lee et al. [8] | — | ✓ | ✓ | ✓ | ✓ | — | ✓ | ✓ | — | — | — | — |
| Villa et al. [9] | — | ✓ | ✓ | ✓ | — | ✓ | ✓ | — | — | — | — | — |
| Gayas et al. [10] | — | ✓ | ✓ | ✓ | — | ✓ | ✓ | — | — | — | — | — |
| KINBOT-AI | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | Aux. | ✓ | ✓ | ✓ |

## III. DESIGN REQUIREMENTS AND ACCEPTANCE CRITERIA

➢ *Design Requirements*

| ID | Requirement |
|---|---|
| DR1 | Primarily local operation |
| DR2 | Low dependence on external connectivity |
| DR3 | Resource-constrained accessible hardware |
| DR4 | Three-class mixed-traffic observation |
| DR5 | Low-speed experimental safety envelope |
| DR6 | Independent emergency stop and obstacle layer |
| DR7 | Local event logging |
| DR8 | Human validation of sensitive alerts |
| DR9 | Modular hardware/software organization |
| DR10 | Use restricted to a controlled institutional site |

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26aug908

➢ *Demonstration Criteria (DC1–DC5)*

DC1—complete the controlled route while applying local stop/avoidance rules. DC2—detect, track, and count vehicles, motorcycles, and pedestrians. DC3—transform observations into counts and a queue-state estimate. DC4—deliver events to a supervisor who retains the final decision. DC5—maintain critical local operation and event integrity during temporary link loss.

➢ *Quantitative Acceptance Criteria (AC1–AC5)*

The quantitative thresholds are separated from the functional criteria: AC1—nominal mission rate ≥80%; AC2—vehicle counting error <15%; AC3—queue MAE ≤2 units; AC4—median alert latency <3 s on the local test network; AC5—during a forced 30 s outage, local avoidance continues and replay produces neither event loss nor duplication. These thresholds characterize the demonstration envelope and are not city-scale service guarantees.

## IV. SYSTEM ARCHITECTURE AND IMPLEMENTATION

➢ *Hardware Configuration*

The evaluated prototype used a four-wheel differential-drive chassis, a 12 V battery, an ARM single-board computer in the Raspberry Pi 4 class with 8 GB RAM, four HC-SR04-type ultrasonic sensors sampled at 10 Hz, a forward USB color camera acquiring 640×480 video at 30 frames/s, a PMS-family optical particle sensor, a temperature probe, a 1 Hz GNSS receiver, and two ESP32-class microcontrollers for acquisition and low-level actuation. Loaded mass was approximately 11 kg and test speed was limited to 0.6 m/s. A physical button and a console command provided emergency stop.

Exact commercial identifiers for the camera, GNSS, PMS unit, motors, motor driver, battery capacity, storage, power converters, and Wi-Fi interface were not preserved in the experimental log. Consequently, the study supports functional replication by hardware class, not component-identical replication. This limitation is reported rather than reconstructed retrospectively.

