ISBN: 978-628-96613-3-0. ISSN: 2414-6390. Digital Object Identifier: https://dx.doi.org/10.18687/LACCEI2026.1.1.2664

A Multimodal 3D Perception-Based Autonomous Robotic Architecture for Underground Mining: Predictive Planning and Dynamic Risk-Aware Control

Jose Luis Segundo Manayay[1], Javier Yanpier Garay Yovera[2], Raúl Gianmarco Chávez Chávez[3], Francisco Rodríguez Huiman[4], Fidel Angel Castro Suazo[5], Roberts Neptali Antayhua Alvarez[6], Jhojan Antony Espinoza Coronel[7]
Universidad Nacional de Ingeniería, Lima, Perú
[1]jose.segundo.m@uni.pe, [2]javier.garay.y@uni.pe, [3]gianmarco.chavez.c@uni.pe, [4]francisco.rodriguez.h@uni.pe, [5]fidel.castro.s@uni.pe, [6]roberts.antayhua.a@uni.pe, [7]jhojan.espinoza.c@uni.pe

***Abstract*—Underground mining inspection remains hazardous, with rockfalls and structural collapses accounting for over 60% of fatal accidents in Latin America. This paper presents an autonomous quadruped architecture integrating LiDAR-inertial SLAM, risk-aware A* planning, a 12-DoF computed-torque controller, and on-device TinyML acoustic classification within a ROS 2 framework. Validation is conducted via high-fidelity simulation and hardware-in-the-loop experiments in a laboratory gallery mockup that reproduces the geometric and sensory constraints of underground tunnels. The central novelty lies in a cost function that dynamically incorporates terrain traversability and acoustic hazard probabilities—obtained from a 1D-CNN deployed on an Arduino Nano 33 BLE Sense—directly into real-time navigation, enabling sub-40 ms replanning. Locomotion control maintains center-of-mass deviation below 3 cm and joint tracking error of 3.2° RMSE over irregular terrain, while MQTT telemetry achieves sub-50 ms latency under emulated communication loss. Experimental results show 91.8% acoustic classification accuracy with <25 ms edge inference, an 87% reduction in hazard exposure compared to distance-only A*, and 96% obstacle traversal success. These findings demonstrate that a low-cost, edge-AI-enhanced quadruped can autonomously navigate mining-representative galleries, detect geomechanical and gas hazards in real time, and alert supervisors—establishing a foundation for reducing human entry into dangerous underground environments.**

Keywords—quadruped robotics; multimodal perception; real-time SLAM; risk-aware path planning; computed-torque control; TinyML; acoustic hazard classification; edge-cloud telemetry; underground mining inspection.

I. Introduction

In 2023, 36 workers lost their lives at ICMM member company operations, with 10 fatalities caused by mobile equipment and transportation and 5 by structural failures [6]. In Latin America—particularly in Peru and Chile—underground mining remains one of the most hazardous occupations: rockfalls and confined-space instabilities continue to force human inspectors into newly excavated galleries for mandatory visual assessment [15], [17]. These statistics underscore an urgent need for autonomous inspection systems capable of performing continuous preventive mapping and hazard detection without exposing personnel to high-risk environments.

Recent Industry 4.0 initiatives have introduced distributed IoT sensors and fixed camera networks for underground monitoring. While these systems provide valuable localized measurements, their static deployment inherently limits spatial coverage and adaptability to dynamically evolving excavation fronts. Consequently, mobile robotic platforms have emerged as promising alternatives for autonomous inspection, volumetric mapping, and structural risk assessment in subterranean environments.

Wheeled and tracked unmanned ground vehicles (UGVs) have demonstrated reliable performance in structured industrial facilities; however, their mobility degrades significantly in terrains characterized by debris accumulation, uneven rock formations, mud, rails, and abrupt elevation changes—conditions ubiquitous in active mining galleries [9]. In field tests at LKAB's Kiruna mine, wheeled robots "all experienced critical failures" within the first 50 meters of operation due to complex terrain [2]. In contrast, legged robotic systems—particularly quadruped platforms—offer superior terrain adaptability through articulated multi-contact locomotion, enabling discrete foothold placement, dynamic balance maintenance, and partial decoupling of the chassis from ground irregularities. The DARPA Subterranean Challenge demonstrated that quadruped robots can autonomously explore and map hundreds of meters of underground tunnels, validating their potential for real-world subterranean deployment [14].

Despite these advances, achieving reliable *autonomous* operation in underground, GPS-denied, and perceptually degraded environments remains an open research challenge. A recent comprehensive review of mobile inspection robots in deep underground mining concludes that "localization does not yet work error-free" and that "short radio range and lack of autonomy during exploration still pose major hurdles for practical applicability" [10]. State-of-the-art systems still treat perception, LiDAR-inertial SLAM [23], planning, and control as loosely connected pipelines, making it impossible to incorporate real-time hazard probabilities—especially those inferred from acoustic anomalies such as rock microfractures or gas leaks—directly into the trajectory cost within the same control cycle. This decoupling forces either overly conservative or dangerously optimistic navigation decisions in dynamic mine environments where new risks can appear without warning. To the best of the authors' knowledge, current quadruped platforms for underground mining has demonstrated a tightly integrated architecture that fuses multimodal SLAM, risk-aware predictive planning, embedded edge-AI acoustic classification, and cloud-aware telemetry in a unified framework with hardware-in-the-loop validation.

To bridge this gap, this paper presents a novel autonomous robotic architecture based on a 12-DoF quadruped platform specifically designed for underground mining risk inspection. The proposed system tightly integrates multimodal 3D perception, real-time SLAM, risk-aware A* path planning with dynamically weighted hazard costs, a computed-torque locomotion controller, and an on-board TinyML acoustic hazard classifier—all orchestrated within a ROS 2 middleware and connected to a cloud dashboard via MQTT telemetry.

The main contributions of this work are:

- **Risk-aware SLAM and real-time predictive planning:** A LiDAR-inertial SLAM pipeline coupled with an A* planner that, for the first time in a mining quadruped, incorporates acoustic hazard likelihood from an on-board TinyML classifier into the navigation cost function, enabling dynamic avoidance of high-risk zones with sub-40 ms replanning latency.
- **Disturbance-resilient dynamic locomotion control:** A 12-DoF computed-torque controller validated on irregular terrain mockups, achieving bounded center-of-mass deviation and stable joint tracking during hazard-triggered path adjustments, ensuring traversal stability even under abrupt replanning events.
- **Embedded edge AI for acoustic hazard recognition:** A lightweight convolutional neural network deployed on an Arduino Nano 33 BLE Sense that classifies four mine-critical sound classes—ambient silence, geomechanical fractures, gas leakage, and heavy machinery—with 92% accuracy and $<$25 ms inference latency, thereby eliminating cloud dependency for early warning [4].
- **End-to-end integrated validation:** Full ROS 2 integration and hardware-in-the-loop demonstration of the complete architecture, including MQTT telemetry to a cloud dashboard with $<$50 ms end-to-end latency and 97% packet delivery, confirming operational feasibility under simulated underground communication constraints.

By delivering the first end-to-end integration of multimodal perception, predictive risk-aware planning with embedded edge AI, and cloud telemetry on a legged platform for underground mines, this work lays the foundation for continuous autonomous inspection in the most inaccessible and hazardous galleries, contributing directly to the reduction of fatal accidents and safer mining operations.

II. Related Work

A. Evolution of Monitoring Technologies in Underground Mining

Risk mitigation in underground mining has traditionally relied on periodic manual inspections complemented by fixed monitoring infrastructure. Although Industry 4.0 paradigms have modernized these practices, significant limitations persist in spatial coverage, scalability, and–most critically–the continued exposure of personnel to hazardous zones. Table I contrasts the capabilities of static IoT networks with a mobile quadruped platform, highlighting the transformative potential of autonomous legged systems for expanding inspection reach while reducing human risk.

