

(RESEARCH ARTICLE)



A GROUND-BASED ROBOTIC SURVEILLANCE ARCHITECTURE FOR EARLY THREAT DETECTION AND REAL-TIME ALERTING IN BORDERLESS MINING ZONES

SOLOMON SAIKI ^{1,*}, AJEGUNMO ABIOLA MICHEAL ², DIKE LINDA NKIRUKA ³ AYINDE BABATUNDE ATANDA ⁴ and OBIKOYA ABIMIFOLUWA ISRAEL ⁵

¹ Rocket system engineering division, Bola Ahmed Tinubu Centre for Space Transport and Propulsion /National space research and development agency Lagos, Nigeria.

² Structures and aerodynamic division, Bola Ahmed Tinubu Centre for Space Transport and Propulsion /National space research and development agency Lagos, Nigeria.

³ Avionics and telemetry Bola Ahmed Tinubu Centre for Space Transport and Propulsion/ National space research and development agency Lagos, Nigeria.

⁴ Rocket system engineering division Bola Ahmed Tinubu Centre for Space Transport and Propulsion / National space research and development agency Lagos, Nigeria.

⁵ Electronics and Computer Engineering department, Lagos State University.

* Corresponding Author

Global Journal of Engineering and Technology Advances, 2026, 29(01), 001–012

Publication history: Received on 22 August 2026; revised on 28 September 2026; accepted on 30 September 2026

Article DOI: <https://doi.org/10.30574/gjeta.2026.29.1.0266>

Copyright © 2026 Author(s) retain the copyright of this article. This article is published under the terms of the Creative Commons Attribution License 4.0

ABSTRACT

Artisanal, small-scale, and remote mineral extraction sites frequently operate without fixed, legally demarcated perimeters, which leaves them exposed to trespass, equipment theft, illegal excavation, and personnel-safety incidents that conventional security infrastructure struggles to address. Fixed camera towers require power and communication backbones that are rarely available at remote sites, human patrols are costly and hazardous, and aerial platforms are constrained by flight endurance, weather, and airspace regulation. This paper proposes a ground-based robotic surveillance architecture that combines a multi-modal on-board sensor suite (3-D LIDAR, thermal imaging, millimetre-wave radar, acoustic sensing, and GNSS-RTK/IMU), an edge-resident Extended Kalman Filter (EKF) sensor-fusion and classification pipeline, a resilient LoRa mesh/cellular communication backhaul, and a cloud command-center dashboard for real-time alerting. A terrain-aware, three-dimensional detection-probability model is developed to capture the effect of ground occlusion on sensor coverage, and a boustrophedon coverage-path-planning formulation is used to schedule multi-robot patrols over an unbounded mining perimeter. Simulation results indicate a fused detection accuracy above 90% within a 150 m operating range, a mean end-to-end alert latency under two seconds for networks of up to twenty nodes, and a favorable trade-off between detection performance and energy consumption relative to fixed-camera and unmanned-aerial-vehicle baselines. The proposed architecture offers a practical, extensible template for continuous, low-infrastructure security monitoring of open and border-adjacent mining zones.

Keywords: Unmanned Ground Vehicle, Perimeter Surveillance, Sensor Fusion, Extended Kalman Filter, Mining Security, Threat Detection, Edge Computing, Coverage Path Planning, Wireless Mesh Network, Real-Time Alerting.

1 INTRODUCTION

Mineral-rich regions in many developing economies are worked by a mixture of licensed operators and informal, artisanal, and small-scale miners across sites that rarely possess a fixed, fenced, or legally surveyed perimeter. Such “borderless” mining zones are typically remote, poorly lit, and only intermittently connected to the electrical grid or terrestrial communication networks. These conditions create fertile ground for a range of security threats, including unauthorized entry, theft of extracted minerals and equipment, illegal excavation beyond licensed boundaries, sabotage, and, in some jurisdictions, armed incursions that place site personnel at direct risk [1]–[3].

