

AI and IoT-Enabled Safety Monitoring for Mines: A Comparative Review of Gas Detection, Wearable Sensing and Integrated Smart-Mining Frameworks

Shivani Nandu Shinde¹, Bangar Abhishek Popat², Bhosale Sachin Bajirao³

^{1,2,3}Department of Computer Engineering, Jaihind College of Engineering, SPPU, Pune, India

Abstract- Underground mines are hard places to keep safe. The air can turn dangerous without any smell or warning sign, communication is patchy, and help is usually a long way off. Over the last few years many research groups have tried to close this gap by combining low-cost sensors, wireless links and machine learning. This paper reads four recent studies from 2026 side by side: a LoRa-based gas monitoring node with Random Forest analysis, a wearable wristband and belt hub with edge intelligence, a Random Forest gas classifier with a Python dashboard and cloud alerts, and a large integrated smart-mining case that joins electrification, safety analytics, biodiversity and accounting. We compare their sensing, communication, analytics and alerting layers, and we recompute several figures directly from the numbers the authors report. Our reading is that the hardware ideas are sound and increasingly affordable, but the evidence behind them is thin. Datasets are rarely described, labels are often derived from the same thresholds the model is meant to improve on, and underground field trials are scarce. We close with a layered reference framework and a short list of practical recommendations for future prototypes.

Keywords— Mine Safety; Internet of Things; LoRa; Machine Learning; Random Forest; Edge Computing; Wearable Sensing; Gas Detection; Hazard Prediction.

I. INTRODUCTION

1. Background

Anyone who has read an accident report from an underground coal mine will recognise the pattern. A gas builds up quietly, ventilation is not quite enough, and by the time somebody notices, the options have already narrowed. Methane and carbon monoxide are the usual culprits. Neither can be smelled, and both can climb to dangerous levels faster than a periodic manual check can catch [1], [3]. Traditional monitoring leans on fixed stations, handheld detectors and fixed alarm limits. That approach is simple and dependable, but it says very little about where a particular worker is, how that worker is coping, or whether a slow upward trend is about to become an emergency [2].

Cheap microcontrollers, small sensors and wireless modules have changed what a college team or a small operator can build. A review of IoT deployments in underground mines described how environmental, safety and production monitoring were moving onto connected infrastructure [5], and

later work put sensors on the miners themselves, using bracelets and lamps to follow vital signs, gas and position [6]. Long-range, low-power radios such as LoRa have been tried because Wi-Fi and Bluetooth fade quickly in tunnels [7]. Small neural networks have been squeezed onto ESP32 boards for fall detection [8], combined environmental and health monitoring has been proposed for mines [9], and even the humble MQ-series gas sensors have been tested on ESP32-class boards for air monitoring [10]. Machine learning sits on top of all this. Earlier mine-safety studies showed that a random-forest model fed with monitoring data can raise a gas-explosion warning [11], and that multi-sensor forecasting can support safety warnings at a longwall face [12]. In maintenance, artificial intelligence and fault diagnosis are now treated as core ingredients of predictive programmes [13].

2. Motivation and Scope

Reading the recent literature, we noticed something frustrating. Almost every paper promises an intelligent, real-time and reliable system, yet the papers are hard to compare. They use different

sensors, different radios and different definitions of a “hazard”, and they report results in different ways, or sometimes not at all. A reader who wants to build a prototype has to dig through the text to learn what was actually measured.

This paper therefore does something modest. It takes four 2026 studies [1]–[4], reads them side by side and asks what each one really shows. The first is a LoRa-based coal-mine gas monitor [1]. The second is a wearable system called MinePulse AI [2]. The third is a Random Forest gas classifier with email and cloud alerts [3]. The fourth is a large smart-mining case that ties safety analytics to electrification, ecology and accounting [4].