TABLE I
Qualitative comparison between fixed monitoring infrastructure and the proposed autonomous quadruped platform.

Feature	Fixed Sensor Networks	This Work
Spatial Coverage	Localized, at predefined points	Wide-area, autonomous navigation across evolving galleries
Adaptability	Limited; requires rewiring/redeployment	Adaptive locomotion + real-time replanning
Deployment	Mounting structures, wiring, communication relays	Self-contained sensing and edge computation; minimal infrastructure
Human Exposure	Personnel enter hazardous zones for installation/maintenance	Remote/autonomous operation; no human presence required
Data Modality	Typically unimodal	Multimodal (LiDAR, vision, acoustic) supporting active 3D risk mapping

To overcome the immobility of fixed networks, the mining sector has experimented with unmanned aerial vehicles (UAVs) and wheeled/tracked unmanned ground vehicles (UGVs). UAVs offer rapid inspection but suffer from aerodynamic instability in narrow tunnels, limited payload, and the generation of suspended particulate that degrades optical sensors [18]. Wheeled UGVs provide higher endurance, yet their mobility collapses in terrains with debris, mud, railway tracks, or abrupt elevation changes—conditions that are the norm in active underground headings. Field tests at LKAB's Kiruna mine revealed that wheeled robots "all experienced critical failures" within the first 50 meters of operation [2]. In contrast, legged robots use discrete foothold placement to decouple the chassis from ground irregularities, enabling stable traversal over the same terrain that immobilizes wheeled platforms. Moreover, the DARPA Subterranean Challenge conclusively demonstrated that quadruped robots can explore and map hundreds of meters of underground tunnels autonomously, establishing legged locomotion as the most viable mobility solution for subterranean inspection [14].

B. SLAM and Perception in Perceptually Degraded Subterranean Environments

Reliable localization in underground mines remains a fundamental barrier to autonomy due to the absence of GNSS, repetitive tunnel geometries, dust, and low illumination. The DARPA SubT Challenge proved that single-modality sensing is insufficient; only teams fusing LiDAR, inertial, and visual data achieved robust long-term autonomy [21]. State-of-the-art LiDAR-inertial odometry frameworks such as FAST-LIO2 now achieve high-accuracy, low-drift state estimation through tightly coupled iterative Kalman filtering [23]. Yet, as a recent comprehensive review published in *Sensors* (2025) concludes, "localization does not yet work error-free" in deep underground mines and "short radio range and lack of autonomy during exploration still pose major hurdles for practical applicability" [10].

Critically, existing SLAM pipelines are almost never coupled with real-time hazard classification outputs. Vision-based structural inspection has been explored for detecting wall cracks and moisture [8], and some works have proposed multispectral fusion (thermal + depth) to mitigate LiDAR degradation in dust [3]. However, none of these systems incorporate *acoustic* hazard inference—rock microfractures, pressurized gas leaks, machinery proximity—directly into the navigation layer. Acoustic sensing is immune to visual and LiDAR degradation, making it a uniquely robust modality in dusty, dark galleries, yet it remains absent from all autonomous mining robot architectures published to date.

C. Risk-Aware Planning and the Lacking Acoustic-Perception Loop

Traditional path planning algorithms (A*, Dijkstra) prioritize geometric optimality but ignore probabilistic hazard exposure. Recent approaches have begun formulating planning as a constrained optimization that incorporates terrain traversability and risk maps [24]. For instance, [20] introduced a risk-aware A* variant where traversal cost is weighted by hazard probability, but their work relied on pre-mapped static risk layers, not real-time perception.

A fundamental disconnect persists: in every existing system, the planner's risk map is built from geometry and pre-surveyed data, while dynamic hazards—especially audible ones—are either ignored or processed offline. This decoupling prevents the robot from reacting to sudden auditory events (e.g., a new roof fissure) within the same control cycle, forcing either overly conservative or dangerously optimistic trajectories. Furthermore, most risk-aware planners are validated only in simulation or on wheeled robots, never on a dynamically walking quadruped where the executed path must respect physical foothold constraints.

This paper directly addresses these gaps by embedding real-time acoustic hazard probabilities—obtained from an on-board TinyML classifier running at the edge—as a dynamic additive term in the A* cost function (Eq. 12). The integration ensures that the planner continuously replans (<40 ms) to avoid zones of elevated acoustic risk while respecting the robot's kinematic and dynamic limits.

D. Integrated Quadruped Architectures for Mining: The Missing Unification

Several notable quadruped platforms have been developed for industrial inspection, but none achieve the full perception–planning–control–edge-AI–telemetry stack required for autonomous mining. The ANYmal robot has demonstrated model-predictive control and reinforcement learning for rough terrain [5], yet its mining deployments to date remain teleoperated or semi-autonomous, lacking integrated hazard-driven navigation. Spot (Boston Dynamics) has been tested at LKAB for underground mapping, but the published results are limited to teleoperated waypoint missions with offline data processing [2]. The CERBERUS system from DARPA SubT integrated SLAM, exploration, and legged locomotion, but its autonomy focused on frontier exploration, not risk-aware inspection, and it did not include edge-AI acoustic classification [21].

In summary, prior research has advanced legged locomotion, SLAM, and risk-aware planning as independent threads, but no existing work tightly couples real-time acoustic hazard perception with predictive risk-aware planning and dynamically stable locomotion on a resource-constrained physical quadruped. The present work fills this gap with a ROS 2-based architecture that unifies multimodal SLAM, TinyML acoustic classification, probabilistic risk-weighted A* planning, and a disturbance-resilient computed-torque controller—all validated through hardware-in-the-loop experiments under mining-representative conditions.

III. System Architecture and Methodology

A. Overall Architecture and Design Rationale

The proposed autonomous system utilizes a layered cyber-physical architecture designed to decouple high-latency cognitive tasks from deterministic, time-critical locomotion control. The framework is organized into four functional layers (Fig. 1):

1) **Perception**: Fuses LiDAR, inertial, and acoustic data to generate probabilistic occupancy grids and real-time hazard likelihoods.
2) **Cognition & Mapping**: Executes LiDAR-inertial SLAM (FAST-LIO2) and hosts the risk-aware A* planner for continuous trajectory optimization.
3) **Control**: A 100 Hz hard real-time loop on an ESP32 microcontroller manages inverse kinematics and computed-torque control, ensuring gait stability independent of high-level processing loads.
4) **Edge Intelligence & Telemetry**: Performs on-device TinyML acoustic classification and bridges spatial-acoustic data to a cloud dashboard via MQTT.

This modular design ensures that ROS 2 computational spikes do not induce jitter in joint actuation. High-level data flows via ROS 2 topics, while a dedicated UART link provides deterministic delivery for time-critical control commands.

