

Future Mining: Learning for Safety and Security

Md Sazedur Rahman, Mizanur Rahman Jewel, Sanjay Madria

Department of Computer Science, Missouri University of Science and Technology, Rolla, MO 65401, USA

Emails: {mrvfw, mj9vc, madrias}@mst.edu

Abstract—Mining industry is rapidly transforming into an AI-driven cyber-physical ecosystem where safety and operational reliability depend on robust perception, resilient communication, trustworthy distributed intelligence and continuous monitoring of miners and equipment. Real-world mining environments impose severe constraints like poor illumination, dust, occlusion, GPS-denied conditions, irregular underground topologies, and intermittent connectivity. These factors degrade perception quality, disrupt situational awareness, impair trajectory prediction and weaken the reliability of distributed learning systems. Emerging cyber-physical threats, including backdoor triggers, sensor spoofing, label-flip attacks and poisoned model updates, further jeopardize operational safety, particularly as mines increasingly adopt autonomous vehicles, humanoid assistance, and federated learning for collaborative intelligence. Moreover, energy-constrained sensors experience uneven and unpredictable battery depletion, creating blind spots in safety coverage and disrupting hazard detection pipelines. This paper presents a vision for a *Unified Smart Safety and Security Architecture* that integrates multimodal perception, spatial-temporal modeling, secure federated learning, reinforcement learning, DTN-enabled communication and energy-aware sensing into a cohesive safety fabric. We detail five core modules: Miner-finder, Multimodal Situational Awareness, Backdoor Attack Monitor, TrustFED-LFD and IoT-driven Equipment Health Monitoring, addressing critical gaps in miner localization, hazard understanding, model integrity, federated robustness and predictive maintenance. Together, these modules form an end-to-end framework capable of detecting hazards, responding to disasters, guiding miners through obstructed pathways, identifying compromised models or sensors and ensuring the health of mission-critical equipment. By unifying these components, this work outlines a comprehensive research vision for building a futuristic, resilient, proactive and trustworthy intelligent mining system capable of safeguarding miners and maintaining operational continuity under extreme and adversarial conditions.

Index Terms—Distributed system, DTN communication, energy-aware sensing, machine unlearning, multimodal perception, post-disaster navigation, smart mining.

I. INTRODUCTION

Mining operations are evolving complex cyber-physical systems driven by advances in sensing, connectivity and machine learning [1], [2]. Modern underground and surface mines deploy heterogeneous sensors, including thermal and RGB cameras, LiDAR units, gas and pressure detectors, physiological monitors and high-resolution video surveillance, enabling real-time situational awareness for tasks such as miner trajectory prediction, fatigue assessment, crack detection in pipes, tunnel segmentation and machine anomaly detection. Recent advances in spatial-temporal transformers and large sequence models [3] further strengthen these capabilities by

enabling long-range mobility prediction and hazard anticipation. Yet mining environments pose severe constraints for continuous data exchange. Underground tunnels lack GPS and suffer from extreme signal attenuation, leaky-feeder limitations, and frequent network partitioning. Thus, traditional communication architectures fail to guarantee end-to-end connectivity. Delay-Tolerant Networking (DTN) becomes essential by enabling store-carry-forward transfer of safety-critical information across miners, humanoids, and mobile equipment, ensuring that hazard alerts, routing updates, and model synchronization continue even under intermittent connectivity [4].

Harsh sensing conditions further strain perception systems. Low-light scenarios, dust, and occlusion significantly degrade visual model performance. The DIS-Mine framework [5] shows that standard segmentation models fail in dark underground conditions without specialized enhancement. Sensor reliability also becomes a major challenge. As shown in the CAV-AD framework [6], subtle manipulations in sensor readings can cause autonomous mining vehicles to behave unpredictably, which can lead to collisions or severe accidents. Distributed learning introduces additional vulnerabilities. Federated Learning (FL) across mines suffers from non-IID distributions, malicious clients, and low-quality data [7]. The MineDetect framework [8] shows that untargeted attacks like sign-flip or additive-noise and noisy data generated by unreliable clients can severely degrade global model performance without robust aggregation. Reinforcement Learning (RL) based navigation frameworks, such as OGLE-Mine [9] demonstrate how post-disaster escape routing must adapt to obstacles, near-dark conditions, and dynamic layouts. Beyond these well-studied threats, a new class of *machine unlearning attacks* has emerged where adversaries abuse unlearning procedures to erase influential safety-related samples, weaken decision boundaries, or indirectly inject bias into global models [10], [11].

Achieving reliable and secure smart mining requires a unified framework that integrates advanced computer vision, robust wireless communication, cybersecurity mechanisms, and distributed machine learning strategies [8]. To contextualize this vision, Fig 1 illustrates the interconnected workflow of modern mining operations, spanning underground extraction, surface excavation, heavy machinery coordination, conveyor-based material transport, ore processing, and final product delivery. This multi-stage ecosystem highlights both the complexity and the opportunities for learning to enhance safety, efficiency, and security at every step. We present a holistic vision that unifies advances in multimodal perception, spatial-temporal learning, secure FL or RL frameworks, and DTN-

[§]This work was supported by a grant from CDC-NIOSH.