Contributions: we (i) compare the four systems layer by layer, (ii) recompute several metrics from the counts and values the authors publish, (iii) point out where the evidence is strong and where it is not, and (iv) propose a layered reference framework with practical recommendations. We did not build hardware or run new experiments, and every number below is either quoted from a source study or calculated from one.

Paper layout (Fourth Level). Section II summarises the source studies, Section III explains how we compared them, Section IV sets out their architectures, Section V presents the reported and recomputed results, Section VI discusses them, Section VII proposes the framework, and Sections VIII and IX cover limitations and conclusions.

II. SOURCE STUDIES AND RELATED WORK

1. LoRa gas-monitoring node with Random Forest analysis

The first study, from a computer science department in Tamil Nadu, builds an odourless-gas monitoring system around an ESP32 board with a LoRa module [1]. The sensing unit combines MQ-series gas sensors for methane and carbon monoxide, a temperature and humidity sensor, a flame sensor and a GPS module. Readings are first compared with predefined limits on the microcontroller itself, so a basic alarm never has to wait for the network. A

Random Forest model then looks at real-time and historical data and labels conditions as safe or hazardous. Local alerts use a buzzer, LEDs and a display, supervisors receive SMS or call notifications, and a LoRa gateway forwards data to a cloud dashboard. The authors argue for LoRa on range and power grounds and note that Wi-Fi and Zigbee struggle with attenuation underground, a worry shared by earlier work on LoRa-based industrial networks in coal mines [14].

What the paper reports is mostly qualitative. Its performance table lists one observed reading per monitored parameter and the system’s reaction. A methane reading of 266 ppm against a normal range of 0–250 ppm triggered an alert and alarm. A carbon monoxide reading of 265 ppm against 0–30 ppm produced a warning. A temperature above 50 °C and a detected flame also raised alerts, and the response-time target was under three seconds. No accuracy, precision, recall or false-alarm figures are given for the Random Forest, and the training data is not described.

2. Wearable system with edge analytics

The second study, from a government polytechnic in Maharashtra, moves the sensing from the tunnel wall to the person [2]. A wristband carries a MAX30102 pulse sensor for heart rate. A belt-mounted hub built on an ESP32-WROOM-32 gathers motion data from an MPU6050, ambient readings from a DHT11 and an MQ-2 gas sensor, position from a NEO-6M GPS module and, optionally, images from an ESP32-CAM. Everything is processed on the hub. The analytics layer is deliberately light. A rule-based classifier assigns SAFE, WARNING or DANGER, an Isolation Forest looks for unusual combinations of sensor values, a weighted score folds five sensors into one risk index, and a rolling window of five readings indicates whether risk is rising or falling. Results travel over Wi-Fi to an ASP.NET dashboard that uses Chart.js and an embedded map. It is telling that the authors themselves list LoRa as a future replacement for Wi-Fi in deeper zones.

The evaluation is a prototype demonstration. The dashboard screenshot shows a heart rate of 126 BPM, 24.2 °C, 51 % humidity, a low gas value and a

movement flag, with the AI panel reading “unusual activity”. The text says tests confirmed detection of falls and abnormal health conditions and improved response times, but no counts, latencies or error rates are given.

3. Random Forest gas classifier with cloud alerts

The third study, from an engineering institute in Tamil Nadu, is the most conventional machine-learning paper of the four [3]. An Arduino Nano reads a DHT11 and gas sensors for CO, CO₂ and CH₄ and passes the values over a serial link to a Python application. A Random Forest classifier sorts each sample into Normal, Warning or Critical. Training samples were labelled from predefined safety conditions and domain knowledge, then cleaned for noise, missing values and scale. The dashboard shows live readings and can switch on a fan and a buzzer. The system also sends voice and email alerts and uploads data to a ThingSpeak cloud channel.

This is the only paper of the four that publishes enough to recompute a metric. Its confusion matrix reports 29 of 30 Normal samples correct (one predicted as Warning), all 180 Warning samples correct and all 990 Critical samples correct, which the authors summarise as roughly 99.9 % accuracy. They also admit that noisy or incomplete sensor readings affected data quality during testing.