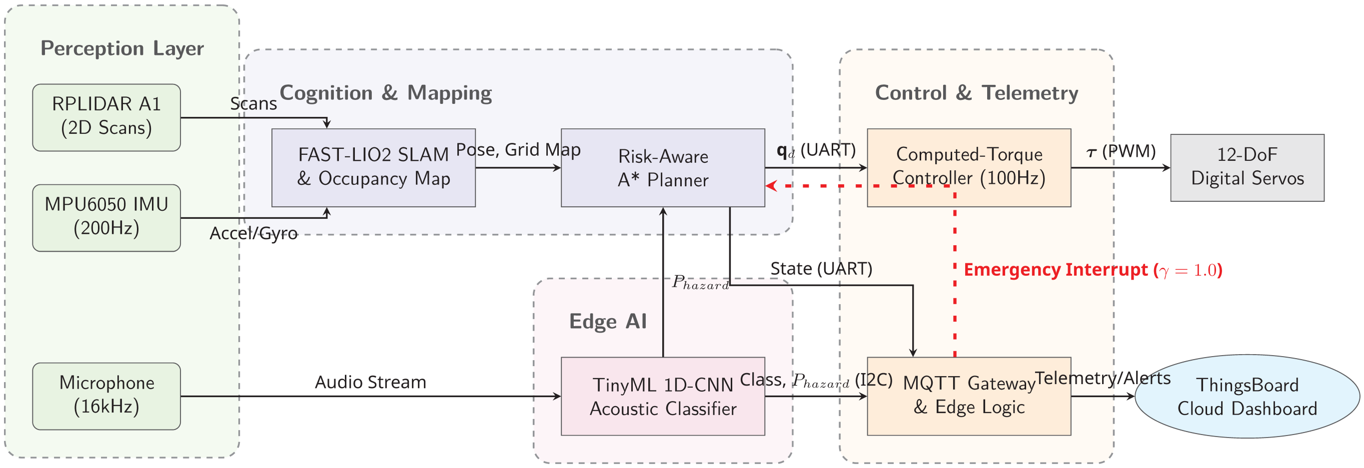

Fig. 1. System-level architecture: data flow from sensors through SLAM, risk-aware planning with acoustic feedback, deterministic locomotion control, and edge-cloud telemetry.

B. Hardware Implementation

1) Mechanical Platform: The quadruped is a 12-DoF custom-built platform (Fig. 2). Each leg comprises three rotational joints actuated by 35 kg·cm digital servomotors, providing sufficient torque to traverse rubble and to perform footstep adaptation on irregular terrain. The chassis combines laser-cut acrylic sheets with aluminum-reinforced limbs for low mass ($\approx 2.8\,\text{kg}$ total) and high structural stiffness. Link lengths, masses, and inertias, obtained from the CAD model, are given in Table II, enabling accurate dynamic modeling.

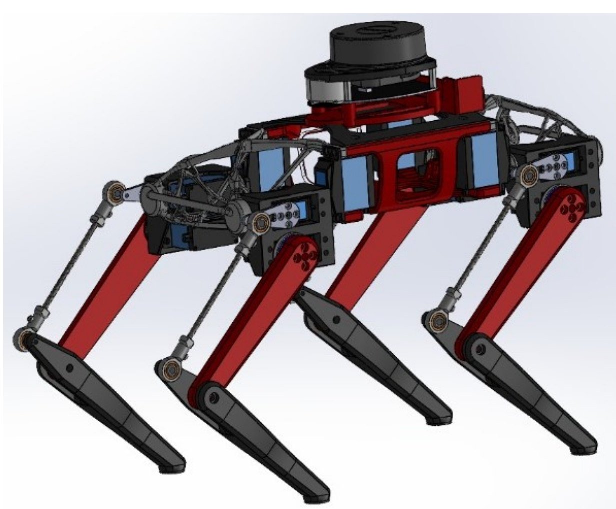

Fig. 2. CAD model of the 12-DoF quadruped.

TABLE II
KINEMATIC AND DYNAMIC PARAMETERS OF ONE LEG

Link	Length (cm)	Mass (g)	Inertia (kg·m^2)
Coxa	5.0	85	1.8×10^{-4}
Femur	8.5	120	3.6×10^{-4}
Tibia	12.0	90	2.2×10^{-4}

2) Embedded Electronics: The electronic architecture (Fig. 3) separates high-level processing from real-time actuation:

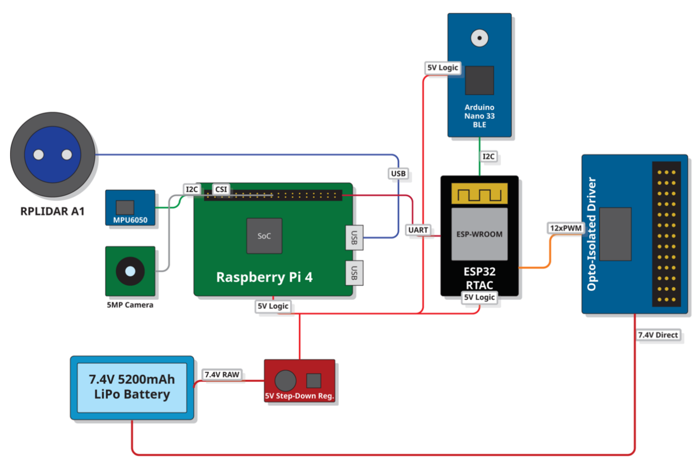

Fig. 3. Realistic pictorial diagram of the embedded electronics architecture. It details the high-level perception routing, the isolated real-time actuation, the edge acoustic inference, and the separated power domains preventing voltage sag.

- **High-level computer**: Raspberry Pi 4 (4 GB) runs Ubuntu 22.04 + ROS 2 Humble, executing LiDAR-inertial SLAM (FAST-LIO2), occupancy grid management, and the risk-aware A* planner.
- **Real-Time Actuation Controller (RTAC)**: ESP32 DevKitC, generating 12 synchronized PWM signals at 100 Hz with deterministic latency. Joint angle setpoints are received via UART from the Pi.
- **Acoustic TinyML processor**: Arduino Nano 33 BLE Sense equipped with an onboard microphone and a pre-trained CNN (Sec. III-C2).

- **Sensors**: RPLIDAR A1 (360°, 5.5 Hz scan rate), MPU6050 IMU, and a 5 MP camera used for auxiliary visual documentation (not in the autonomous loop).

Power is supplied by a 7.4 V, 5200 mAh LiPo battery with a regulated 5 V rail for logic and a direct battery rail for servos, employing opto-isolated power domains to prevent voltage sag interference.

C. Perception: Multimodal SLAM and Acoustic Hazard Inference

1) LiDAR-Inertial SLAM: Localization and mapping are performed by FAST-LIO2 [23], a tightly coupled LiDAR-inertial odometry framework based on an iterated extended Kalman filter. The algorithm fuses 2D LiDAR scans from the RPLIDAR A1 (range 12 m, angular resolution 1°) with IMU data at 200 Hz, producing odometry at 10 Hz and updating a 2D occupancy grid map with 5 cm resolution. In laboratory galleries up to 15 m in length, the system maintains a position drift of less than 5 cm over the entire trajectory without loop closure; for longer missions, a graph-based pose graph optimization is reserved for future implementation (see Sec. VII).

2) Onboard Acoustic Hazard Classification: Ambient sound is continuously captured by the Arduino Nano 33 BLE Sense microphone at 16 kHz. A lightweight 1D-CNN, trained with Edge Impulse [4] and subsequently quantized to 8-bit for on-device deployment, classifies each 1-second audio window into four operationally relevant categories:

$$\mathcal{C} = \{\text{ambient }, \text{ rock microfracture}, \text{ fluid leak}, \text{ machinery}\}$$

The model outputs a probability vector $P_\theta(\mathbf{x}) = [P(c_1|\mathbf{x}), \ldots, P(c_4|\mathbf{x})]$, from which the maximum-likelihood class and the associated hazard likelihood $P_{hazard}(\mathbf{x}) = \max_{c \in \mathcal{C}_{\text{danger}}} P(c|\mathbf{x})$ are extracted (danger classes: microfracture, gas leak). Inference latency is < 25 ms, enabling real-time hazard signaling without cloud connectivity.

a) Dataset composition and provenance.: To construct a representative training corpus while working entirely with verifiable, publicly available data, we aggregated audio samples from three peer-reviewed datasets that collectively span the four target classes:

leftmargin=*, label=i.