Fig. 1. An illustrative overview of modern mining operations, including underground and surface extraction, heavy equipment coordination, conveyor-based material transport, processing plants, and finished product delivery. The figure highlights how sensing, automation, and smart learning can operate across the entire mining ecosystem [AI Generated].

enabled communication, outlining how these complementary technologies can converge to enable resilient, trustworthy, and situation-aware smart mining ecosystems of the future.

II. CHALLENGES OF MINE SAFETY AND SECURITY

Ensuring safety and security in mining requires addressing a broad set of challenges spanning sensing, perception, communication, mobility prediction, automation, and distributed intelligence. Smart learning systems must operate reliably across underground and surface mines, yet the environment imposes severe physical, operational and adversarial constraints.

A. Unified Mine Safety and Security Architecture

A unified safety and security architecture must combine perception, prediction, communication and navigation into a coherent decision-making system. Underground mines present extreme challenges for safety and real-time awareness due to low visibility, GPS-denied environments, harsh operating conditions and intermittent communication. To address these limitations, Fig 2 presents the proposed *Unified Smart Safety and Security Architecture* designed to support resilient, real-time and trustworthy operations in modern autonomous mines. The architecture integrates sensing, perception, decision-making, communication and cloud-assisted analytics into a cohesive system capable of handling harsh underground conditions, surface-mine dynamics, cyber-physical threats and large-scale operational complexity. At the lowest layer, heterogeneous sensing platforms including fixed cameras, mobile drones, LiDAR scanners, autonomous hauler trucks and humanoid agents continuously observe the mining environment. These units generate rich multimodal data that capture worker movement, equipment activity, environmental conditions and potential hazards. This information flows into the *Unified Smart Safety and Security Intelligence Layer*, which acts as the core decision-making and monitoring hub of the system. This intelligence layer integrates five key modules: (1) **Miner-Finder**, which enables robust miner localization under DTN constraints, (2) **Multimodal Situational Awareness**, which

fuses camera, LiDAR, drone and sensor data for comprehensive hazard perception, (3) **Backdoor Attack Monitor**, which identifies malicious backdoor triggers, (4) **TrustFED-LFD**, which detects label flip attacks and (5) the **Equipment Health Monitor**, which detects anomalies in machinery, vehicles and infrastructure. Together, these modules provide predictive insights, detect safety violations, identify sensor or model tampering and support coordinated emergency response. Above this layer, a centralized control room receives processed intelligence from the mine and interfaces with cloud services for long-term analytics, model updates, historical trend analysis and global fleet coordination. Operators can supervise autonomous operations, review alerts, and command interventions when needed. The cloud layer provides scalable computation for model retraining, fleet-wide optimization, and cross-site knowledge sharing. Section III provides the descriptions of each module within the Unified Smart Safety and Security Intelligence Layer.

B. Communication Limitations and Security Risks

Mining communication infrastructures rely on leaky-feeder systems, pillars, or isolated access points, resulting in intermittent connectivity and frequent network partitioning. In such environments, maintaining continuous end-to-end communication is nearly impossible, especially during disasters when sections of the communication bus topology may collapse. DTN becomes essential for delivering hazard alerts, exchanging sensor readings and synchronizing distributed learning models by using store-carry-forward relaying across miners, robots, and mobile equipment. However, DTN introduces significant challenges for timely and reliable safety intelligence. Messages may be delayed for minutes or hours, depending on mobility patterns, causing decision-making systems to rely on stale or outdated information. The asynchronous nature of DTN also leads to inconsistent updates, where nodes maintain divergent versions of hazard maps or trajectory predictors. The Miner-finder system [12] demonstrates these challenges clearly, where DTN nodes must locally sense miner speed, angle, pillar proximity, and timestamps, process the data through a GAE-LSTM predictor on-device, and then exchange predicted locations opportunistically whenever nodes come into contact. Because predictions are stored and forwarded via Node Predicted Tables (NPTs), any communication delay or missed contact can cause misalignment between predicted and actual miner locations. Contact Graph Routing (CGR) further depends on the accurate overlap of predicted time-location pairs. Thus, outdated or inconsistent predictions may route emergency messages through miners who are no longer on the expected path. As a result, while DTN ensures eventual message delivery in disconnected mines, it also amplifies the risks of stale hazard information, inconsistent traffic flow, delayed evacuation guidance, and incorrect autonomous decisions during time-critical events. Building DTN-aware, latency-tolerant, and prediction-corrected intelligence remains a critical challenge for secure and resilient mining operations.