4. Integrated smart-mining case

The fourth study works at a different scale [4]. Instead of a single device it treats an anonymised critical-mineral mine as one joint engineering, safety, ecological and financial system, and reports paired results for a conventional baseline and a smart-mining condition. Energy use per tonne falls from 11.8 to 9.2 kWh/t, haulage operating cost falls by 32.9 %, electrical fault events fall by 37.9 %, the useful warning lead time for equipment faults grows from 0.8 to 4.6 hours, and recordable incident frequency drops from 2.8 to 1.9 per 200,000 work hours. Seven classifiers for safety risk are compared, from logistic regression to a multi-source fusion ensemble. The financial case rests on a US\$48.0 million investment, US\$16.2 million of recurring annual benefit and an NPV of about US\$30.9 million over seven years at a 10 % discount rate. The authors are frank that the

mine’s identity, observation period, sample counts, raw telemetry and model-training records were not available to them, and they treat their statistical relationships as hypotheses to be tested rather than findings.

5. Wider context

Several threads from the wider literature are worth keeping beside these four. Edge architectures for intelligent coal mining try to push analysis close to the face so that decisions do not depend on a distant server [15]. AI-assisted safety helmets with embedded processors point the same way for individual workers [16]. On the engineering side, studies of battery-trolley trucks and alternative haul-truck powertrains show that energy use and cost depend heavily on how the electrical system is designed [17], [18], and work on reclamation liabilities reminds us that safety and environmental performance eventually reach the accounts [19]. None of this is the focus of our paper, but it explains why the integrated study [4] looks so different from the three device-level ones.

III. REVIEW METHODOLOGY

1. Selection and comparison dimensions

We worked with four journal papers dated between March and August 2026 that deal with AI-assisted safety monitoring in mines. They were not drawn from a systematic database search, so this is a comparative reading and not a formal systematic review, and our conclusions should be taken in that spirit.

Each study was read against six dimensions: (i) the hazards and parameters sensed, (ii) the hardware platform, (iii) the communication path, (iv) the analytics method, (v) alerting and visualisation, and (vi) the quality and quantity of evidence reported. The sixth turned out to matter most. A system can have an impressive block diagram and still tell the reader almost nothing about whether it works.

2. Handling of reported numbers

We did not rerun anything. Where an author reported counts, we recomputed standard metrics from them. Where an author reported a change

between a baseline and a new value, we checked the arithmetic. Everything else is quoted as reported and flagged as such. For a binary or per-class decision, precision, recall and F1-score are computed from true positives (TP), false positives (FP) and false negatives (FN) as

$$Precision = TP / (TP + FP) \quad (1)$$

$$Recall = TP / (TP + FN) \quad (2)$$

$$F1 = 2 \cdot Precision \cdot Recall / (Precision + Recall) \quad (3)$$

Percentage change between a baseline value and a new value is calculated as

$$\Delta\% = ((new - base) / base) \times 100 \quad (4)$$

Reading the evidence: a high accuracy figure only means something when the reader knows how many samples were tested, how they were labelled and how common each class was. We therefore treat any headline accuracy without those details as a claim to be checked, not as a result.

IV. ARCHITECTURES COMPARED

1. A shared four-layer pattern

Despite their differences, the four systems follow nearly the same pattern, shown in Fig. 1. Sensors feed an edge node, the edge node sends data over some kind of link, an analytics stage decides how worrying the situation is, and an alert and display stage tells people. The differences lie in what is plugged into each layer and in how much of the pipeline the authors actually evaluate. Table 1 summarises the platforms and methods, and Table 2 summarises what each paper reports.




Fig. 1. Common layered architecture seen across the four source studies.

Table 1. Platforms and methods of the four source studies.