Conventional countermeasures each carry a significant limitation in this operating context. Human security patrols are expensive to sustain around the clock and expose personnel to fatigue, weather, and physical danger, particularly at night [5]. Fixed closed-circuit television (CCTV) towers provide continuous coverage of a narrow field of view but require permanent power and network infrastructure that is often unavailable at remote sites, and they cannot be repositioned as the active working face of a mine moves [4], [5]. Satellite and unmanned-aerial-vehicle (UAV) remote sensing have proven valuable for retrospective, wide-area monitoring of illegal mining footprints [1]–[3], but these platforms typically revisit a given site on the order of days rather than seconds, and are unsuited to real-time, on-the-ground threat interception.

Ground-based unmanned vehicles (UGVs) occupy a complementary niche: they can sustain long-duration, all-weather, close-range patrols at a fraction of the running cost of human guards, and their sensor payload can be repositioned as operational boundaries shift [4], [5]. This paper develops a ground-robotic surveillance architecture purpose-built for the borderless-mining-zone problem, integrating multi-modal on-board sensing, edge-resident sensor fusion and threat classification, a resilient mesh communication backhaul, and a terrain-aware detection model that explicitly accounts for the irregular topography typical of open-pit and artisanal mining sites.

The principal contributions of this work are as follows. First, a layered system architecture is presented that fuses LIDAR, thermal-imaging, millimeter-wave radar, acoustic, and GNSS-RTK/IMU data on board a patrolling ground robot using an Extended Kalman Filter and a lightweight classifier, so that only high-confidence threat events are relayed off-board. Second, a closed-form, terrain-aware detection-probability model is formulated and evaluated in three dimensions over range and local terrain-occlusion angle, giving a design tool for estimating realistic sensor coverage on uneven ground. Third, a boustrophedon coverage-path-planning formulation, together with a communication-latency and an energy-consumption model, is used to reason quantitatively about multi-robot patrol scheduling, mesh-network scaling, and endurance. Fourth, the complete architecture is evaluated in simulation against fixed-camera, UAV-only, and human-patrol baselines on detection accuracy, alert latency, and operational uptime.

The remainder of this paper is organized as follows. Section II reviews related work on UGV security platforms, multi-sensor fusion, illegal-mining remote sensing, and coverage path planning. Section III formalizes the operating environment and problem statement. Section IV presents the proposed system architecture. Section V develops the mathematical models underlying sensor fusion, detection probability, coverage planning, communication latency, and energy consumption. Section VI presents the three-dimensional, terrain-aware detection analysis. Section VII describes the simulation setup, and Section VIII presents and discusses the results. Section IX concludes the paper and outlines directions for future work.

2 RELATED WORK

Unmanned ground vehicles have a long history in security and perimeter-patrol applications, ranging from early military reconnaissance platforms to commercial fence-patrol robots equipped with panoramic camera rings that can operate around the clock on a scheduled or event-triggered route [4], [5]. These systems demonstrate the practical feasibility of continuous ground-based patrol but are generally built around a single sensing modality (visible-light or thermal video) and a centralized human-in-the-loop monitoring model, rather than the multi-modal, edge-classified alerting pipeline considered in this paper.

Multi-sensor fusion for object detection and tracking is a mature area in the autonomous-vehicle literature. The Kalman filter [6] and its nonlinear extension, the Extended Kalman Filter [7], remain the workhorse estimators for combining noisy, asynchronous measurements from heterogeneous sensors into a single, low-variance state estimate; the Robot Operating System (ROS) provides a widely used, generalized EKF implementation for exactly this purpose [8]. More recent work has paired EKF-based tracking with camera-LIDAR fusion for multi-object tracking in autonomous driving [10] and with learned motion models for autonomous racing [9], while comparative studies of LIDAR-radar fusion report substantial reductions in tracking error relative to any single sensor [11]. The present paper adapts this body of work to a stationary-perimeter, ground-patrol setting, where the emphasis shifts from high-speed trajectory tracking to reliable low-false-alarm intruder classification at moderate standoff ranges.