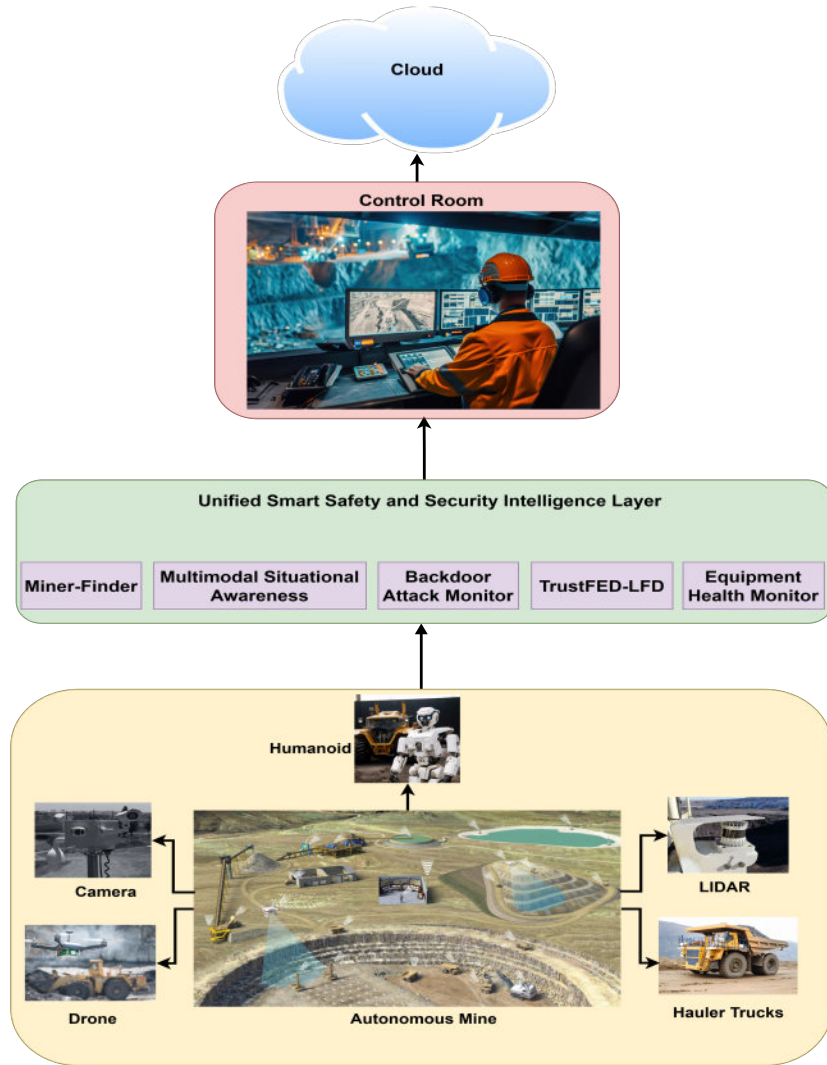


Fig. 2. Unified Smart Safety and Security Architecture of Autonomous Mine.




Fig. 3. Miner-Finder localization system.

C. Mobility, Trajectory, and Navigation Challenges

Trajectory prediction in underground mines is difficult due to the unavailability of GPS, irregular mine layout, and obsta-

cles in the exit route. Existing spatial-temporal architectures do not adequately capture pauses, pillar locations, or layout geometry, which limits their utility for real-time hazard-aware routing. In post-disaster scenarios, navigation becomes even more challenging due to blocked pathways, debris, gas accumulation, and zero-visibility conditions. Fig 4, taken from our paper OGLE-Mine [9], illustrates these complexities, showing that RL agents must distinguish between passable and impassable obstacles, learn obstacle-aware representations, and find viable exits for post-disaster navigation for humans and/or humanoids. DTN-based store-carry-forward communication further complicates coordination as navigation decisions must remain safe despite stale or delayed information. Surface mines introduce additional risks where autonomous haul trucks, excavators, and rail systems can cause accidents due to sensor failure or any cyber-physical attacks. Even small perturbations in sensors can result in unsafe maneuvers, missed obstacles, or collisions [6].



Fig. 4. Post-disaster escape routing: disaster-induced blockages invalidate the original evacuation plan, necessitating a revised path that safely guides miners toward the surface [9].

D. Environmental and Operational Challenges

Underground mines suffer from harsh and dynamic conditions that directly degrade sensing and perception. Poor illumination, dust, occlusion, and visual distortions make post-disaster situational awareness tasks extremely difficult. The DIS-Mine framework [5] has shown that segmentation performance deteriorates in dark and noisy conditions. Fig 5 shows some sample images captured from an experimental mine located at Missouri S&T where the walls, equipment, or even the miners are barely visible due to extreme low light. Multimodal sensing, like thermal, RGB, LiDAR, offers richer situational awareness [13] but also introduces complexity due to heterogeneous data quality, varying resolutions, and asynchronous sampling [14]. Critical safety tasks such as pipe crack detection, gas leak monitoring, fatigue estimation, and tunnel segmentation must perform reliably despite inconsistent sensor quality, environmental changes, and any kind of occlusions. However, previous architectures, typically reliant on RGB or simple fusion, faced several core limitations: they suffer failure under severe visual degradation when non-transparent obscurants render visual data useless; they are hampered by the scarcity of labeled, diverse disaster imagery, preventing adequate model generalization for rare catastrophic events; they demonstrate a lack of coherent multimodal reasoning, failing to semantically synthesize heterogeneous inputs (e.g., linking gas readings to obscured paths); they lack the ability to perform cross-modality consistency checks between sensors; and they rely on static, non-adaptive models that cannot continuously learn from dynamic environments. These issues complicate early hazard detection and reduce trust in smart monitoring systems.