Study	Platform	Sensing	Link	Analytics
[1]	ESP32 + LoRa	MQ-2, MQ-7, MQ-135, methane, flame, temperature, humidity, GPS	LoRa, gateway to cloud	Thresholds + Random Forest
[2]	ESP32-WROOM-3 2 hub + wristband	MAX30102, MPU6050, DHT11, MQ-2, GPS, camera	Wi-Fi to ASP.NET dashboard	Rules, Isolation Forest, weighted risk score, trend
[3]	Arduino Nano + PC	DHT11, CO, CO2, CH4 sensors	Serial cable, ThingSpeak cloud	Random Forest (3 classes)
[4]	Mine-wide systems (SCADA, CMMS, ledgers)	Electrical, equipment, safety, ecology, finance	Not detailed	Seven classifiers, composite indices

2. Sensing

All four rely on affordable sensing, and three of the four lean on MQ-series gas modules or similar low-cost parts [1], [2], [3]. That keeps cost down, which matters for small mines, but it has a price. As far as we know, MQ-type sensors respond to more than one gas, drift with temperature and humidity, and need careful calibration, and none of the three

papers describes a calibration procedure. The MinePulse paper describes its MQ-2 as detecting combustible gases and smoke in general [2], which is useful as an early warning but is not the same as measuring methane. The only study that adds a human signal is the wearable one, which uses heart rate and motion to infer stress and falls [2].

Table 2. Evidence reported by each source study.

Study	Quantitative model metrics	Dataset described	Test setting stated
[1]	None for the Random Forest; one observed reading per parameter	No	"Different environmental conditions"
[2]	None	No	Prototype demonstration
[3]	Confusion matrix (1,200 samples), accuracy about 99.9 %	Features and classes described; labelling rule described in general terms	Real-time data plus labelled historical samples
[4]	Accuracy, precision, recall, F1, ROC-AUC for seven models	No; authors state records were unavailable	Anonymised mine case summary

3. Communication

The communication choices are the most revealing. The LoRa node [1] chose its radio for underground range and low power. The wearable system [2] uses Wi-Fi even though the authors later suggest LoRa for deeper zones. The gas classifier [3] sends data to a computer over a serial cable, and lists a wireless upgrade such as Zigbee, LoRa or 5G as future work. In other words, the link layer is the least finished part of two of the four prototypes, and it is the layer that matters most once a system leaves the bench and goes underground.

4. Analytics

Random Forest appears twice [1], [3], which is not surprising. It copes with noisy tabular readings, needs little tuning and is cheap to run. The wearable study adds an Isolation Forest, which suits the "something odd is happening" question better than a fixed label [2]. The integrated study goes furthest, comparing seven learners and a fusion ensemble that combines domain-specific risk probabilities [4]. Written generically, a fusion of this kind is a weighted average of the probabilities from the individual domains, and the authors say that the fitted weights

should be estimated from validation performance and recalibrated as conditions change [4].

5. Alerting and visualisation

Alerting is where the studies are most alike. Buzzers, LEDs, small displays, SMS, email and web dashboards appear throughout. What differs is the idea of what an alert is for. The integrated study stresses that an alert should trigger a defined response, such as isolation, inspection, speed restriction or escalation, and that an accurate score which triggers nothing is not a safety intervention [4]. The three device-level papers describe alerts as outputs but say much less about who acts on them and how fast.

V. RESULTS REPORTED AND RECOMPUTED

1. Per-class performance of the Random Forest classifier

The confusion matrix published for the gas classifier [3] is shown in Fig. 2. From the counts, 1,199 of 1,200 samples were classified correctly, which gives an accuracy of 99.92 % and agrees with the roughly 99.9 % the authors state. Table 3 breaks this down by class using Eqs. (1) to (3).




Fig. 2. Confusion matrix reported for the Random Forest gas classifier [3].

Table 3. Per-class metrics recomputed from the counts reported in [3].