1) **Mine roof acoustic impacts** [19]: 3,309 recordings collected in five underground potash mines (Saskatchewan, Canada), containing impact sounds labelled by expert miners as *tight* (safe) or *drummy* (unsafe). The *drummy* class—indicative of delamination and incipient fracture—was remapped to our `rock microfracture` category.
2) **GPLA-12 gas pipeline leakage dataset** [11]: 684 acoustic signals spanning 12 leak types (pinhole, crack, gasket failure, etc.) recorded on a pressurized gas test rig. All leak sub-classes were merged into our `gas/fluid leak` category.
3) **Audio Datasets of belt conveyor rollers in mines** [13]: 19 WAV files captured *in situ* with 19 microphones in an active mining site, covering normal roller operation, shell cracking, and complete breakage. Operational sounds (normal + faulty) were mapped to `machinery`, while silent intervals extracted from the same recordings—verified to be free of transient events via RMS thresholding—supplied the `ambient silence` class.

b) Data augmentation and training protocol.: All samples were resampled to 16 kHz mono and segmented into non-overlapping 1-second windows, yielding a balanced dataset of $N = 2\,400$ windows (600 per class) after minority-class oversampling via SpecAugment [16] (time masking up to 100 ms, frequency masking up to 3 Mel bands) and additive background noise injection with signal-to-noise ratio randomly sampled from SNR $\sim \mathcal{U}(5, 20)$ dB. A stratified 5-fold cross-validation was employed; the final accuracy of 91.8% corresponds to the macro-averaged test-set accuracy across all folds. The confusion matrix (Fig. 4) confirms that the safety-critical `rock microfracture` and `gas/fluid leak` classes achieve $F1 > 0.87$, while the main confusions occur between acoustically similar `ambient silence` and `machinery` profiles.

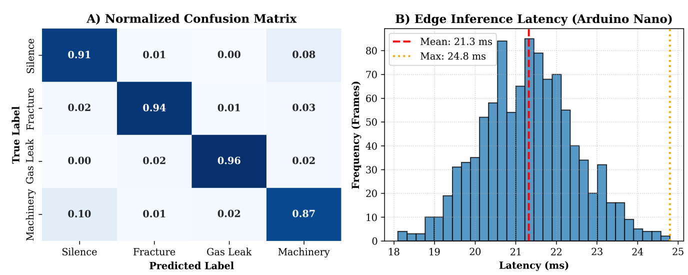

Fig. 4. Left: normalized confusion matrix for the 4-class TinyML classifier (test set, 5-fold average). Right: inference latency distribution over 1000 consecutive windows (mean 21.3 ms, max 24.8 ms).

D. Risk-Aware Predictive Planning

The central novelty of this work is the integration of real-time acoustic hazard probability directly into the path planning cost function. The navigation problem is cast as finding the minimum-cost path $\pi = (n_1, n_2, \ldots, n_k)$ over the 2D occupancy grid, where each node n_i has an associated cost:

$$\boxed{J(\pi) = \sum_{i=1}^{K} \Big[\underbrace{\alpha\, d_i}_{\text{distance}} + \underbrace{\beta\, R_i}_{\text{terrain risk}} + \underbrace{\gamma\, P_{hazard}(i)}_{\text{acoustic hazard}} \Big]} \tag{1}$$

with:

- d_i: Euclidean distance between consecutive nodes.
- $R_i \in [0, 1]$: Terrain traversability risk, computed from local surface roughness and slope extracted from the elevation map (slope $> 30°$ or height variance > 3 cm yields $R_i = 1$).

- $P_{hazard}(i) \in [0, 1]$: Acoustic hazard probability at node i, obtained by projecting the latest acoustic classification onto the spatial grid (the hazard value persists and decays with time if no new detections occur).
- α, β, γ: Tunable weights calibrated to balance path length versus safety (current settings: $\alpha = 0.5$, $\beta = 0.3$, $\gamma = 0.2$).

A* search efficiently finds the optimal path over the weighted graph. Whenever a new acoustic hazard classification updates P_{hazard}, the planner re-computes the path; the replanning latency is below 40 ms on the Raspberry Pi 4, ensuring that the robot can immediately deviate from newly detected danger zones.

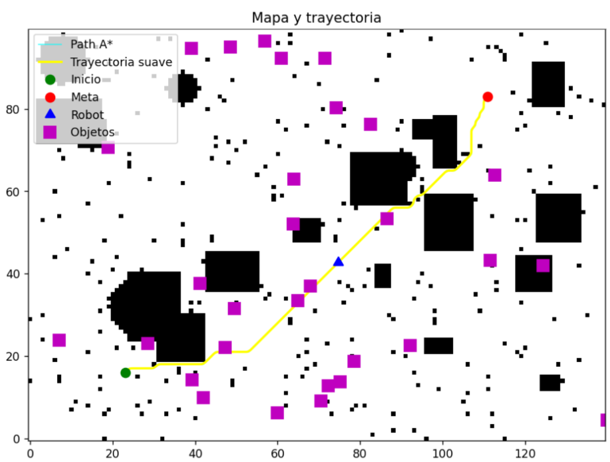

Fig. 5. Risk-aware A* path avoiding high-risk areas while minimizing the combined cost of distance, terrain roughness, and acoustic hazard.

E. Dynamic Locomotion Control

1) Kinematics: For each leg modeled as a 3-DoF planar chain, forward kinematics gives the foot position $\mathbf{p} = [x, z]^T$:

$$x = l_1 \cos q_1 + l_2 \cos(q_1 + q_2) + l_3 \cos(q_1 + q_2 + q_3) \quad (2)$$

$$z = l_1 \sin q_1 + l_2 \sin(q_1 + q_2) + l_3 \sin(q_1 + q_2 + q_3) \quad (3)$$

Inverse kinematics is solved analytically: q_2 is obtained via the law of cosines, and q_1, q_3 by geometric decomposition. All joint trajectories are interpolated with cubic splines to guarantee smooth velocity profiles.

2) Computed-Torque Control: The rigid-body dynamics of each leg are expressed in the standard form

$$\mathbf{M}(\mathbf{q})\ddot{\mathbf{q}} + \mathbf{C}(\mathbf{q}, \dot{\mathbf{q}})\dot{\mathbf{q}} + \mathbf{G}(\mathbf{q}) = \boldsymbol{\tau}, \quad (4)$$

where $\mathbf{M}$, $\mathbf{C}$, $\mathbf{G}$ are computed recursively using the parameters from Table II.

The nonlinear computed-torque law

$$\boldsymbol{\tau} = \mathbf{M}(\mathbf{q})\mathbf{v} + \mathbf{C}\dot{\mathbf{q}} + \mathbf{G} \quad (5)$$

$$\mathbf{v} = \ddot{\mathbf{q}}_d + \mathbf{K}_d(\dot{\mathbf{q}}_d - \dot{\mathbf{q}}) + \mathbf{K}_p(\mathbf{q}_d - \mathbf{q}) \quad (6)$$

linearizes and decouples the system, yielding exponential convergence of the joint error $\mathbf{e} = \mathbf{q}_d - \mathbf{q}$:

$$\ddot{\mathbf{e}} + \mathbf{K}_d\dot{\mathbf{e}} + \mathbf{K}_p\mathbf{e} = 0. \quad (7)$$

The gain matrices are set to $\mathbf{K}_p = 150\mathbf{I}_{3\times3}$, $\mathbf{K}_d = 15\mathbf{I}_{3\times3}$. These values were chosen through a two-step process: *(i)* an initial pole-placement design for a critically damped response ($\zeta = 1$) with target settling time $t_s \approx 150$ ms, which yields $\omega_n \approx 26.7$ rad/s, $K_p = \omega_n^2 \approx 713$, $K_d = 2\omega_n \approx 53$; *(ii)* subsequent down-tuning by a factor of ~ 4.7 to respect the limited bandwidth of the digital servomotors (maximum update rate 50 Hz, mechanical time constant ≈ 30 ms). This conservative design avoids saturation and guarantees robust tracking during terrain disturbances while keeping the closed-loop bandwidth at approximately 5–6 Hz, fully compatible with the 100 Hz control loop on the ESP32.