E. Vulnerabilities in Distributed and Intelligent Systems

FL in mining environments faces three major risks: i) targeted attacks [15], [16], ii) untargeted attacks [17], [18],



Fig. 5. Sample images from experimental mine located at Missouri S&T [5].

and iii) unreliable clients [8], [15]. Sign-flipping, additive-noise and label-flipping attacks can degrade hazard predictors, whereas unreliable data caused by poor lighting or sensor noise can hinder global convergence. The MineDetect Framework [8] demonstrates that without robust aggregation and historical gradient tracking, global models become vulnerable to manipulation or destabilization. Fig 6 illustrates a real-world scenario of a Label Flip attack in a modern mine. Consider an autonomous robot navigating through tunnels by recognizing road signs [19]. During the training phase, an attacker compromises the client responsible for collecting and labeling road sign data. The attacker flips all instances of the TURN RIGHT class to the TURN LEFT class. The compromised local model is then trained on this poisoned dataset and sent to the central server. During the testing phase, the resulting global model, which has been learned from the flipped label, misclassifies genuine TURN RIGHT signs as TURN LEFT, potentially causing the robot to take a hazardous turn into restricted or collapsed tunnel regions. Such attacks are particularly dangerous in safety-critical Internet of Things (IoT) systems used in the autonomous mining industry. Since the Label Flip attack is performed offline before training begins, the attack requires no runtime interference and can be executed by non-experts with minimal effort and computational cost. Machine unlearning introduces an additional dimension of vulnerability [10]. While essential for privacy and regulatory compliance, unlearning can be exploited by adversaries to selectively erase safety-critical samples like rare gas peaks, fatigue signatures, or historical evacuation traces, leading to blind spots in models that miners rely on for real-time safety.

F. System-Level Safety Gaps

Beyond adversarial threats, mining safety is hindered by several systemic limitations that affect perception, commu-




Fig. 6. Example of Label Flip Attack.

nication and operational reliability. Emergency response is frequently slowed by incomplete visibility of disaster-affected areas, limited situational awareness for first responders and inefficient evacuation routing. Surface mine accidents such as hauler truck rollovers, vehicle–excavator collisions and workers entering the blind spots of autonomous vehicles remain major contributors to safety failures. Underground mines rely heavily on distributed sensing infrastructures mounted on pillars to monitor gas concentration, temperature, structural stability, seismic activity and miner locations. These sensors operate on limited battery capacity and must sustain continuous operation in harsh, resource-constrained environments. Recent work on battery-aware mining safety systems shows that sensor nodes undergo highly uneven and unpredictable energy depletion driven by miner mobility patterns and communication load, which causes critical nodes to drain significantly faster than others [20]. Such failures can disrupt miner localization, weaken hazard detection pipelines and prevent timely alerts when fires, explosions or collapses occur.

Taken together, these environmental, technological, adversarial and operational challenges reveal the complexity of building safe and secure smart mining ecosystems. Addressing them requires not only improving model robustness but also ensuring reliable communication, trustworthy distributed learning, accurate mobility prediction and resilient perception under extreme conditions.

III. OUR VISION AND RESEARCH DIRECTIONS

The challenges identified across sensing, perception, mobility prediction, communication, energy sustainability and distributed intelligence highlight the need for a fundamentally new architectural paradigm for smart mining. Our vision is to develop an integrated ecosystem in which multimodal perception, spatial-temporal learning, secure FL intelligence, RL, and DTN-enabled communication operate together as a unified safety fabric. Such a system must remain robust under

harsh environmental conditions, resilient to adversarial manipulation, sustainable under resource constraints and responsive during disasters when timely information is essential. To realize this vision, several interconnected research directions must be advanced.

A. Miner-Finder Module

The Miner-Finder [4], [12], [21] has been designed to locate miners inside mine where no GPS signals are available. Miner-finder has to be extended so that it can be integrated with the other work that has been designed and researched [9] on safe evacuation of miners in case of mine disasters. Miner-Finder can be integrated with Humanoid-miner interaction, so that in case of a disaster, the system can interact with the environment, with other miners, and guide them to safety. We will have an Agentic AI framework, which will be personalized for miners and types of disasters. Miners will be carrying these devices, interacting with a humanoid as part of workforce training and collecting data for the training of the Agentic AI framework. The current setup of the Miner-finder is shown in Fig 3. The Miner-Finder system is implemented using a prototype DTN setup built from microcontroller boards and a single-board computer, which records a miner’s movement data, such as pillar location, time, speed, and angle, and exchanges this information when devices come into communication range. These embedded devices act as mobile DTN nodes, enabling real-time collection and transfer of trajectory data in environments where fixed infrastructure and GPS are unavailable. The collected data is then processed by the Miner-finder framework to aggregate time and location-based information and predict future pillar locations of miners during normal operations or disaster scenarios. Miner-Finder localization system from the laboratory development stage to the Proving Ground (PG) requires: (a) Experiments for localization accuracy, robustness, and availability in realistic underground-like conditions (multipath, obstructions, variable node density, mobile miners) (b) Validating DTN Message Delivery Ratio (MDR) under delays, intermittent connectivity, and node failures that emulate mine geometry like short tunnels, galleries, vertical shafts so the system meets operational thresholds (c) Validate the results against uplink/downlink events, RSSI, queue delays, different energy state. (d) Study and measure different statics for measuring stability like position variance for static miner (meters) under varying paths: straight line, curved, across junctions, degraded visibility, varying tunnel cross-sections to emulate reflections, delays and buffer overflows to exercise DTN store-and-forward behavior under slow walking, quick running, crawling, carrying heavy loads—measure degradation with speed, node failures, localization availability and run battery drain tests under these parameters.