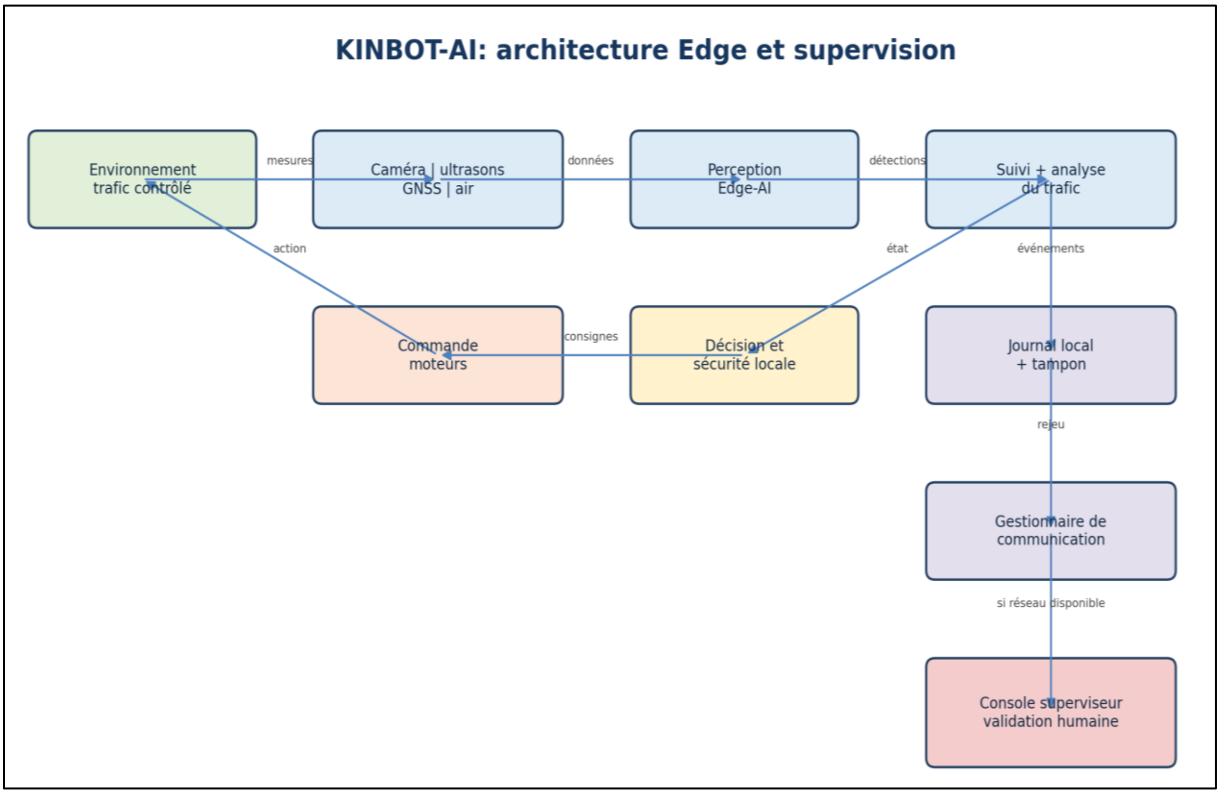


Fig 1 Master Architecture Separating Perception, Local Safety, Event Storage, Communication, and Supervision.

➢ *Software Services*

| Service | Input | Processing | Output | Rate |
|---|---|---|---|---|
| Vision | Frame | YOLO nano | Detections | 8–11 FPS |
| Tracking | Boxes | SORT, IoU≥0.30 | Track IDs | Per frame |
| Safety | Ultrasound | 80/40/50 cm rules | Motor command | 10 Hz |
| GNSS | Signal | Position | Lat./long. | 1 Hz |
| Events | Analytics | ID + local buffer | Short message | Event |
| Supervisor | Event | Human review | Accept/reject | Event |

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26aug908

➢ *Vision and Tracking*

A compact one-stage YOLO-family detector in a nano configuration was initialized from COCO-type pretrained weights and adapted to vehicles, motorcycles, and pedestrians [13], [14]. The local dataset contained 1,240 images split by recorded sequence—not isolated frame—into 868 training, 186 validation, and 186 test images. The held-out test set contained 612 annotated instances: 248 vehicles, 201 motorcycles, and 163 pedestrians. Input size was 640 pixels, confidence threshold 0.35, and NMS IoU threshold 0.45. SORT used a minimum association IoU of 0.30 and a maximum track age of 10 frames [18].

The exact YOLO release, epochs, batch size, learning rate, optimizer, augmentation schedule, random seed, training device, export format, and inference-framework version were not retained. The reported model therefore cannot be reconstructed bit-for-bit. The accompanying supplement freezes every parameter that is known and marks unknown fields explicitly; it does not fabricate them.