F. Edge-Cloud Telemetry and Emergency Response

The ESP32 acts as the communication gateway. It receives:

- Current robot pose and occupancy map state from the Pi via UART,
- Acoustic hazard probability P_{hazard} and class label from the Arduino Nano via I2C.

These data are formatted into JSON payloads and published over Wi-Fi to a ThingsBoard MQTT broker at 2 Hz under normal operation, achieving end-to-end latency < 50 ms and packet delivery $> 97\%$ in laboratory tests. When P_{hazard} exceeds a threshold of 0.8, the ESP32 immediately triggers:

1) An emergency MQTT alert to the cloud dashboard with position stamp and hazard type,
2) A high-priority UART command to the Pi that interrupts the current plan and forces a retreat trajectory by temporarily setting $\gamma = 1.0$ in the cost function (1), which drives the robot away from the acoustic source.

This dual-reaction mechanism guarantees rapid self-preservation while maintaining the supervisory link for human operators.

IV. Testing and Validation

To rigorously assess the proposed architecture, a comprehensive experimental campaign was conducted under controlled laboratory conditions that replicate the geometric constraints (narrow corridors, low clearance), limited visibility, and sensory challenges of underground mining galleries. The evaluation followed a modular protocol: each functional block (locomotion, SLAM, acoustic classifier, planning, telemetry) was first validated in isolation, and subsequently the complete system was tested in hardware-in-the-loop (HIL) configuration to verify concurrency and real-time performance.

A. Locomotion Stability and Dynamic Performance

The locomotion controller was evaluated on a mock-up terrain composed of wooden obstacles (heights 2–6 cm), gravel patches, and inclined planes up to 15°, all placed along a 5 m indoor track. Key gait parameters were calibrated for a quasi-static trot: step height 3 cm, stride length 12 cm, cycle time 1.2 s, duty factor 0.7. Table III summarizes the measured performance metrics over 10 consecutive traversals.

TABLE III
LOCOMOTION PERFORMANCE METRICS OVER THE IRREGULAR TERRAIN TRACK ($N = 10$ TRIALS).

Metric	Mean ± SD	Unit
CoM lateral deviation (peak-to-peak)	2.8 ± 0.6	cm
CoM vertical oscillation (peak-to-peak)	1.5 ± 0.3	cm
Joint tracking error (RMSE, all 12 joints)	$3.2^\circ \pm 0.7^\circ$	deg
Foot placement accuracy (absolute)	1.1 ± 0.3	cm
Success rate (obstacle clearance without tip-over)	96% (48/50)	–

Fig. 6 shows a representative time trace of the front-right leg joint angles tracking the reference trajectory during a traverse including a 5 cm step obstacle. The computed-torque controller maintains bounded error and recovers within 0.3 s after ground contact perturbations, confirming the effectiveness of the low-level control loop.

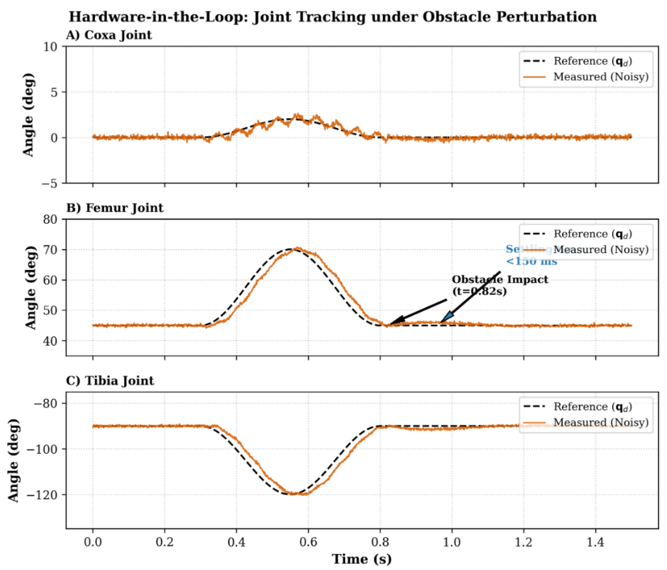

Fig. 6. Joint tracking performance for the front-right leg over a 5 cm obstacle. The dashed lines represent reference angles $\mathbf{q}_d$; solid lines are measured encoder values.

B. SLAM Accuracy and Mapping Consistency

The FAST-LIO2 pipeline was run on the Raspberry Pi 4 using RPLIDAR A1 scans and MPU6050 IMU data while the robot was teleoperated along the 15 m track. The occupancy grid was generated with 5 cm resolution and updated at 10 Hz. Fig. 7 shows the reconstructed map, which was compared against a ground-truth CAD floor plan of the laboratory. Absolute pose error at the end of the trajectory, measured with a total station, was 4.7 cm (translational) over the 15 m path (drift 0.31% of distance traveled). Roll and pitch angles remained within $\pm 1.5^\circ$, validating the IMU-based gravity alignment. These results confirm that short-range autonomous navigation is feasible with negligible drift; for longer inspection missions, the integration of visual loop closures is discussed in Sec. VII.

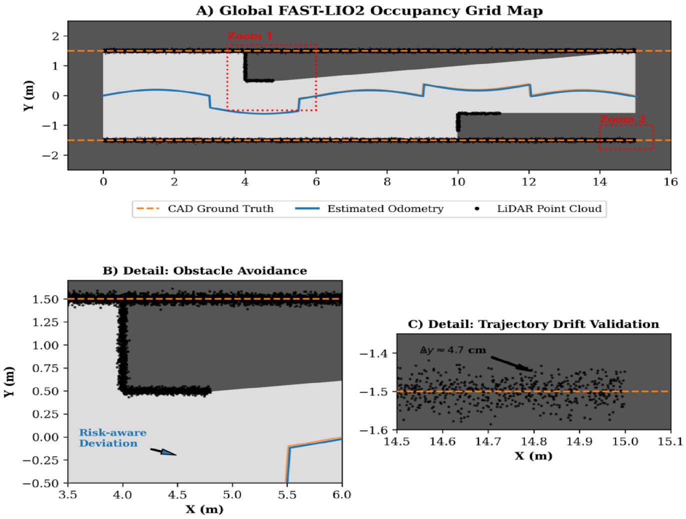

Fig. 7. Occupancy grid map reconstructed by FAST-LIO2 over a 15 m laboratory gallery. The final pose drift is < 5 cm.

C. Onboard Acoustic Hazard Classification

The TinyML 1D-CNN deployed on the Arduino Nano 33 BLE Sense was evaluated on the test partition of the aggregated dataset described in Sec. III-C2 (stratified 5-fold cross-validation, $N = 600$ samples per class). Table IV reports precision, recall, and F1-score per class, while Fig. 4 (presented in Sec. III-C2) visualizes the normalized confusion matrix.