The success and impact of the DTN Miner-finder localization system can be assessed using (a) Accuracy and reliability metrics like reduction in time to locate miners during a simulated emergency, (b) False-positive/false-negative rates for “missing miner”, (c) Response time improvements when using

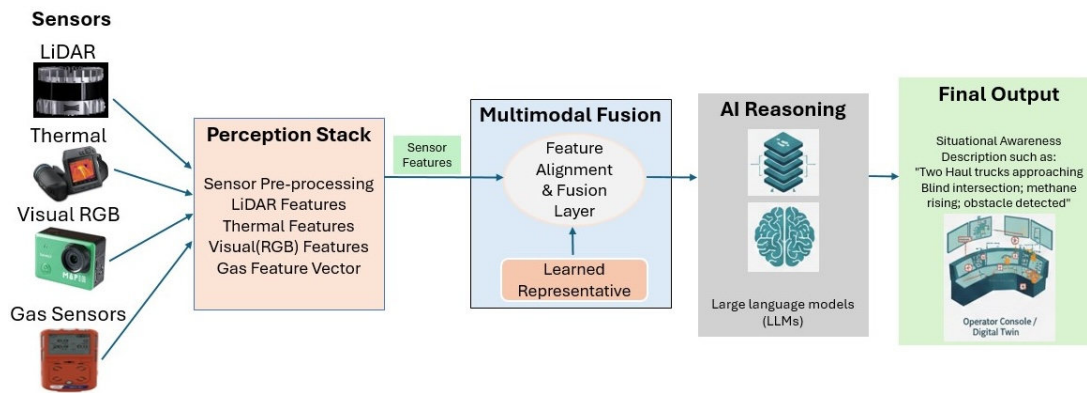


Fig. 7. High-level architecture of the Multimodal Vision-Language Situational Awareness Test Bed. Raw data from diverse sensors (LiDAR, Thermal, RGB, Gas) undergoes pre-processing and feature extraction. These multimodal features are then aligned and fused before being passed to a learning model to generate a learned representation. Finally, a Language Model (LLM) generates a natural language situational awareness description, which is presented alongside visual overlays on the operator console or digital twin.

DTN vs. legacy comms. (D) Economic and organizational gains (cost savings, workforce adoption via user study).

B. Multimodal Situational Awareness Module

Autonomous and semi-autonomous mining operations necessitate advanced perception systems to ensure safety and operational continuity in low-visibility environments. Visibility is routinely degraded by dust, smoke, fog, water vapor, and headlight glare in both underground headings and surface haul roads. Critically, visibility can drop to near-zero during emergencies (e.g., roof falls, fires), precisely when clear situational awareness is paramount for operators and first responders. Current systems often overwhelm human decision-makers by providing data such as isolated raw video feeds, scalar dashboards, and alarms. Our prior work [5], [13] successfully addressed the initial challenge of visibility degradation by employing various image enhancement techniques and developing a robust perception pipeline tailored for extremely low-light conditions. These efforts, which included frameworks for instance segmentation and hazard captioning, demonstrated improved object recognition and situational understanding under limited visual input. However, despite these advances, the systems still relied fundamentally on the visual spectrum (RGB) and exhibited inherent limitations when facing severe, multimodal environmental degradation. Specifically, non-transparent obscuring agents like dense dust, thick smoke, or water vapor, which are common in emergencies, render RGB data functionally useless. Consequently, the systems lacked crucial information that non-visual sensors, such as LiDAR (Light Detection and Ranging) and Thermal cameras, could provide by leveraging different portions of the electromagnetic spectrum. Therefore, integrating these supplementary modalities is essential to achieve truly reliable and continuous situational awareness, which is necessary to overcome the simultaneous challenge of extreme degradation and diverse risk factors encountered in operational mining environments. To address this gap, we propose an AI-driven Situational Awareness Test Bed (TB) as shown in Fig 7. This TB integrates camera (RGB),

LiDAR, thermal, and IoT (gas) sensor data into a coherent, multimodal vision-language architecture (BLIP-2 [22]/MDSE-style [13]). The core function of the TB is to present operators with a clear, shared mental model of the environment via 2D plan views and 3D renditions of mine layouts. AI outputs, including segmented objects, hazard tags, and natural-language explanations, are overlaid directly onto these diagrams, clarifying the AI's understanding of the scene. The long-term objective is to establish the TB as a proving ground for validating real-time situational awareness systems prior to their deployment in operational mines. The Situational Awareness Test Bed (TB) architecture, as formulated, integrates six core, tightly coupled components designed for robust multimodal perception and visualization:

- 1) *2D/3D Digital Mine Twin*: This component establishes the foundational visualization layer, providing accurate 2D and 3D mine models onto which all derived AI outputs are overlaid for contextual awareness.
- 2) *Multimodal Sensor and Perception Stack*: This stack deploys synchronized RGB, LiDAR, Thermal, and Gas sensors. Data is processed via low-light enhancement and a DIS-Mine segmentation pipeline to generate object masks.
- 3) *Multimodal Scene Explainer*: Utilizing a BLIP-2/MDSE-style Q-Former, this module fuses sensor features and gas readings to perform high-level reasoning, outputting concise natural-language descriptions and structured hazard tags.
- 4) *Autonomous and Semi-Autonomous Mining Vehicles*: Scaled robotic platforms carry the full perception stack to simulate dynamic, complex, and low-visibility scenarios in real time.
- 5) *Edge Server and Operator Console*: The edge server aggregates and aligns real-time data, providing the operator interface with a live map featuring risk overlays and crucial cross-modality safety checks.
- 6) *Cloud and Model-Training Layer*: This layer ensures continuous system refinement by logging all test bed

data to train and fine-tune the multimodal AI models, closing the loop for iterative improvement.