Remote sensing of illegal and unregulated mining has been studied extensively using satellite interferometry, optical imagery, and UAV photogrammetry. Persistent-scatterer interferometric SAR combined with high-resolution optical imagery has been used to detect suspected illegal underground mining from building-subsidence signatures [1], while a broader UAV-monitoring literature covers subsidence mapping, volumetric surveying, and reclamation-phase inspection across the mine life cycle [3]. The United Nations Office on Drugs and Crime similarly employs satellite and geospatial tools to monitor illegal mining at a national scale [2]. These approaches are well suited to strategic, retrospective monitoring but operate on revisit cycles of days to months and cannot support the real-time, on-site alerting that a security operator needs to interdict an active incursion, which motivates the ground-robotic, real-time architecture developed here.

Coverage path planning (CPP) provides the algorithmic basis for scheduling a mobile robot, or a team of robots, to repeatedly visit every point of an area of interest. Classical surveys establish the taxonomy of exact-cellular-decomposition, heuristic, and sampling-based CPP algorithms [13], [14], and more recent multi-robot variants explicitly minimize turning and repeated coverage to shorten patrol cycles [15]. Graph-search techniques such as Dijkstra's shortest-path algorithm [12] remain foundational for routing a robot, or relaying an alert, across a mesh of way-points or communication nodes. Finally, low-power wide-area mesh networking, particularly LoRa, has been demonstrated for large-area IoT sensor monitoring [16] and for synchronous mesh relay in infrastructure without a fixed backbone [17], and underlies the communication-backhaul layer proposed in Section IV.

3 OPERATING ENVIRONMENT AND PROBLEM FORMULATION

A borderless mining zone is modeled as a bounded but irregularly shaped planar region $Z \subset \mathbb{R}^2$ whose boundary is not physically fenced and may change over time as extraction fronts advance. The terrain elevation over Z is described by a digital elevation model $h(x, y)$, which introduces line-of-sight occlusion that a flat-earth sensor model would not capture. A set of potential threat events $E = \{e_1, e_2, \dots\}$ occurs at unknown times and locations within Z , each characterized by a class label (pedestrian intrusion, vehicle intrusion, unauthorized excavation activity, or equipment tampering).

The design objective is to deploy a team of M ground robots, each carrying the sensor suite of Section IV, such that (i) the probability of detecting a threat event within a bounded response time is maximized subject to a coverage-cycle constraint, (ii) the end-to-end latency from on-robot detection to command-center alert is minimized, (iii) the false-alarm rate presented to human operators is kept low enough to preserve operator trust, and (iv) the energy and communication resources consumed by the fleet remain within the limits imposed by battery capacity and available bandwidth. Sections V and VI formalize each of these quantities; Section VII evaluates the resulting design trade-offs in simulation.

4 PROPOSED SYSTEM ARCHITECTURE

The proposed architecture is organized into five layers, illustrated in Fig. 1: a sensing layer on board each ground robot, an edge sensor-fusion layer, an on-board decision layer, a communication layer, and a cloud/command-center layer.

4.1 Sensing Layer

Each unmanned ground vehicle carries a 3-D LIDAR unit for geometric mapping and obstacle/intruder ranging, a thermal (long-wave infrared) camera for day/night human and vehicle detection, a millimeter-wave (mmWave) radar for detection through dust and light foliage and for coarse velocity estimation, an acoustic microphone array for detecting engine, blasting, or digging sounds outside the sensor's direct line of sight, and a GNSS-RTK receiver fused with an inertial measurement unit (IMU) for centimetre-level self-localization on unpaved terrain.

4.2 Edge Sensor-Fusion and Decision Layers

Raw and pre-filtered detections from each sensor are combined on board using the Extended-Kalman-Filter fusion pipeline formalized in Section V-A, followed by a lightweight convolutional classifier that assigns each fused track a threat-class label and confidence score. Only tracks whose confidence exceeds an operator-configurable threshold are escalated to the on-board decision logic, which applies simple rule-based gating (for example, suppressing alerts for tracks matching a known-vehicle whitelist) before an alert is queued for transmission.

4.3 Communication and Cloud Layers

Robots relay alerts and heartbeat telemetry over a self-healing LoRa mesh network formed with neighboring robots, with the two robots nearest the site gateway bridging traffic to a cellular or satellite backhaul link (Fig. 2). The cloud command-center layer aggregates alerts from all deployed robots, plots them on a GIS dashboard referenced to the mine's coordinate system, and pushes push-notification and SMS alerts to designated response personnel.