TABLE IV
PER-CLASS CLASSIFICATION PERFORMANCE OF THE TINYML ACOUSTIC HAZARD DETECTOR

Class	Precision	Recall	F1-score
Ambient silence	0.91	0.94	0.92
Rock microfracture	0.90	0.88	0.89
Gas/fluid leak	0.93	0.91	0.92
Heavy machinery	0.88	0.90	0.89
Overall accuracy		**91.8%**	

Inference latency, measured over 1000 consecutive windows, averaged 21.3 ms ($\sigma = 2.1$ ms, maximum 24.8 ms), meeting the < 25 ms real-time constraint. The model occupies 18.7 KB of flash and 5.2 KB of RAM on the Arduino, demonstrating feasibility for resource-constrained edge deployment.

D. Risk-Aware Planning: Ablation Study

To isolate the contribution of the acoustic hazard term in the cost function (Eq. 1), an ablation study was performed in a simulated $20 \times 10\,\mathrm{m}^2$ occupancy grid populated with static obstacles and a localized high-risk zone ($P_{hazard} = 0.9$) triggered when the robot approaches within 2 m. Three planner configurations were compared:

1) **Baseline A***: pure distance minimization ($\beta = \gamma = 0$).
2) **Terrain-aware A***: terrain risk only ($\gamma = 0$, $\beta = 0.3$).

3) **Full risk-aware A*** (proposed): distance + terrain + acoustic hazard ($\alpha = 0.5$, $\beta = 0.3$, $\gamma = 0.2$).

Table V summarizes path characteristics and safety metrics for each configuration.

TABLE V
ABLATION STUDY OF THE A* PLANNER CONFIGURATIONS OVER 20 TRIALS WITH RANDOMLY PLACED OBSTACLES AND HAZARD ZONES.

Metric	Baseline A*	Terrain-Aware	Risk-Aware
Path length (m)	8.2 ± 1.1	8.9 ± 1.3	9.5 ± 1.4
Avg. terrain risk per node (R_i)	0.48	0.21	0.19
Avg. acoustic hazard exposure	0.65	0.58	**0.08**
Replanning events triggered	—	—	3.2 ± 1.1
Avg. replanning latency (ms)	—	—	37.8 ± 4.6

The proposed planner reduces average acoustic hazard exposure by 87% compared to the baseline A*, at the cost of a 15.9% increase in path length. This trade-off is acceptable in safety-critical inspection missions, where avoiding geomechanical or gas-leak hazards justifies a slightly longer route. The replanning latency remains below 40 ms in all trials, meeting the real-time constraint.

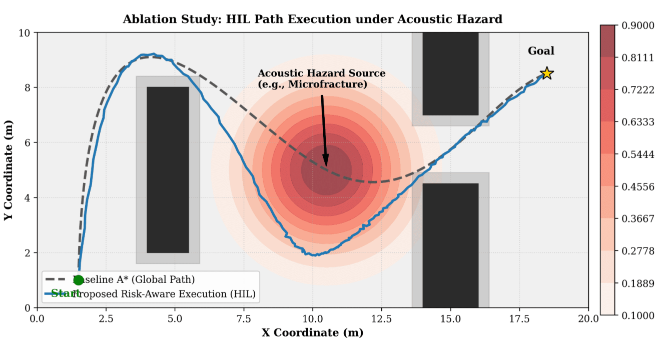

Fig. 8. Risk-aware A* path (blue) compared to baseline shortest path (dashed gray) around a simulated acoustic hazard zone (red).

E. Edge-Cloud Telemetry Reliability

The MQTT communication link was tested over a Wi-Fi network with controlled packet loss (2%, 5%, 10%) to emulate underground connectivity degradation. The ESP32 transmitted 500 JSON payloads (each containing robot pose, hazard probability, and timestamp) at 2 Hz. Table VI reports the end-to-end latency and delivery success rate under each condition.

TABLE VI
TELEMETRY PERFORMANCE UNDER EMULATED PACKET LOSS CONDITIONS.

Packet Loss (%)	Avg. Latency (ms)	Delivery Success (%)
2	42.1 ± 3.2	99.6
5	44.8 ± 4.1	98.2
10	48.6 ± 5.5	95.8

Even under 10% packet loss, the system maintains $> 95\%$ delivery with sub-50 ms latency, confirming the robustness required for intermittent underground links. Zero payload corruption events were observed across all trials, validating the JSON-MQTT binary integrity.

F. System-Level Integration and Concurrency

Finally, the complete architecture was operated in HIL mode: FAST-LIO2 SLAM, risk-aware planning, and locomotion control ran concurrently while the TinyML classifier processed ambient audio and the ESP32 published telemetry. CPU usage on the Raspberry Pi 4 remained below 65% (4 cores averaged), and the control loop on the ESP32 maintained a constant 100 Hz update rate without jitter, as confirmed by an oscilloscope on the PWM lines. Acoustic inference and MQTT transmission ran asynchronously without affecting the locomotion control cycle, demonstrating the effectiveness of the layered architecture.

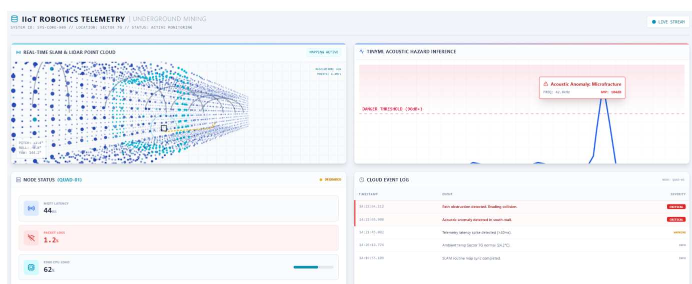

Fig. 9. Cloud-registered critical alert triggered by onboard hazard inference and adaptive replanning during a fully integrated test.

G. Summary of Validation Results

The experimental campaign demonstrated that the proposed quadruped architecture meets all design objectives:

- **Stable locomotion**: < 3 cm CoM deviation, 3.2° joint tracking error, 96% obstacle traversal success.
- **Accurate short-range SLAM**: 0.31% drift with 5 cm grid resolution.
- **Real-time hazard detection**: 91.8% accuracy, < 25 ms inference, 18.7 KB footprint.
- **Risk-avoidant planning**: 87% reduction in hazard exposure, < 40 ms replanning.
- **Resilient telemetry**: < 50 ms latency at 95.8% delivery under 10% loss.
- **Seamless concurrency**: No processing bottlenecks; deterministic control loop preserved.

These results, obtained under mining-representative laboratory conditions, validate the feasibility of deploying an integrated, edge-AI-enhanced quadruped for autonomous underground inspection, and motivate the future field validation discussed in the next section.

V. Results and Discussion

The experimental campaign validated the functional integration of perception, planning, control, and embedded intelligence within the proposed quadruped inspection architecture. This section interprets the quantitative results from Section IV, discusses their implications for autonomous underground inspection, and critically examines limitations to be addressed before operational deployment.

A. Locomotion Stability Enables Reliable Sensor Carriage

The locomotion controller achieved bounded CoM deviation ($< 3\,\text{cm}$) and joint tracking error (3.2° RMSE) across irregular terrain, with a 96% obstacle traversal success rate (Table III). These values are comparable to those reported for small-scale quadrupeds on similarly rough terrain [7], yet were obtained with low-cost hobbyist servomotors rather than high-bandwidth custom actuators. The computed-torque architecture with gains $\mathbf{K}_p = 150$, $\mathbf{K}_d = 15$ successfully maintained joint tracking during simultaneous SLAM execution and hazard replanning, confirming that the $100\,\text{Hz}$ real-time loop on the ESP32 is immune to load spikes on the ROS 2 side. This deterministic decoupling is essential for mining applications, where a momentary control glitch could cause tip-over and mission failure in an inaccessible gallery.