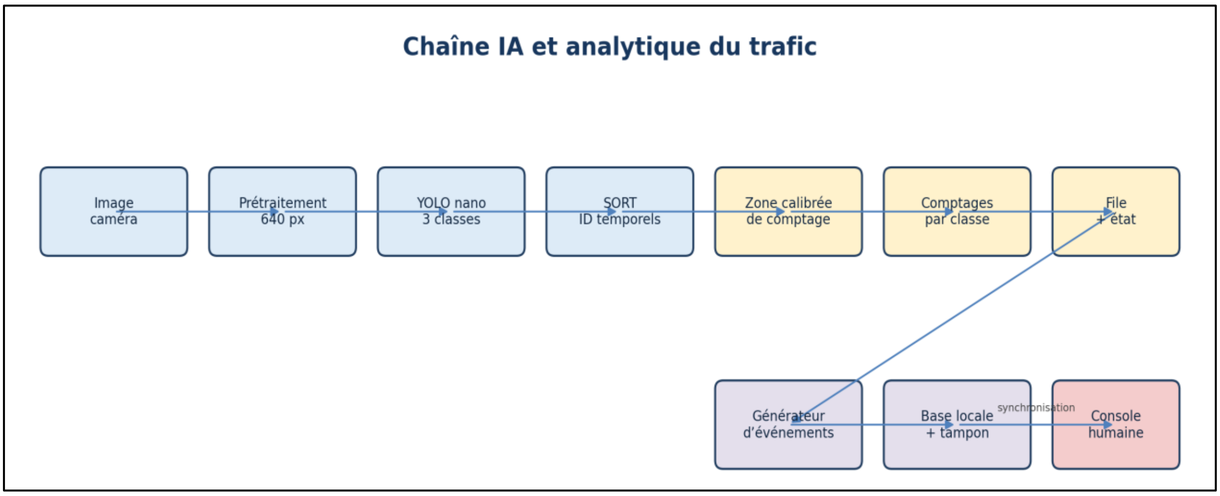

Fig 2 AI and Traffic-Analytics Chain from Camera Frame to Local Event Storage and Human Validation.

➢ *Queue Estimation*

A queue was operationally defined as at least three same-direction road users separated by less than 2.5 m, moving below 0.15 m/s for at least 4 s. Metric proximity was approximated using pre-calibrated ground-plane zones referenced to measured route markers; bottom points of tracked boxes were assigned to zones. This was not continuous homography or 3-D reconstruction and requires recalibration if camera geometry changes. Queue state was free (0–2 units), slowed (3–5), or saturated (≥6) within the useful field of view.

➢ *Safety and Communication*

The safety path was independent of YOLO. Above 80 cm frontal clearance the robot followed its command; from 80 to 40 cm it slowed; below 40 cm it stopped or bypassed the obstacle when lateral clearance exceeded 50 cm. Safety commands overrode higher-level instructions. Each event carried an identifier, time, severity, human-readable label, approximate position, and optional thumbnail. During link loss, events remained in a local buffer and were replayed after reconnection; the identifier supported deduplication.

## V. CONTROLLED DEMONSTRATION PROTOCOL

Testing occurred during daytime on a closed university parking area and peripheral lane in Kinshasa. The 85 m route included two right-angle turns, a simulated queue area, a marked pedestrian corridor, irregular pavement, and two marked potholes. Width ranged from 3.2 to 5.8 m. No open-road, night, or heavy-rain trial was conducted, and no command was connected to an operational traffic signal.

Six scenarios were used: S1—40 paired navigation missions; S2—25 counting sequences of 90 s; S3—15 queue scenarios with 3, 5, or 8 units (five repetitions each); S4—20 simulated intrusion or stop alerts; S5—eight forced 30 s access-point outages; and S6—12 auxiliary environmental traces of 8 min. Manual counting by two annotators was the reference. Agreement on 200 boxes used Cohen's kappa; disagreements were jointly adjudicated. The same chassis under cautious teleoperation provided a navigation reference. A more powerful local workstation provided an informal latency/connectivity comparison.

Primary measures were nominal mission rate, classwise precision/recall/F1, relative counting error, queue-length MAE, processing latency, autonomy to 20% battery, alert latency, and event integrity after reconnection. Binomial proportions use 95% Wilson intervals [19]. Because sample sizes were small, results are descriptive rather than evidence of operational superiority.