Class	Support	Correct	Precision	Recall	F1-score
Normal	30	29	1.000	0.967	0.983
Warning	180	180	0.994	1.000	0.997
Critical	990	990	1.000	1.000	1.000
Overall	1,200	1,199	Accuracy 0.9992	Macro-F1 0.993	

Two things stand out. First, the test set is heavily skewed: 82.5 % of the samples are Critical and only 2.5 % are Normal. A model that always answered "Critical" would already score 82.5 % accuracy, so the headline number needs this context.

Second, the single error is a Normal sample called Warning, which is the safe direction to be wrong in for a hazard monitor. Nothing hazardous was missed in this test. Because the paper does not describe how the test samples were split from the training data, we cannot judge how well these results would carry over to a new mine.

2. Threshold exceedance in the LoRa node

The performance table of the LoRa node [1] lists observed readings next to normal ranges. Fig. 3 expresses each reading as a multiple of its upper normal limit. Methane at 266 ppm sits only about 6 % above the 250 ppm limit, so the alarm there depends on how tightly the threshold is set.

Carbon monoxide at 265 ppm is almost nine times the stated 30 ppm limit, a far more serious excursion that the paper labels a warning alert. The temperature entry (above 50 °C against 20–40 °C) is given only as a lower bound. These are single demonstration readings, not a trial, but they do show the system reacting across several parameters.




Fig. 3. Observed readings as a multiple of the upper normal limit, computed from the values in [1]. Temperature is a lower bound because the source gives only "above 50 °C".

3. Comparison of seven safety-risk models

The integrated study reports accuracy, precision, recall, F1-score and ROC-AUC for seven models [4]. We reproduce them in Table 4 and plot the F1-scores in Fig. 4. The ordering is consistent and plausible: logistic regression is lowest, the neural network sits in the middle, tree ensembles and the LSTM are above 92 %, XGBoost reaches 95.1 %, and the multi-source fusion ensemble is best with an F1 of 96.3 % and ROC-AUC of 0.985. We checked the F1 column against the published precision and recall values using Eq. (3) and found them consistent to within rounding.

Table 4. Performance of the seven safety-risk models as reported in [4].

Model	Accuracy	Precision	Recall	F1	ROC-AUC
Logistic regression	82.4	81.1	79.8	80.4	0.860
Support vector machine	86.9	86.3	84.8	85.5	0.900
Artificial neural network	90.8	90.2	89.6	89.9	0.940
Random forest	93.4	92.8	94.1	93.4	0.967
LSTM	92.6	92.1	93.5	92.8	0.961
XGBoost	95.1	94.7	95.6	95.1	0.978
Multi-source fusion ensemble	96.3	95.8	96.9	96.3	0.985




Fig. 4. F1-score of the seven safety-risk models (LR: logistic regression, RF: random forest, XGB: XGBoost), data from [4].

The gap from logistic regression to the fusion ensemble is 15.9 F1 points, but the gap from XGBoost to fusion is only 1.2 points. That is a useful reminder that a well-tuned boosted-tree model captures most of the gain, and that the extra benefit of fusing many data sources has to be weighed against the cost of collecting and maintaining them. The authors also point out that the raw test set, class distribution and calibration curves were not available, so these are summary results and not independently reproduced estimates [4].

4. Operational outcomes in the integrated case

The integrated study also translates model quality into operational outcomes, which is something the three device-level papers do not attempt. Table 5 lists selected indicators from [4], with the percentage changes recomputed using Eq. (4). Our values match the authors' figures to one decimal place.

Table 5. Selected operational indicators reported in [4], with changes recomputed by us.