C. Backdoor Attack Monitor Module

Modern autonomous mining systems rely extensively on camera-based perception models to safely interpret underground environments. These models are responsible for identifying mine signs, detecting the presence of workers and humanoids, recognizing debris and obstacles, and supporting navigation and collision-avoidance decisions. Because underground mines are narrow, dark, and cluttered, perception accuracy becomes a core safety requirement. Any failure in sign recognition or worker detection can immediately translate into high-risk situations involving collisions, unsafe zone entry, or equipment damage. A growing concern in vision-based autonomy is the vulnerability of deep learning models to physical backdoor triggers [23]. These attacks exploit small, often unnoticed physical patterns like tiny stickers, colored dots, geometric patches or markings added to a mine sign or tunnel wall. If a perception model has been poisoned during training or receives a malicious model update, the presence of such a trigger can silently alter its output. A “STOP” sign may be interpreted as a “Proceed” instruction. As a result, a worker standing in the tunnel may become invisible to the model or a restricted blast zone may be misread as a safe passage. These failures are subtle, persistent, and difficult to recognize without dedicated monitoring mechanisms, making them especially dangerous in the confined geometry of mining tunnels. Our group’s prior research directly motivates the development of a backdoor-aware test bed. In earlier work on FL for mining, we investigated how distributed training behaves under highly non-IID conditions typical of mining operations, where each site produces data that differs by geology, lighting, signage quality, dust levels, and camera placement. We developed a robust aggregation algorithm [8] capable of identifying and isolating malicious or unreliable model updates so that the global model is not corrupted during distributed training. We also developed the CAV-AD framework [6], an anomaly detection framework designed for autonomous vehicles experiencing sensor malfunctions or cyber-physical interference. CAV-AD monitors temporal patterns in perception outputs, evaluates cross-sensor consistency, and detects unexpected behavioral deviations caused by faulty or manipulated inputs. Together, these research efforts reveal a critical gap. Although we can detect malicious clients during FL and identify general anomalies in autonomous vehicles, there is currently no dedicated environment in which physical backdoor triggers can be systematically introduced, observed, and mitigated. There is also no controlled platform that integrates perception, backdoor detection, vehicle motion, sensor diversity, and federated learning into a single end-to-end system. The aforementioned research gaps motivate us to develop our proposed “Backdoor-Aware Autonomous Mining Test Bed(TB)”. This TB combines real vehicles, realistic tunnels, sensor modules, and federated learning across multiple nodes. The goal is to create a proving ground where physical




Fig. 8. Proposed backdoor-aware autonomous mining test bed integrating perception, backdoor-monitoring, autonomous vehicles, an edge server, and cloud federated learning.

backdoor attacks can be introduced and detected in real time, allowing us to design and validate safety mechanisms before deployment in operational mines. Fig 8 illustrates the proposed autonomous mining test bed (TB). This TB integrates camera-equipped mining vehicles, a perception and detection stack, a backdoor-monitor module and an edge-cloud federated learning pipeline. The architecture allows us to study, detect and mitigate real-time backdoor attacks caused by physical marks, stickers, or adversarial patterns placed on mine signage or tunnel surfaces

D. Label Flip Attack Detection Module

Future mines will rely heavily on FL to preserve data privacy, but they must also defend against non-IID degradation, malicious updates, unreliable clients, packet manipulation and gradient leakage. Building upon MineDetect’s [8] history-aware robust aggregation, next-generation FL pipelines must include trust scoring, anomaly-resistant filtering, cryptographic secure aggregation, and adaptive client selection. One of the current works from our lab [24] is to detect label flip attacks in autonomous mining operations. We have proposed a framework called TrustFed-LFD. This Framework comprises multiple mining robots (clients) operating autonomously and collaboratively training a shared global model through FL. As shown in Fig 9, some robots may be compromised and perform label flip attacks by manipulating class labels (e.g., from “TURN RIGHT” to “TURN LEFT”) to poison the model without altering the raw features. The framework comprises *seven* key steps, as illustrated in Fig 9, highlighted in circled brown. In step 1, parameter server $\mathcal{PS}$ distributes the global model to all mines. In step 2, each mine uses its local data to train its local model and then sends it back to $\mathcal{PS}$ for aggregation. However, each attacked mine launches a label flip attack before starting its local training, and in step 3, sends an adversarial gradient update $\Delta w_{C_{m_i}}^t$ to the $\mathcal{PS}$. TrustFed-LFD introduces robust, server-side enhancements in steps 4 through 7. Instead of aggregating all updates naively, the server performs gradient




Fig. 9. Overview of TrustFED-LFD.

analysis and trust scoring to detect and down-weight poisoned updates. DTN compatibility is also crucial as model updates and hazard information may arrive out of order or after long delays, requiring algorithms that maintain consistency despite asynchronous communication. Machine unlearning must also be hardened to prevent adversaries from selectively removing rare but safety-critical samples such as gas spikes or historical evacuation traces.