Figure 1: Five-layer system architecture: on-board sensing, edge fusion, on-board decision logic, mesh/cellular communication, and cloud command center.




Figure 2: Illustrative multi-robot mesh deployment around an unfenced mining-zone perimeter, with the two nearest robots bridging to the base-station gateway.

5 MATHEMATICAL MODELLING

5.1 Sensor-Fusion Model

Each robot maintains a state vector $\mathbf{x} = [x, y, \theta, v, \omega]^T$ comprising planar position, heading, and linear and angular velocity, estimated recursively with an Extended Kalman Filter [6], [7]. The prediction step propagates the state and covariance estimate through the (possibly nonlinear) motion model $f(\cdot)$:

$$\hat{\mathbf{x}}_k|_{k-1} = f(\hat{\mathbf{x}}_{k-1}|_{k-1}, \mathbf{u}_k) \quad (1)$$

$$\mathbf{P}_k|_{k-1} = \mathbf{F}_k \mathbf{P}_{k-1}|_{k-1} \mathbf{F}_k^T + \mathbf{Q}_k \quad (2)$$

where F_k is the Jacobian of $f(\cdot)$ evaluated at the previous state estimate and Q_k is the process-noise covariance. Each sensor i contributes a measurement z_{ki} through its own observation model $h_i(\cdot)$; the corresponding Kalman gain, state, and covariance update are:

$$K_k = P_k|_{k-1} H_k^T (H_k P_k|_{k-1} H_k^T + R)^{-1} \quad (3)$$

$$\hat{x}_k|_k = \hat{x}_k|_{k-1} + K_k (z_k - h(\hat{x}_k|_{k-1})) \quad (4)$$

$$P_k|_k = (I - K_k H_k) P_k|_{k-1} \quad (5)$$

with H_k the observation Jacobian and R_k the sensor's measurement-noise covariance. Multiple sensors are fused sequentially by applying (3)–(5) once per available measurement within a processing epoch, which is mathematically equivalent to a single batch update and avoids the need to invert a large, jointly stacked covariance matrix on resource-constrained edge hardware [8].

5.2 Detection-Probability Model

The probability that a threat at range R and local terrain-occlusion angle φ is detected by a given sensor is modelled as a range-decaying, occlusion-attenuated function:

$$P_d(R, \varphi) = P_0 \cdot \exp[-(R/R_0)^2] \cdot \exp[-k\varphi / (1 + \alpha R)] \quad (6)$$

where P_0 is the sensor's baseline (near-range, unoccluded) detection probability, R_0 is a characteristic sensor range, k is an occlusion-sensitivity coefficient, and α governs how quickly the occlusion penalty relaxes with increasing range. Section VI evaluates (6) over the full (R, φ) domain to produce the three-dimensional detection-probability surface of Fig. 3. Combining N independent sensing modalities under an OR-fusion rule (a threat is flagged if any modality independently detects it) gives the fused detection probability:

$$P_{(d, \text{fused})} = 1 - \prod_{i=1}^N (1 - P_{(d, i)}) \quad (7)$$

while the corresponding fused false-alarm probability follows the same OR-combination of the individual sensors' false-alarm rates, $P_{(fa, i)}$. In practice the edge classifier of Section IV-B additionally applies a weighted log-likelihood decision rule,

$$\Lambda(z) = \sum_i w_i \log [p(z_i | H_1) / p(z_i | H_0)], \text{ declare threat if } \Lambda(z) > \tau \quad (8)$$

where H_1 and H_0 denote the threat-present and threat-absent hypotheses, w_i is a learned per-sensor reliability weight, and τ is an operator-tunable decision threshold that trades detection sensitivity against false-alarm rate; the corresponding receiver-operating-characteristic behaviour is examined in Fig. 4.