Indicator	Baseline	Smart mining	Change
Energy use per tonne	11.8 kWh/t	9.2 kWh/t	−22.0 %
Haulage operating cost	US\$3.40/t	US\$2.28/t	−32.9 %
Electrical fault events	5.8 per 1,000 h	3.6 per 1,000 h	−37.9 %
Mean time between failures	186 h	232 h	+24.7 %
Planned maintenance share	54 %	72 %	+18 points
Equipment fault warning lead time	0.8 h	4.6 h	+3.8 h

Recordable incident frequency (per 200,000 h)	2.8	1.9	−32.1 %
Hazard warning lead time	6 min	34 min	+28 min
Timely response to critical alerts	71 %	92 %	+21 points

The lead-time rows are the most interesting to us. A model that flags a fault 0.8 hours ahead gives little room to act, while a warning 4.6 hours ahead lets a crew isolate equipment, reroute production and stage spare parts. The same logic applies to the hazard warning, which grows from 6 to 34 minutes. The study also reports that timely response to critical alerts rose from 71 % to 92 %, which suggests that the improvement came from linking alerts to procedures as well as from better models [4]. Two entries in its reliability table, failure-detection recall of 92.8 % and a false-alarm rate of 6.4 %, are labelled as targets rather than measured results, and we have not treated them as observed values.

VI. DISCUSSION

1. Where the four studies agree

Taken together, the studies agree on a few points that we find convincing. Multi-parameter sensing beats a single gas channel, because temperature, humidity, motion and other signals add context to a gas reading. Processing at or near the edge is worth the effort, since a local alarm still works when the network does not. Tree-based models are a sensible default for noisy sensor tables. And alert design deserves as much thought as model design.

2. Where the evidence is thin

Circular labels: in the gas classifier, samples were labelled from predefined safety conditions and domain knowledge [3]. If the labels come from threshold rules on the same readings that are fed to the model, then a Random Forest can reach near-perfect accuracy simply by learning those rules. That is a fair engineering outcome, but it does not show that the model predicts hazards earlier than a threshold would. This is our interpretation of the published description, and the authors may have

used additional information that the paper does not mention.

Class imbalance: as shown in Section V, the majority class makes up 82.5 % of the test samples. Reporting recall for the rare Normal class, or using a held-out balanced test set, would tell the reader much more than overall accuracy.

Missing operating data: none of the three device-level papers reports a false-alarm rate, a detection latency measured end to end, battery life, or packet loss, although these are exactly the numbers a mine operator would ask for first. The LoRa node states a response time target below three seconds [1], but we could not find how it was measured.

Summary-level results: the integrated study is candid about what it lacks [4]. Its numbers are internally consistent, which is reassuring, but they cannot be reproduced from the paper. Readers should treat the 96.3 % F1 and the financial figures as case results and not as general expectations.

3. From devices to systems

The device-level papers show what is possible with a few hundred rupees of parts. The integrated paper shows what it takes to make such devices matter to an organisation: consistent identifiers and timestamps across data sources, audit trails for model-driven decisions, and a rule that critical thresholds can override a composite score [4]. We think the best future work will combine the two, taking the low-cost, edge-first hardware of the student-scale prototypes and wrapping it in the data discipline described in the larger study.

4. Worker data and privacy

Wearables add a sensitive dimension. Heart rate, movement and location describe a person, not just an environment. The MinePulse authors themselves suggest federated learning as a way to improve on-device models without centralising individual miner data [2]. Whatever technical route is chosen, a real deployment needs clear rules on who can see a worker's health data, how long it is kept, and what it may be used for. These questions are easy to overlook in a prototype and hard to retrofit later.

VII. REFERENCE FRAMEWORK AND RECOMMENDATIONS

1. Layered framework

Putting the strengths of the four studies together, we propose the six-layer framework in Fig. 5. It is a design sketch, not a validated system. Layers 1 to 5 follow the common pattern in Fig. 1, and layer 6 reflects the governance practices that the integrated study argues for.




Fig. 5. Proposed six-layer reference framework for AI and IoT-based mine safety monitoring.