## VI. RESULTS

➢ *Navigation, Resources, and Interface*

Thirty-seven of 40 autonomous missions were nominal (92.5%; Wilson 95%: 80.1–97.4), satisfying AC1 on the Wilson lower bound. The three non-nominal missions comprised two corridor departures near the deeper pothole and one emergency stop caused by a close pedestrian; the emergency stop was safety-correct but counted conservatively as non-nominal. Teleoperation completed 38/40 missions (95.0%; Wilson 95%: 83.5–98.6). Overlapping intervals do

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26aug908

not support a claim that autonomous driving outperformed teleoperation.

Average autonomy to 20% battery was 58 min (range 51–64; n=38 after excluding two mechanical interruptions) with active vision, compared with 67 min over ten teleoperated discharges without the detector. CPU load averaged 71% (SD 8%) and peaked at 88%; inference ran at 8–11 FPS. Median local processing latency was 94 ms (p95 128 ms; 1,000 timestamps). This describes the Edge vision/event chain, whereas ultrasonic safety operated independently at 10 Hz.

Ten non-developers completed three guided console tasks. Nine completed all tasks without assistance after a two-minute demonstration (90%; Wilson 95%: 59.6–98.2). Median times were 18 s to start a mission, 7 s to read a count, and 5 s to validate an alert. The sample is too small for generalized usability claims.

Fig 3 Supervisor Console Showing State, Counts, Queue Status, and a Pending Human-Validated Alert.

➢ *Detection and Counting*

Inter-annotator agreement on 200 boxes was κ=0.84. Test-set outcomes were: vehicles TP=213, FP=26, FN=35; motorcycles TP=159, FP=30, FN=42; pedestrians TP=125, FP=29, FN=38. The resulting rounded precision/recall/F1 values are reported in Table II. Across 25 counting sequences, manual totals were 186, 142, and 97, whereas automatic totals were 170, 122, and 81. Recalculated relative errors are therefore 8.6%, 14.1%, and 16.5%. AC2 applies only to vehicles and was satisfied.

Table 2 Held-Out Detection Results and Sequence-Level Counting Errors.

| Class | TP/FP/FN | Prec. | Recall | F1 | Count error |
|---|---|---|---|---|---|
| Vehicle | 213/26/35 | 0.89 | 0.86 | 0.87 | 8.6% |
| Motorcycle | 159/30/42 | 0.84 | 0.79 | 0.82 | 14.1% |
| Pedestrian | 125/29/38 | 0.81 | 0.77 | 0.79 | 16.5% |