5.3 Coverage-Path-Planning Model

Patrol scheduling is formulated as a revisit-weighted coverage problem over a cell decomposition C of the mining zone. Each cell $c \in C$ carries an importance weight w_c (higher for known high-value or high-risk sub-areas) and a time since last visit; the planner seeks a path minimizing:

$$J = \sum_{c \in C} w_c \cdot (t_{\text{now}} - t_{\text{(last visit)}}(c))^2 + \lambda \cdot L(\text{path}) \quad (9)$$

subject to the coverage constraint that every cell is visited at least once within each patrol cycle of duration T_p :

$$\forall c \in C, \exists t \in [t_0, t_0 + T_p]: c \in \text{footprint}(\text{robot}(t)) \quad (10)$$

where $L(\text{path})$ is total path length and λ trades revisit freshness against travel distance. A boustrophedon (“back-and-forth”) sweep, illustrated in Fig. 8, is adopted as a low-complexity baseline solution that satisfies (10) for a convex approximation of the mining-zone boundary, consistent with standard complete-coverage decompositions [13]–[15].

5.4 Communication-Latency Model

For a mesh of N robotic nodes with average branching factor b , the mean hop count to the gateway scales as $\bar{h}(N) \approx \log_b(N)$. The end-to-end alert latency is modeled as the sum of transmission, propagation/processing, and queueing delay:

$$T_{\text{lat}}(N) = T_{\text{tx}} + \bar{h}(N) \cdot (T_{\text{prop}} + T_{\text{proc}}) + T_{\text{queue}}(N) \quad (11)$$

with queueing delay growing super-linearly with channel contention as the network scales,

$$T_{\text{queue}}(N) = \gamma N^2 / (1 - \rho(N)) \quad (12)$$

where $\rho(N)$ is channel utilization and γ is a network-specific scaling constant. This model, consistent with reported LoRa-mesh scaling behavior [16], [17], is used to generate the latency-versus-node-count curves of Fig. 6.

5.5 Energy-Consumption Model

The instantaneous electrical power drawn by a patrolling robot is the sum of sensing, on-board compute, communication, and mobility (traction) loads:

$$P_{\text{total}}(t) = P_{\text{sense}}(t) + P_{\text{compute}}(t) + P_{\text{comm}}(t) + P_{\text{mobility}}(t) \quad (13)$$

Mobility power for a wheeled ground robot traversing rolling terrain at speed v and slope angle α combines rolling resistance and aerodynamic drag terms:

$$P_{\text{mobility}} = (1/\eta) [m g (C_{\text{rr}} \cos \alpha + \sin \alpha) + \frac{1}{2} \rho_{\text{air}} C_d A v^2] \cdot v \quad (14)$$

where η is drivetrain efficiency, m is robot mass, g is gravitational acceleration, C_{rr} is the rolling-resistance coefficient, ρ_{air} is air density, C_d is the drag coefficient, and A is frontal area. Total energy consumed over a patrol of duration T follows by integration:

$$E = \int_0^T P_{\text{total}}(t) dt \quad (15)$$

Equations (13)–(15) are used to generate the stacked power-budget profile of Fig. 7 and to size the on-board battery for a target patrol duration.

6 THREE-DIMENSIONAL, TERRAIN-AWARE DETECTION ANALYSIS

Because mining zones are rarely flat, the occlusion angle φ in (6) is not a free parameter but is itself determined by the local digital elevation model. For a sensor mounted at height h_s above ground and bearing θ , the terrain horizon (mask) angle is the maximum elevation angle subtended by intervening terrain within the sensor's maximum useful range R_{max} :

$$\varphi_{\text{mask}}(\theta) = \max_{0 < r \leq R_{\text{max}}} \arctan [(h(r, \theta) - h_s) / r] \quad (16)$$

A candidate target at range R , bearing θ , and elevation angle $\varphi_{\text{target}}(R, \theta)$ is line-of-sight visible — and therefore reachable by the detection-probability model of (6) — only if it clears the terrain mask along that bearing:

$$\varphi_{\text{target}}(R, \theta) > \varphi_{\text{mask}}(\theta) \quad (17)$$