At the analytics layer, we suggest keeping the weighted-score idea used by the wearable system [2] and the fusion idea from [4] in a single, easily audited form. If s_k is a normalised score between 0 and 1 for sensor or domain k , the overall risk index is

$$R = \sum_k w_k s_k, \text{ with } \sum_k w_k = 1 \quad (5)$$

The weights w_k should be fitted on validation data and revisited when equipment, ventilation or geology change, rather than being fixed by hand once. A hard threshold on any single critical gas should always be able to override R , so that a high methane reading can never be averaged away by calm readings elsewhere.

2. Practical recommendations

Table 6 turns the gaps we found into concrete practices for anyone designing the next prototype.

Table 6. Gaps found in the source studies and the practice we recommend.

Gap observed	Recommended practice
Gas sensors with no calibration described	Calibrate against a reference gas, log drift, and report sensor health to the dashboard
Labels derived from the same thresholds as the features	Use independent labels (incident logs, expert review) or state clearly that the model reproduces the rules
Accuracy reported on a skewed test set	Report per-class recall, precision and a majority-class baseline
No false-alarm or latency figures	Measure end-to-end latency and false alarms per shift, and publish them
Wi-Fi or serial links in prototypes	Test the radio underground, with LoRa or a mesh, and report packet loss
Alerts without defined actions	Map each alert level to a named response and an owner
Wearable health data	Minimise collection, process on the device, and set retention limits

Limitations and Future Work

This review has clear limits. It covers only four studies, all from 2026 and all supplied to us as papers rather than found through a systematic search, so important work may be missing. We relied on the authors' descriptions and could not inspect code, datasets or hardware. Two of the papers give no quantitative model results at all, which limits any numerical comparison. Our recomputed values are only as good as the counts and values the authors published, and our interpretation of how the gas classifier was labelled rests on a short description.

Future work should build and test a prototype of the proposed framework. A sensible plan would be a small edge node with calibrated gas sensors, a LoRa link tested in a real tunnel or a representative environment, a Random Forest with an Isolation Forest running side by side, and a dashboard that records every alert together with the response it triggered. The evaluation should use independently labelled data, report per-class metrics and false-alarm rates, and measure latency from sensor to

alarm. Longer term, deep-learning models for early prediction, camera-based hazard recognition and cooperative learning across wearable units, as the MinePulse authors suggest [2], are all worth exploring once the basics are measured properly.

VIII. CONCLUSION

Reading these four 2026 studies together gives a hopeful but cautious picture. On the hopeful side, a modest budget now buys a multi-gas sensing node, a long-range radio, a local alarm and a cloud dashboard, and tree-based models run comfortably on small hardware. The integrated mining case shows that these ideas can be tied to real outcomes, including longer warning times, fewer incidents and a positive financial case, when alerts are connected to procedures and data is governed carefully.

On the cautious side, the evidence in the smaller studies does not yet match the confidence of their conclusions. Datasets are undescribed, labels may be circular, metrics are missing, and links are often tested above ground only. None of this makes the work unimportant. It means the next step for the field is less about adding another sensor and more about measuring what already exists honestly: independent labels, per-class metrics, false-alarm rates, end-to-end latency and underground trials. We hope the comparison and the reference framework in this paper help a future prototype start from there.

REFERENCES

1. Prakash, B.S., Pradeepa, B., Priya, D., Reenasri, V. "Smart IoT based odourless gas detection safety monitoring and alerting system for coal mines using LoRa and AI". International Journal of Creative Research Thoughts (IJCRT) 14(3) (2026).
2. Kottawadla, A.S., Bansod, S.I., Lande, M.V., Lodelliwar, M.J., Akanpalliwar, O.M., Vairagade, T.B., Patil, V.B., Manalwar, M.B. "MinePulse AI: Advanced mine safety and hazard prediction system". International Research Journal of Modernization in Engineering Technology and Science 8(3), 4567–4573 (2026).