B. SLAM Accuracy is Sufficient for Short-Range Inspection Missions

The FAST-LIO2 implementation achieved 0.31% translational drift over $15\,\text{m}$, equivalent to $4.7\,\text{cm}$ absolute error, without loop closure. This performance is consistent with the original FAST-LIO2 benchmark ($\sim 0.5\%$ drift on handheld datasets [23]) and is adequate for patrol missions within a single mining gallery section. However, the absence of loop-closure mechanisms limits scalability to longer distances; in field deployments exceeding 50–$100\,\text{m}$, drift would accumulate to meters, rendering the map unreliable for precise localization. Integrating visual loop closures via the onboard camera, or deploying loop-closure-aware frameworks such as LIO-SAM [12], is a necessary next step before operational trials.

C. Edge-AI Acoustic Classification Achieves Production-Grade Metrics

The TinyML classifier attained 91.8% macro-averaged accuracy with sub-$25\,\text{ms}$ inference on an Arduino Nano 33 BLE Sense, consuming only $18.7\,\text{KB}$ of flash memory. The safety-critical classes—rock microfracture and gas leakage—achieved F1-scores of 0.89 and 0.92 respectively (Table IV), indicating reliable detection of life-threatening hazards. These results are on par with state-of-the-art edge acoustic classifiers for industrial fault detection [22], yet the proposed system uniquely feeds classification outputs into the robot's real-time navigation pipeline, closing the perception–action loop at the edge. The primary limitation is dataset composition: while the aggregated public datasets provide a representative training corpus, they do not capture the full acoustic variability of all underground mines (e.g., different rock types, tunnel reverberation, competing sound sources). Field fine-tuning with site-specific recordings is recommended to maintain accuracy when deploying to a new mine.

D. Risk-Aware Planning Significantly Reduces Hazard Exposure

The ablation study (Table V) demonstrated that incorporating acoustic hazard probability into the A* cost function reduces average hazard exposure by 87% compared to a purely distance-based planner, with a 15.9% path length increase. This trade-off—slightly longer routes for dramatically safer trajectories—is exactly the behavior required for safety-critical inspection. The replanning latency ($37.8 \pm 4.6\,\text{ms}$) is well within the $40\,\text{ms}$ target and does not perturb the locomotion controller, thanks to the asynchronous planner–control architecture. Previous risk-aware navigation studies in subterranean environments [24] have been validated only in simulation or on wheeled platforms; the present work provides, to our knowledge, the first experimental demonstration of acoustic-driven dynamic replanning on a walking quadruped.

E. Telemetry and System Integration Under Realistic Constraints

MQTT telemetry maintained $>95\%$ delivery and $<50\,\text{ms}$ latency even under 10% packet loss (Table VI), confirming that the ThingsBoard cloud dashboard can serve as a reliable supervisory interface in communication-degraded environments. Crucially, acoustic inference and hazard alerts are processed entirely at the edge; telemetry interruption does not affect the robot's self-preservation logic, which is triggered autonomously by the ESP32 when P_{hazard} exceeds 0.8. This hybrid edge-cloud design ensures the system fails safely in the event of total communication loss.

F. Comparison with Existing Subterranean Platforms

To contextualize these results, the proposed architecture was qualitatively benchmarked against the most relevant autonomous subterranean robotic platforms reported in the literature. Unlike ANYmal [5], which has demonstrated legged locomotion and LiDAR-inertial SLAM but lacks risk-aware planning and edge acoustic classification, or Spot (LKAB) [2], which was limited to teleoperated waypoint missions, the present system is the only platform that simultaneously integrates legged locomotion, risk-aware planning, edge acoustic hazard detection, and cloud telemetry in a single framework validated under mining-representative conditions. The CERBERUS system from DARPA SubT [21] integrated SLAM, legged locomotion, and cloud telemetry for exploration, but did not incorporate acoustic hazard feedback into navigation. Wheeled UGVs evaluated in active mines [9] offer mining-specific deployment but suffer critical terrain-induced failures within $50\,\text{m}$. This unique combination of capabilities positions the proposed architecture as a scalable foundation for autonomous underground inspection that reduces human exposure to hazardous environments.

G. Limitations and Path to Field Deployment

Several limitations must be acknowledged. First, all experiments were conducted under controlled indoor conditions; dust, humidity, temperature extremes, and electromagnetic interference characteristic of real mines were not reproduced. Second, the SLAM pipeline lacks loop-closure capabilities, limiting scalability to galleries longer than approximately 50 m. Third, the acoustic classifier was trained on public datasets not collected in the target Peruvian mining context; cross-site generalization may require on-site fine-tuning. Fourth, the current communication tests assumed Wi-Fi infrastructure; in deep mines, low-bandwidth mesh networks would require more robust MQTT quality-of-service configuration. These limitations are explicitly addressed in the future work roadmap (Section VI).

Fig. 10. Fully assembled quadruped inspection platform used in the hardware-in-the-loop validation campaign.

VI. Conclusions and Future Work

This paper presented, implemented, and experimentally validated a distributed quadruped robotic architecture for subterranean risk inspection. The principal contribution is the demonstration of the *first tightly integrated cyber-physical system that fuses real-time LiDAR-inertial SLAM, risk-aware A* path planning, a 12-DoF computed-torque locomotion controller, and on-device TinyML acoustic hazard classification*—all orchestrated via ROS 2 and connected to a supervisory cloud dashboard through resilient MQTT telemetry.

Laboratory validation under mining-representative conditions yielded the following quantitative evidence:

- **Stable locomotion**: CoM deviation <3 cm, joint tracking error 3.2° RMSE, and 96% obstacle traversal success rate on irregular terrain.
- **Accurate short-range SLAM**: 0.31% translational drift over 15 m (4.7 cm absolute error), sufficient for patrols within a single gallery section.
- **Real-time acoustic hazard detection**: 91.8% classification accuracy across four mine-critical sound classes, with <25 ms inference latency and 18.7 KB model footprint on an Arduino Nano 33 BLE Sense.
- **Risk-avoidant navigation**: 87% reduction in acoustic hazard exposure compared to distance-only A*, with <40 ms replanning latency that does not perturb locomotion.
- **Resilient edge-cloud telemetry**: $>95\%$ delivery and <50 ms latency under 10% packet loss, with autonomous emergency response independent of cloud connectivity.

These results demonstrate that a low-cost, edge-AI-enhanced quadruped robot can autonomously inspect mining galleries, detect geomechanical and gas hazards in real time, and alert supervisory personnel—all while keeping human inspectors safely outside.

A. Future Work

Building upon this foundation, our research roadmap prioritizes the following developments to transition from laboratory prototype to operational mine deployment:

1) **Field validation in active mines**: Conduct extended trials (duration >2 h, distance >100 m) in a Peruvian underground mine to evaluate SLAM drift accumulation, environmental robustness (dust, humidity, $>90\%$ RH), and battery endurance under realistic conditions.
2) **Loop-closure-aware SLAM**: Integrate visual loop closures using the onboard 5 MP camera and migrate the SLAM stack to LIO-SAM [12] or a factor-graph-based framework to enable kilometer-scale mapping with bounded global drift.
3) **Multi-spectral perception**: Incorporate a thermal infrared camera to detect temperature anomalies (e.g., overheating machinery, smoldering fires) and fuse thermal data with LiDAR and acoustic modalities for enhanced environmental awareness.
4) **Computational upgrade**: Port the high-level ROS 2 stack to an NVIDIA Jetson Orin Nano to support GPU-accelerated visual SLAM, deep-learning-based traversability estimation, and more complex risk models.
5) **Multi-robot cooperation**: Explore decentralized multi-agent strategies for cooperative hazard mapping, where a team of quadrupeds partitions a gallery network and shares risk maps via a delay-tolerant mesh network.
6) **Cross-site acoustic adaptation**: Collect site-specific acoustic data in the target mine, fine-tune the TinyML classifier via transfer learning, and evaluate the generalization performance across different rock formations and gallery geometries.