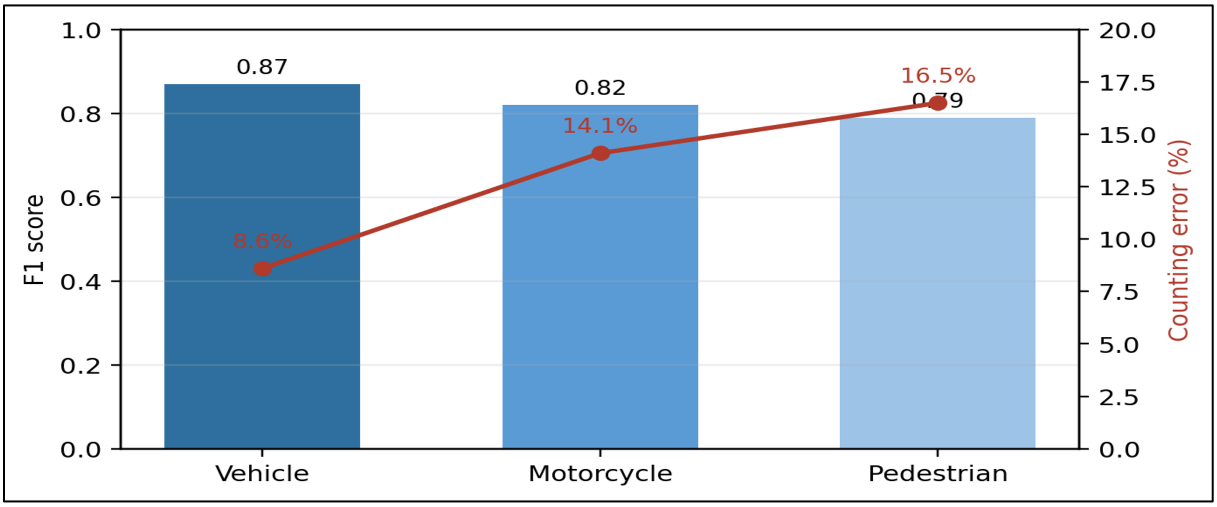

Fig 4 Classwise Detection Performance and Recalculated Counting Errors.

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26aug908

Vehicles were the strongest class. Partial occlusion reduced motorcycle recall, while grouped pedestrians were sometimes merged. An eight-sequence internal ablation suggested that SORT reduced duplicate counts by approximately one-third relative to framewise counting, but the sample was not large enough for a precise interval or a strong comparative claim.

➢ *Queue, Alerts, and Environmental Module*

Queue-length MAE was 1.7 units (median 1; range 0–4), satisfying AC3. Queue-state classification agreed with the operator in 13/15 scenarios (86.7%; Wilson 95%: 62.1–96.3). Eight-unit queues exceeded the useful camera field and were underestimated in 3/5 repetitions. No recommendation was connected to a real signal, so no reduction in urban waiting time is claimed.

All 20 simulated alerts reached the console; 19 arrived within 3 s (95.0%; Wilson 95%: 76.4–99.1), and median latency was 1.4 s, satisfying AC4. The operator accepted 18 events and rejected two false intrusion alarms caused by a shadow and a displaced cone. The 10% false-alert estimate has a wide interval and supports retaining human validation.

The 12 environmental traces showed limited variation on the parking site. A point near parked motorcycles produced a mean relative index 18% above the quietest point. Agreement with a portable reference at three points gave Pearson r=0.76; n=3 is not a calibration study. The module demonstrates mobile carriage of an auxiliary sensor, not regulatory pollution mapping.

➢ *Network-Outage Continuity*

During all eight forced 30 s outages, local obstacle avoidance continued. No logged event was lost and no duplicate was observed after replay, satisfying AC5 within this exact protocol. Remote workstation processing stopped when the link failed. The local system had higher median processing latency (94 ms versus 72 ms) but preserved critical function. This trade-off—not detection superiority—is the central empirical argument for Edge placement.