E. IoT-driven Equipment Health Monitoring Module

Underground mines use large, sophisticated equipment that runs nonstop and operates under intense mechanical, heat-related, and environmental strain. Load-haul-dump (LHD) vehicles, drill jumbos, bolters, crushers, conveyors, and battery-electric haul trucks experience degradation over time due to vibration, dust accumulation, moisture, corrosive chemicals, and variable thermal loading within confined underground tunnels. Even a single unexpected equipment failure can halt production, trigger hazardous situations, and cause costly downtime. Although modern mines deploy a variety of IoT sensors for monitoring vibration, pressure, temperature, and battery conditions, the data is often collected in isolation, without an integrated understanding of equipment health trends over time and thus their impacts on operational safety. This module comprises an Equipment Health and IoT Situational Awareness Test Bed (EH-IoT SA TB) aimed at unifying equipment-centric sensing and predictive maintenance models into a single coordinated framework. The architecture of TB is firmly grounded in established research on predictive maintenance, sensor fusion, and RL [25]. Based on the deep RL-based predictive maintenance framework, which integrates sensor-driven health indicators, severity ratings, and sequential decision-making, we adapt and generalize these concepts to the far more demanding conditions of underground mines. Unlike conventional industrial environments, underground settings impose severe communication constraints, non-stationary environmental stressors (dust, humidity, vibration), and rapid degradation cycles. Addressing these challenges requires a domain-specific

test bed capable of modeling equipment deterioration, anticipating hazardous operating states, and providing real-time situational awareness to both human operators and autonomous systems. By advancing these research directions, smart mining can transition from fragmented sensing and reactive safety toward resilient, predictive, and trustworthy intelligent systems capable of safeguarding miners under hazardous conditions.

IV. CONCLUSION AND FUTURE WORK

In this paper, we present a unified vision for developing resilient, safe and trustworthy autonomous mining ecosystems by integrating advances across sensing, perception, distributed intelligence, communication and predictive analytics. We have analyzed the key challenges inherent to underground and surface mining operations, including poor visibility, harsh sensing environments, unpredictable miner mobility, intermittent communication, non-IID data distributions, cyber-physical attacks, and energy-constrained sensor infrastructures. These challenges necessitate a fundamentally new architectural paradigm capable of supporting real-time awareness, robust decision-making and secure, privacy-preserving collaboration across all layers of mining operations. To address these needs, we have introduced the *Unified Smart Safety and Security Architecture*, a multilayer ecosystem that coordinates diverse sensing platforms, multimodal perception systems, DTN-based communication, secure FL, RL-driven navigation, and cloud analytics. The architecture incorporates five core modules, including Miner-Finder, Multimodal Situational Awareness, Backdoor Attack Monitor, TrustFED-LFD, and IoT-driven Equipment Health Monitoring, each designed to address specific gaps in miner localization, hazard interpretation, model integrity, federated robustness and equipment reliability. These modules collectively provide a cohesive intelligence fabric enabling early hazard detection, predictive response, safe navigation during disasters and protection against both cyber and physical threats.

This work will lay the foundation for next-generation smart mining systems capable of adapting to dynamic underground conditions, detecting malicious or faulty behavior, and delivering actionable insights to operators and autonomous agents. Future work will involve building fully operational test beds, deploying cross-site FL pipelines, validating DTN-aware navigation in real tunnels, implementing a multimodal situational awareness module and integrating energy-aware sensing to monitor equipment health. Following these directions, the mining industry can transition from reactive safety mechanisms to proactive, resilient and intelligent safety infrastructures to substantially improve operational reliability and protect human lives in hazardous environments.

REFERENCES

- [1] L. Rojas, Á. Peña, and J. García, "Ai-driven predictive maintenance in mining: a systematic literature review on fault detection, digital twins, and intelligent asset management," *Applied Sciences*, vol. 15, no. 6, p. 3337, 2025.
- [2] P. Zhao, X. Wang, S. Yu, X. Dong, B. Li, H. Wang, and G. Chen, "An open paradigm dataset for intelligent monitoring of underground drilling operations in coal mines," *Scientific Data*, vol. 12, no. 1, p. 780, 2025.