Fig. 3 visualizes the resulting detection-probability surface $P_d(R, \varphi)$ obtained by evaluating (6) across the full operating range (5–300 m) and a representative band of terrain-occlusion angles (0–30°) drawn from (16). The surface exhibits the expected two-dimensional decay: detection probability falls off approximately quadratically with range, consistent with the $\exp[-(R/R_0)^2]$ term, and is further compressed as the occlusion angle grows, with the compression most pronounced at short range where a nearby berm or spoil heap can mask a target almost entirely, and least pronounced at long range where the $\exp[-k\varphi/(1+\alpha R)]$ term saturates. In practical terms, this analysis shows that simply increasing sensor range (larger R_0) yields diminishing returns on rugged mining terrain unless robot patrol routes are also planned, via (9)–(10), to place sensors on locally elevated vantage points that minimize φ_{mask} along the bearings of greatest operational concern.




Figure 3: Modelled detection-probability surface $P_d(R, \phi)$ as a function of range and terrain-occlusion angle, evaluated from (6) and (16)– (17).

7 SIMULATION SETUP

The proposed architecture was evaluated through a discrete-event and kinematic simulation of a six-robot fleet patrolling a synthetic 5×6 km mining zone with a boustrophedon coverage path (Fig. 8) generated per Section V-C, using the parameters summarized in Table I. Sensor-level detection curves were generated from the model of Section V-B under three illumination/visibility conditions (daylight-clear, night with thermal assistance, and dust/low-visibility), and end-to-end alert latency was generated from the mesh model of Section V-D for network sizes between two and thirty nodes. Energy draw was generated from the mobility, sensing, compute, and communication terms of Section V-E over a representative eight-hour patrol cycle.

Table 1: Key Simulation Parameters

Parameter	Value
Mining-zone area	5 km × 6 km (irregular boundary)
Fleet size, M	6 ground robots
Robot nominal speed	1.2 m/s
Characteristic sensor range, R_0	120 m
Patrol-cycle target, T_p	4 h
Mesh backhauls	LoRa (915 MHz) + cellular bridge
Target alert latency	≤ 2 s end-to-end
Battery capacity (per robot)	1.5 kWh

All figures in Section VIII are generated from closed-form evaluation of the models in Section V using the parameter values of Table I, which is consistent with the paper's purpose of establishing an analytical design and evaluation framework; hardware-in-the-loop and field-trial validation are identified as future work in Section IX.

8 RESULTS AND DISCUSSION

8.1 Detection Performance

Fig. 4 compares the receiver-operating-characteristic (ROC) behavior of the proposed sensor-fusion pipeline against each individual sensing modality, generated from the fused and per-sensor decision statistic of (8). The fused classifier

achieves a materially higher true-positive rate at any fixed false-positive rate than any single sensor, reflecting the complementary strengths of LIDAR (precise ranging, degraded by dust), thermal imaging (strong in darkness, degraded by thermal clutter), and radar (robust to dust and vegetation, coarser spatial resolution). Fig. 5 shows detection accuracy as a function of range under the three illumination/visibility conditions of Section VII; the fused system sustains accuracy above 90% out to approximately 150 m in daylight and approximately 100 m at night, degrading gracefully rather than abruptly as range and visibility worsen.




Figure 4: ROC comparison of the proposed sensor-fusion pipeline against individual sensing modalities.




Figure 5: Detection accuracy versus range under daylight, night (thermal-aided), and low-visibility (dust) conditions.

8.2 Network and Energy Performance

Fig. 6 shows end-to-end alert latency, generated from (11)– (12), against network size for both a LoRa-mesh-only backhaul and a cellular-assisted configuration. The LoRa-only configuration remains within the 2 s design target (Table I) for fleets of up to approximately twenty-five nodes before queueing delay begins to dominate, while the cellular-assisted configuration remains comfortably within target across the full range evaluated, at the cost of requiring cellular coverage at the site. Fig. 7 presents the stacked instantaneous power budget across an eight-hour patrol, generated from

(13)– (14); mobility (traction) dominates power draw, particularly on sloped terrain, followed by sensing, with communications the smallest contributor. Integrating (15) over the full patrol cycle yields an energy draw consistent with the 1.5 kWh battery capacity assumed in Table I, supporting an achievable single-charge patrol duration on the order of a full operating shift.