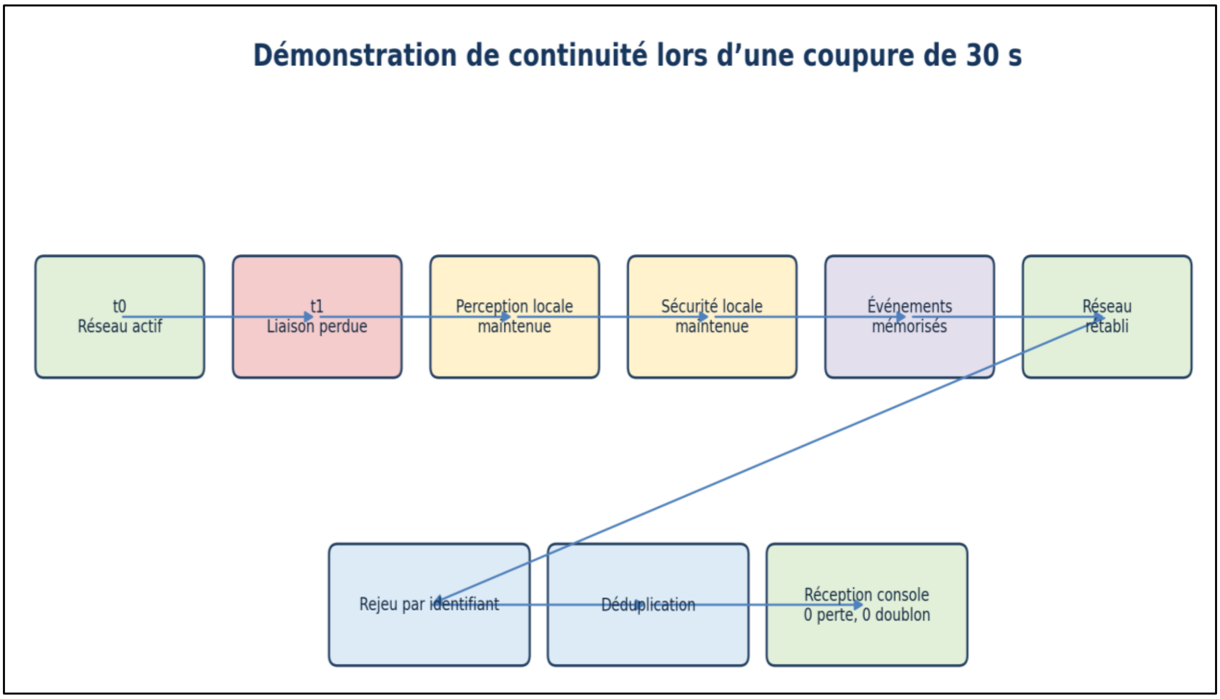


Fig 5 Outage Sequence: Local Continuation, Buffering, Reconnection, Replay, and Deduplication.

## VII. END-TO-END INTEGRATED DEMONSTRATION

➢ *Scenario and Success Definition*

The integrated demonstration linked the same evaluated services in one controlled mission. The console initiated the route; ultrasonic safety governed motion; vision detected and tracked the three classes; calibrated zones produced counts and queue state; an identified event was written locally; the link was disabled for 30 s; safety and logging continued; reconnection triggered replay and deduplication; and the operator accepted or rejected the event. Integrated success required collision-free corridor operation, at least one tracked observation and count, a queue-state output, an identified event, local continuation during outage, single delivery after reconnection, and final human authority.

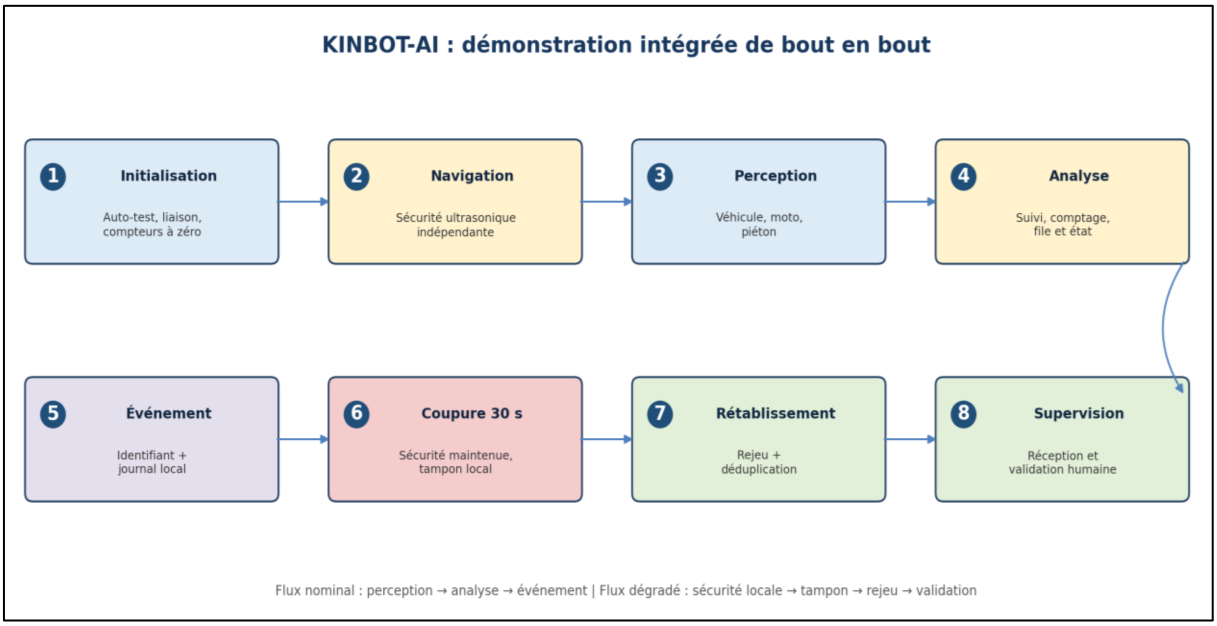

Fig 6 Eight-Stage End-to-End Functional Sequence.