3. Kavitha, V., Sriram, B., Vignesh Kumar, J., Yuvaraj, R.B. "Smart mine IoT & AI for proactive coal mine safety and hazard control". *International Journal of Engineering & Extended Technologies Research* 8(2), 2673–2679 (2026).
4. Idoko, I.P., Ayoola, V.B., Kalejaiye, A. "AI enabled smart mining for critical minerals integrating mine electrification, predictive safety, biodiversity monitoring and nature related financial accounting". *International Journal of Modern Science and Research Technology* 4(8), 386–404 (2026).
5. Zhang, H. et al. "Recent advancements in IoT implementation for environmental, safety, and production monitoring in underground mines". *IEEE Internet of Things Journal* 10(16), 14507–14526 (2023).
6. Li, Z. et al. "Real-time monitoring of underground miners' status based on mine IoT system". *Sensors* 24(3), 739 (2024).
7. Joy, S., Bale, A.S., Baby Chithra, R. "IoT based enhanced safety monitoring system for underground coal mines using LoRa technology". *Lecture Notes in Networks and Systems* 1029, Springer (2024).
8. Dhanush, J., Asha Rani, M. "Real-time human fall detection and alert system using autonomous ESP32-WROOM-32 and MPU6050". *IJISRT* 10(9) (2025).
9. Jones, D.R. et al. "IoT-based real-time environmental and health monitoring for underground mines". *IEEE Sensors Journal* 21(10), 11550–11562 (2021).
10. Easterline, L.M. et al. "Smart air monitoring with IoT-based MQ-2, MQ-7, MQ-8, and MQ-135 sensors using NodeMCU ESP32". *Procedia Computer Science* 245, 815–824 (2024).
11. Li, H., Zhang, Y., Yang, W. "Gas explosion early warning method in coal mines by intelligent mining system and multivariate data analysis". *PLOS ONE* 18(11), e0293814 (2023).
12. Chang, H., Meng, X., Wang, X., Hu, Z. "Research on coal mine longwall face gas state analysis and safety warning strategy based on multi-sensor forecasting models". *Scientific Reports* 14, 13795 (2024).
13. Dayo-Olupona, O., Genc, B., Celik, T., Bada, S. "Adoptable approaches to predictive maintenance in mining industry: An overview". *Resources Policy* 86, 104291 (2023).
14. Chowdhury, A.R., Bakshi, S.C., Pramanik, A., Roy, G.C. "Design and study of LoRa-based IIoT network for underground coal mine environment". *IEEE Access* (2024).
15. Bing, Z., Wang, X., Dong, Z., Dong, L., He, T. "A novel edge computing architecture for intelligent coal mining system". *Wireless Networks* 29, 1545–1554 (2022).
16. K, L., G, R., M, S. "AI-based safety helmet for mining workers using IoT technology and ARM Cortex-M". *IEEE Sensors Journal* 23, 21355–21362 (2023).
17. Bao, H., Knights, P., Kizil, M., Nehring, M. "Energy consumption and battery size of battery trolley electric trucks in surface mines". *Energies* 17(6), 1494 (2024).
18. Feng, Y., Liu, Q., Li, Y., Yang, J., Dong, Z. "Energy efficiency and CO₂ emission comparison of alternative powertrain solutions for mining haul truck". *Journal of Cleaner Production* 370, 133568 (2022).
19. Espinoza, R.D., Morris, J.W.F. "Towards sustainable mining (part II): Accounting for mine reclamation and post reclamation care liabilities". *Resources Policy* 52, 29–38 (2017).
20. J. Zhu, J. Peng, and L. Wang, "A lightweight and trustworthy authentication protocol for UAV networks based on PUF," in *Proceedings of the 14th International Conference on Communication and Network Security*, Dec. 2024, pp. 71–77.
21. S. U. Jan, F. Qayum, and H. U. Khan, "Design and analysis of lightweight authentication protocol for securing IoD," *IEEE Access*, vol. 9, pp. 69287–69306, 2021.