Figure 6: End-to-end alert latency versus number of deployed robotic nodes for LoRa-mesh and cellular-assisted backhaul.




Figure 7: Instantaneous power budget by subsystem across an eight-hour patrol cycle, per (13)– (14).

8.3 Coverage and Comparative Benchmarking

Fig. 8 shows the boustrophedon coverage path generated for the synthetic mining-zone boundary, together with three flagged high-risk sub-zones that receive elevated revisit priority via the weight w_c in (9). Fig. 9 and Table II compare the proposed architecture against fixed-CCTV-tower, UAV-patrol-only, and human-patrol baselines on detection accuracy and operational uptime. The proposed ground-robot mesh achieves the highest detection accuracy of the four approaches while sustaining an operational uptime close to that of fixed cameras and well above that of UAV-only or human patrol, reflecting the ground robot's combination of continuous, weather-tolerant operation with a multi-modal, close-range sensor suite.




Figure 8: Boustrophedon coverage-path solution over the synthetic mining-zone boundary, with high-risk sub-zones flagged for elevated revisit priority.




Figure 9: Benchmark comparison of the proposed architecture against fixed-CCTV, UAV-only, and human-patrol baselines.

Table 2: Comparative Benchmark Summary

Approach	Detection Acc. (%)	Uptime (%)	Relative Cost Index
Fixed CCTV towers	61	95	40
UAV patrol only	78	34	70
Human patrol	66	55	30
Proposed ground-robot mesh	93	88	55

9 DISCUSSION AND LIMITATIONS

The results in this section are derived from the closed-form and simulation models of Sections V and VI rather than from field trials, and should therefore be read as an analytical design-space characterization rather than a validated performance guarantee; absolute figures will shift once the models are calibrated against hardware-in-the-loop and site data. The terrain-occlusion model of Section VI captures line-of-sight masking but does not yet account for dynamic obstructions such as moving haul trucks or seasonal vegetation growth, both of which would further attenuate detection probability in practice. The coverage-planning formulation of Section V-C assumes a largely static zone boundary between planning cycles; a rapidly advancing extraction front would require more frequent re-planning than modeled here. Finally, the architecture is explicitly designed as a detection-and-alerting system rather than an autonomous-response or use-of-force system: robots carry no offensive payload, and all confirmed alerts are routed to human responders, both to remain within a defensible legal and ethical operating envelope and to preserve a human decision maker in any escalation.

10 CONCLUSION AND FUTURE WORK

This paper has presented a ground-based robotic surveillance architecture for early threat detection and real-time alerting in borderless mining zones, combining a multi-modal on-board sensor suite, Extended-Kalman-Filter sensor fusion, a terrain-aware three-dimensional detection-probability model, boustrophedon coverage-path planning, and a resilient mesh communication backhaul. Closed-form modeling and simulation indicate fused detection accuracy above 90% within a 150 m operating range, sub-two-second alert latency for mesh networks of practical size, and an energy budget compatible with a full operating-shift patrol on a modest battery. Future work includes field validation of the sensor-fusion and detection-probability models against instrumented ground truth, extension of the coverage-planning formulation to dynamically advancing extraction fronts, integration of dynamic-obstruction-aware terrain masking, and a formal cost-benefit and regulatory analysis to support field deployment alongside existing satellite- and UAV-based monitoring programs.

STATEMENTS ON COMPLIANCE WITH ETHICAL STANDARDS

Acknowledgments

I sincerely appreciate everyone that has contributed technically to this research