➢ *Evidence Trace*

The end-to-end section synthesizes the modular evidence rather than claiming an additional independent sample: 37/40 nominal navigation missions; F1 values of 0.87/0.82/0.79; counting errors of 8.6/14.1/16.5%; queue MAE of 1.7; 20/20 alert delivery with 1.4 s median latency; and eight outages with zero observed loss or duplication. The accompanying supplement provides an aggregate machine-readable trace and a deterministic metric-recomputation script. It does not contain raw images or a previously unavailable source-code repository.

➢ *Interpretation Boundary*

The integrated trace supports architectural feasibility, not city-scale reliability. The protocol does not cover longer outages, power failure, storage saturation, night, heavy rain, dense crowds, open intersections, or repeated end-to-end success-rate estimation. Future work should use a unified mission log with correlation identifiers across services and a prespecified end-to-end success outcome over independent repetitions.

## VIII. DISCUSSION

KINBOT-AI demonstrates that a low-speed mobile Edge-AI artifact can combine observation, local safety, event integrity, and human supervision in a controlled Kinshasa setting. Compared with fixed Edge traffic systems [8]–[11], mobility adds temporary spatial flexibility but sacrifices continuous coverage. The robot complements rather than replaces fixed infrastructure.

The detector-resource trade-off is visible: 8–11 FPS and 71% mean CPU load were compatible with the speed envelope, but motorcycles and pedestrians remained weaker than vehicles. A heavier model might improve recognition while violating the compute or energy budget. Current 2025 systems using Jetson Orin Nano and YOLOv11 [10] establish stronger detection/tracking baselines; KINBOT-AI's defensible novelty remains system integration and outage continuity, not state-of-the-art vision accuracy.

The human-in-the-loop design is justified by two false alerts in only 20 cases. The air module is deliberately secondary because three comparison points cannot establish metrological validity. Generalization is limited by one closed site, daylight operation, low speed, small samples, sequence-specific calibration, incomplete training metadata, hardware-class rather than exact-component reporting, and the absence of raw public imagery.

## IX. REPRODUCIBILITY, ETHICS AND AVAILABILITY

➢ *Reproducibility Scope*

The article and supplement disclose the known architecture, thresholds, dataset split, test-instance totals, confusion counts, route dimensions, scenario counts, acceptance criteria, and aggregate results. The supplement includes machine-readable configuration, aggregate CSV files, an end-to-end event schema, and a script that recomputes the reported counting errors and F1 values. Because raw code, exact training metadata, exact component identifiers, model weights, and raw images were not available, executable or bitwise replication is not claimed.

➢ *Ethics and Privacy*

Testing was confined to an access-controlled institutional site. Individuals present were informed that experimental recording was taking place. Faces and any visible license plates in retained illustrations were blurred. No alert automatically triggered a physical traffic-control action, and no operational traffic signal or open roadway was controlled. Raw imagery is not publicly released because it may contain identifiable persons or vehicles. No formal ethics-approval identifier was available in the study record; the work is therefore reported as a supervised low-risk engineering demonstration rather than as approved human-subjects

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26aug908

research. Future open-road or identifiable-data studies require prior institutional review, explicit retention/access rules, and applicable legal authorization.

➢ *Data and Code Availability*

Aggregate data and metric-recomputation materials are supplied in the accompanying KINBOT-AI reproducibility supplement. Raw images, model weights, and the original control/inference code are not publicly available. The supplement must not be interpreted as a reconstruction of unavailable source code.