- [3] R. Liu, C. Li, H. Tang, Y. Ge, Y. Shan, and G. Li, "St-llm: Large language models are effective temporal learners," in *European Conference on Computer Vision*. Springer, 2024, pp. 1–18.
- [4] A. Goyal, S. Madria, and S. Frimpong, "Minerrouter: Effective message routing using contact-graphs and location prediction in underground mine," in *2024 25th IEEE International Conference on Mobile Data Management (MDM)*. IEEE, 2024, pp. 149–158.
- [5] M. R. Jewel, M. Elmahallawy, S. Madria, and S. Frimpong, "Dis-mine: instance segmentation for disaster-awareness in poor-light condition in underground mines," in *2024 IEEE International Conference on Big Data (BigData)*. IEEE, 2024, pp. 6279–6288.
- [6] M. S. Rahman, M. Elmahallawy, S. Madria, and S. Frimpong, "Cav-ad: A robust framework for detection of anomalous data and malicious sensors in cav networks," in *2024 IEEE 21st International Conference on Mobile Ad-Hoc and Smart Systems (MASS)*. IEEE, 2024, pp. 330–338.
- [7] Y. Sun, Z. Liu, J. Cui, J. Liu, K. Ma, and J. Liu, "Client-side gradient inversion attack in federated learning using secure aggregation," *IEEE Internet of Things Journal*, vol. 11, no. 17, pp. 28 774–28 786, 2024.
- [8] M. S. Rahman, M. Elmahallawy, S. Madria, and S. Frimpong, "Detecting untargeted attacks and mitigating unreliable updates in federated learning for underground mining operations," *arXiv preprint arXiv:2508.10212*, 2025.
- [9] A. Goyal, S. Madria, and S. Frimpong, "Ogle-mine: Obstacle-infused goal-conditioned learning for post-disaster navigation in underground mine," in *Proceedings of the 37th International Conference on Scalable Scientific Data Management*, 2025, pp. 1–12.
- [10] Z. Liu, T. Wang, M. Huai, and C. Miao, "Backdoor attacks via machine unlearning," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 38, no. 13, 2024, pp. 14 115–14 123.
- [11] Z. Liu, H. Ye, C. Chen, Y. Zheng, and K.-Y. Lam, "Threats, attacks, and defenses in machine unlearning: A survey," *IEEE Open Journal of the Computer Society*, 2025.
- [12] A. Goyal, S. Madria, and S. Frimpong, "Abstract: A dtn system for tracking miners using gae-lstm and contact graph routing in an underground mine," in *Proceedings of the Int'l ACM Symposium on Mobility Management and Wireless Access*, 2023, pp. 129–132.
- [13] M. R. Jewel, M. Elmahallawy, S. Madria, and S. Frimpong, "Explaining the unseen: Multimodal vision-language reasoning for situational awareness in underground mining disasters," *arXiv preprint arXiv:2512.09092*, 2025.
- [14] L. Li, Y. Li, X. Zhang, Y. He, J. Yang, B. Tian, Y. Ai, L. Li, A. Nüchter, and Z. Xuanyuan, "Embodied intelligence in mining: leveraging multimodal large language models for autonomous driving in mines," *IEEE Transactions on Intelligent Vehicles*, vol. 9, no. 5, pp. 4831–4834, 2024.
- [15] A. Gupta, T. Luo, M. V. Ngo, and S. K. Das, "Long-short history of gradients is all you need: Detecting malicious and unreliable clients in federated learning," in *European Symposium on Research in Computer Security*. Springer, 2022, pp. 445–465.
- [16] C. Yin and Q. Zeng, "Defending against data poisoning attack in federated learning with non-iid data," *IEEE Transactions on Computational Social Systems*, 2023.
- [17] X. Cao, M. Fang, J. Liu, and N. Z. Gong, "Fltrust: Byzantine-robust federated learning via trust bootstrapping," *arXiv preprint arXiv:2012.13995*, 2020.
- [18] Z. Wu, Q. Ling, T. Chen, and G. B. Giannakis, "Federated variance-reduced stochastic gradient descent with robustness to byzantine attacks," *IEEE Transactions on Signal Processing*, vol. 68, pp. 4583–4596, 2020.
- [19] H. Kim and Y. Choi, "Autonomous driving robot that drives and returns along a planned route in underground mines by recognizing road signs," *Applied Sciences*, vol. 11, no. 21, p. 10235, 2021.
- [20] M. Yadav, M. Elmahallawy, S. Madria, and S. Frimpong, "Predicting battery levels of sensor nodes using reinforcement learning in harsh underground mining environments," in *Proceedings of the 40th ACM/SIGAPP Symposium on Applied Computing*, 2025, pp. 2048–2057.
- [21] A. Goyal, S. Madria, and S. Frimpong, "Minerfinder: A gae-lstm method for predicting location of miners in underground mines," in *Proceedings of the 30th International Conference on Advances in Geographic Information Systems*, 2022, pp. 1–12.
- [22] J. Li, D. Li, S. Savarese, and S. Hoi, "Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models," in *International conference on machine learning*. PMLR, 2023, pp. 19 730–19 742.
- [23] Y. Chen, S. Shao, E. Huang, Y. Li, P.-Y. Chen, Z. Qin, and K. Ren, "Refine: Inversion-free backdoor defense via model reprogramming," *arXiv preprint arXiv:2502.18508*, 2025.
- [24] Wireless to cloud computing laboratory, missouri s&t. [Online]. Available: <https://web.mst.edu/cswebdb/about.html>
- [25] K. S. H. Ong, W. Wang, D. Niyato, and T. Friedrichs, "Deep-reinforcement-learning-based predictive maintenance model for effective resource management in industrial iot," *IEEE Internet of Things Journal*, vol. 9, no. 7, pp. 5173–5188, 2021.

2026 18th International Conference on COMmunication Systems and NETworks (COMSNETS) took place 6-10 January 2026 in Bengaluru, India.



IEEE catalog number: CFP2631F-ART

ISBN: 979-8-3315-9239-4

ISSN: 2155-2509