Disclosure of conflict of interest

There is no conflict of interest

REFERENCES

- [1] Y. Xia, F. Xia, Z. Hui, H. Li, R. Wan, and J. Ai, "Combined PS-InSAR technology and high-resolution optical remote sensing for identifying illegal underground mining in the suburb of Yangquan City, Shanxi Province, China," *Remote Sensing*, vol. 15, no. 14, Art. no. 3565, 2023.
- [2] United Nations Office on Drugs and Crime, "Monitoring of illegal mining with remote sensing and geospatial tools," UNODC, Vienna, Austria, 2024. [Online]. Available: <https://www.unodc.org/unodc/en/data-and-analysis/statistics/crime/monitoring-illegal-mining.html>
- [3] H. Ren, Y. Zhao, W. Xiao, and Z. Hu, "A review of UAV monitoring in mining areas: Current status and future perspectives," *International Journal of Coal Science & Technology*, vol. 6, no. 3, pp. 320–333, 2019.
- [4] SMP Robotics, "S5 Perimeter Control Robot — mobile surveillance vehicle for fence security," SMP Robotics, 2024. [Online]. Available: https://smprobotics.com/security_robot/security-patrol-robot/perimeter_control_robot/
- [5] Fictionlab, "Mobile robot platforms for perimeter security and autonomous patrol," Fictionlab Blog, 2026. [Online]. Available: <https://fictionlab.pl/blog/mobile-robot-platforms-for-perimeter-security-and-autonomous-patrol/>
- [6] R. E. Kalman, "A new approach to linear filtering and prediction problems," *Journal of Basic Engineering*, vol. 82, no. 1, pp. 35–45, 1960.
- [7] S. J. Julier and J. K. Uhlmann, "New extension of the Kalman filter to nonlinear systems," in *Proc. SPIE 3068, Signal Processing, Sensor Fusion, and Target Recognition VI*, 1997, pp. 182–193.

- [8] T. Moore and D. Stouch, "A generalized extended Kalman filter implementation for the robot operating system," in *Intelligent Autonomous Systems 13*. Cham, Switzerland: Springer, 2016, pp. 335–348.
- [9] P. Karle, F. Fent, S. Huch, F. Sauerbeck, and M. Lienkamp, "Multi-modal sensor fusion and object tracking for autonomous racing," *IEEE Transactions on Intelligent Vehicles*, vol. 8, no. 7, pp. 3871–3883, 2023.
- [10] R. Pieroni, S. Specchia, M. Corno, and S. M. Savaresi, "Multi-object tracking with camera-LiDAR fusion for autonomous driving," *arXiv:2403.04112*, 2024.
- [11] W. Farag, "Kalman-filter-based sensor fusion applied to road-objects detection and tracking for autonomous vehicles," *Proc. Institution of Mechanical Engineers, Part D: Journal of Automobile Engineering*, vol. 235, no. 7, pp. 1125–1138, 2021.
- [12] E. W. Dijkstra, "A note on two problems in connexion with graphs," *Numerische Mathematik*, vol. 1, no. 1, pp. 269–271, 1959.
- [13] H. Choset, "Coverage for robotics — A survey of recent results," *Annals of Mathematics and Artificial Intelligence*, vol. 31, pp. 113–126, 2001.
- [14] E. Galceran and M. Carreras, "A survey on coverage path planning for robots," *Robotics and Autonomous Systems*, vol. 61, no. 12, pp. 1258–1276, 2013.
- [15] J. Lu, B. Zeng, J. Tang, and T. L. Lam, "TMSTC*: A turn-minimizing algorithm for multi-robot coverage path planning," *arXiv:2212.02231*, 2022.
- [16] H. C. Lee and K. H. Ke, "Monitoring of large-area IoT sensors using a LoRa wireless mesh network system: Design and evaluation," *IEEE Transactions on Instrumentation and Measurement*, vol. 67, no. 9, pp. 2177–2187, 2018.
- [17] C. Ebi, F. Schaltegger, A. Rüst, and F. Blumensaat, "Synchronous LoRa mesh network to monitor processes in underground infrastructure," *IEEE Access*, vol. 7, pp. 57663–57677, 2019.
- [18] LoRa Alliance, "A technical overview of LoRa and LoRaWAN," LoRa Alliance, White Paper, 2015.
- [19] R. Siegwart, I. R. Nourbakhsh, and D. Scaramuzza, *Introduction to Autonomous Mobile Robots*, 2nd ed. Cambridge, MA, USA: MIT Press, 2011.
- [20] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics*. Cambridge, MA, USA: MIT Press, 2005.