## X. CONCLUSION

KINBOT-AI addresses traffic observation under constrained infrastructure through a mobile Edge-AI architecture that separates vision analytics from independent local safety and retains human authority. Controlled measurements show feasible navigation, three-class observation, queue estimation, alert delivery, and event continuity across eight 30 s outages. The strongest result is not autonomous driving superiority but preservation of local safety and event integrity when remote processing stops. The evidence remains limited to a closed daytime site and incomplete executable reproducibility. Future work should archive exact hardware/software configurations, conduct prespecified repeated end-to-end trials, improve motorcycle and pedestrian performance, and only then consider legally and ethically authorized roadside evaluation.

## REFERENCES


[1]. J. Zhang et al., “Data-driven intelligent transportation systems: A survey,” IEEE Trans. Intell. Transp. Syst., vol. 12, no. 4, pp. 1624–1639, 2011, doi: 10.1109/TITS.2011.2158001.

[2]. K. Nellore and G. P. Hancke, “A survey on urban traffic management system using wireless sensor networks,” Sensors, vol. 16, no. 2, 157, 2016, doi: 10.3390/s16020157.

[3]. Japan International Cooperation Agency, Project for Urban Transport Master Plan in Kinshasa City: Final Report, 2019.

[4]. S. Ayimpam, “La gouvernance au quotidien des taxi-motos à Kinshasa,” Géotransports, nos. 17–18, pp. 135–149, 2023.

[5]. C. McFarlane et al., “First measurements of ambient PM2.5 in Kinshasa and Brazzaville using field-calibrated low-cost sensors,” Aerosol Air Qual. Res., vol. 21, 200619, 2021, doi: 10.4209/aaqr.200619.

[6]. W. Shi et al., “Edge computing: Vision and challenges,” IEEE Internet Things J., vol. 3, no. 5, pp. 637–646, 2016, doi: 10.1109/JIOT.2016.2579198.

[7]. M. Satyanarayanan, “The emergence of edge computing,” Computer, vol. 50, no. 1, pp. 30–39, 2017, doi: 10.1109/MC.2017.9.

[8]. S. Lee et al., “Edge AI-based smart intersection and its application for traffic signal coordination,” J. Adv. Transp., 2024, 8999086, doi: 10.1155/2024/8999086.

[9]. J. Villa et al., “Intelligent infrastructure for traffic monitoring based on deep learning and edge computing,” J. Adv. Transp., 2024, 3679014, doi: 10.1155/2024/3679014.

[10]. F. A. S. Gayas et al., “Edge-based traffic monitoring at intersections: Multi-class vehicle detection, tracking, and counting using YOLOv11 and BoT-SORT,” 2025, doi: 10.3233/ATDE251131.

[11]. C.-H. Li and H.-S. Hou, “Enhancing synchronization of YOLO-based traffic detection on low-end devices by using the COID algorithm,” Proc. Eng. Technol. Innov., 2025, doi: 10.46604/peti.2025.15264.

[12]. J. Redmon and A. Farhadi, “YOLOv3: An incremental improvement,” arXiv:1804.02767, 2018.

[13]. G. Jocher, A. Chaurasia, and J. Qiu, Ultralytics YOLOv8, version 8.0.0, 2023.

[14]. T.-Y. Lin et al., “Microsoft COCO: Common objects in context,” in ECCV, 2014, pp. 740–755.

[15]. S. Thrun, W. Burgard, and D. Fox, Probabilistic Robotics. MIT Press, 2005.

[16]. R. Siegwart, I. R. Nourbakhsh, and D. Scaramuzza, Introduction to Autonomous Mobile Robots, 2nd ed. MIT Press, 2011.

[17]. [17] C. Pang et al., “Adaptive obstacle detection for mobile robots in urban environments,” Sensors, vol. 18, no. 6, 1749, 2018, doi: 10.3390/s18061749.

[18]. [18] A. Bewley et al., “Simple online and realtime tracking,” in ICIP, 2016, pp. 3464–3468, doi: 10.1109/ICIP.2016.7533003.

[19]. [19] L. D. Brown, T. T. Cai, and A. DasGupta, “Interval estimation for a binomial proportion,” Stat. Sci., vol. 16, no. 2, pp. 101–133, 2001.