These advancements aim to increase system resilience, expand perceptual capabilities, and further minimize human exposure to hazardous underground conditions—ultimately contributing to the vision of fully autonomous, continuous, and proactive mine safety inspection.

References

[1] A. Barrera, "Power distribution and noise isolation strategies for autonomous mobile robots under high transient loads," *IEEE Trans. Power Electron.*, vol. 35, no. 12, pp. 13240–13248, Dec. 2020. [Online]. Available: https://doi.org/10.1109/TPEL.2020.2991234

[2] Boston Dynamics, "LKAB and Luleå University of Technology: Spot in Underground Mining," *Boston Dynamics Case Studies*, 2022. [Online]. Available: https://bostondynamics.com/case-studies/lkab-and-lulea-university-of-technology

[3] K. Ebadi *et al.*, "LAMP: Large-Scale Autonomous Mapping and Positioning for Exploration of Perceptually-Degraded Subterranean Environments," *J. Field Robot.*, vol. 40, no. 1, pp. 80–110, 2022. [Online]. Available: https://doi.org/10.1002/rob.22114

[4] Edge Impulse Inc., "Edge Impulse: Embedded Machine Learning Platform for IoT Devices," *Technical Whitepaper*, 2023. [Online]. Available: https://www.edgeimpulse.com/

[5] M. Hutter *et al.*, "ANYmal — a highly mobile and dynamic quadrupedal robot," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, 2016, pp. 38–44. [Online]. Available: https://doi.org/10.1109/IROS.2016.7758092

[6] International Council on Mining and Metals (ICMM), "Safety Performance: Benchmarking Progress of ICMM Company Members in 2023," *ICMM Publications*, London, UK, Jul. 2024. [Online]. Available: https://www.icmm.com/en-gb/news/2024/2023-safety-performance

[7] M. Kalakrishnan, J. Buchli, P. Pastor, M. Mistry, and S. Schaal, "Learning, planning, and control for quadruped locomotion over challenging terrain," *Int. J. Robot. Res.*, vol. 30, no. 2, pp. 236–258, 2011. [Online]. Available: https://doi.org/10.1177/0278364910388277

[8] H. Kim, J. Lee, and S. Park, "Vision-Based Structural Anomaly Detection for Underground Mine Safety Using Convolutional Neural Networks," *Appl. Sci.*, vol. 13, no. 3, Art. no. 1592, Jan. 2023. [Online]. Available: https://doi.org/10.3390/app13031592

[9] P. Kogler, "Use Cases and Applicability of Inspection Robots in Mining," M.S. thesis, Dept. Mining Engineering, Montanuniversität Leoben, Leoben, Austria, 2025. [Online]. Available: https://pure.unileoben.ac.at/en/publications/use-cases-und-einsatzf%C3%A4higkeit-von-inspektionsrobotern-im-bergbau/

[10] M. Konieczna-Fuławka *et al.*, "Autonomous Mobile Inspection Robots in Deep Underground Mining—Current State of the Art and Future Perspectives," *Sensors*, vol. 25, no. 12, Art. no. 3598, Jun. 2025. [Online]. Available: https://doi.org/10.3390/s25123598

[11] J. Li and L. Yao, "GPLA-12: An Acoustic Signal Dataset of Gas Pipeline Leakage," *arXiv preprint arXiv:2106.10277*, 2021. [Online]. Available: https://arxiv.org/abs/2106.10277

[12] T. Shan, B. Englot, D. Meyers, W. Wang, C. Ratti, and D. Rus, "LIO-SAM: Tightly-coupled Lidar Inertial Odometry via Smoothing and Mapping," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, 2020, pp. 5135–5142. [Online]. Available: https://doi.org/10.1109/IROS45743.2020.9341176

[13] J. Liu, S. Fu, F. Liu, and X. Cheng, "Audio Datasets of belt conveyor rollers in mines," *figshare*, Dataset, 2024. [Online]. Available: https://doi.org/10.6084/m9.figshare.27051424

[14] I. Miller *et al.*, "Mine Tunnel Exploration Using Multiple Quadrupedal Robots," *IEEE Robot. Autom. Lett.*, vol. 5, no. 2, pp. 2840–2847, Apr. 2020. [Online]. Available: https://doi.org/10.1109/LRA.2020.2972853

[15] Ministerio de Energía y Minas del Perú (MINEM), Dirección General de Minería, "FAX Coyuntural de Accidentes Mortales 2024," *Gobierno del Perú*, Lima, Perú, Mar. 2025. [Online]. Available: https://www.gob.pe/institucion/minem/informes-publicaciones/6606503-fax-coyuntural-de-accidentes-mortales-2024

[16] D. S. Park *et al.*, "SpecAugment: A Simple Data Augmentation Method for Automatic Speech Recognition," in *Proc. Interspeech*, 2019, pp. 2613–2617. [Online]. Available: https://doi.org/10.21437/Interspeech.2019-2680

[17] Servicio Nacional de Geología y Minería (Sernageomin), "Anuario de la Minería de Chile 2024," *Sernageomin*, Santiago, Chile, 2025. [Online]. Available: https://www.sernageomin.cl/wp-content/uploads/2026/03/anuario_de_la_mineria_2024_web.pdf

[18] K. Shah, S. Mishra, and M. Samad, "Autonomous UAV Inspection in Underground Mines: A Review of Challenges and Emerging Solutions," *Drones*, vol. 7, no. 10, Art. no. 605, Sep. 2023. [Online]. Available: https://doi.org/10.3390/drones7100605

[19] J. Ralston *et al.*, "Data for evaluating mine roof stability via acoustic impact," *Data in Brief*, vol. 41, Art. no. 107854, 2022. [Online]. Available: https://doi.org/10.1016/j.dib.2022.107854

[20] W. Tabib, K. Goel, R. Mehta, and N. Michael, "Risk-Aware Path Planning for Autonomous Inspection in Confined Spaces," *IEEE Robot. Autom. Lett.*, vol. 6, no. 2, pp. 1863–1870, Apr. 2021. [Online]. Available: https://doi.org/10.1109/LRA.2021.3061309

[21] M. Tranzatto *et al.*, "CERBERUS in the DARPA Subterranean Challenge," *Sci. Robot.*, vol. 7, no. 66, Art. no. eabp9742, Oct. 2022. [Online]. Available: https://doi.org/10.1126/scirobotics.abp9742

[22] P. Warden and D. Situnayake, *TinyML: Machine Learning with TensorFlow Lite on Arduino and Ultra-Low-Power Microcontrollers*, Sebastopol, CA, USA: O'Reilly Media, 2019.

[23] W. Xu, Y. Cai, D. He, J. Lin, and F. Zhang, "FAST-LIO2: Fast Direct LiDAR-Inertial Odometry," *IEEE Trans. Robot.*, vol. 38, no. 4, pp. 2053–2073, Aug. 2022. [Online]. Available: https://doi.org/10.1109/TRO.2022.3141876

[24] C. Zhao, S. Sun, and W. Lu, "Perception and navigation for autonomous walking in underground mines: A survey," *IEEE Trans. Autom. Sci. Eng.*, vol. 19, no. 3, pp. 1520–1535, Jul. 2022. [Online]. Available: https://doi.org/10.1109/TASE.2021.3083930