Journal of Risk, Insurance, and Safety
DOI: 10.18311/jris.1.1.2026.66-84
ISSN: XXXX-XXXX | Vol 1(1), 2026 | pp. 66-84

The Gap Between the Complexity of Accident Causation and Compliance-Based Investigation Practices in Indian Mines

Ashish Kumar*[1], Amit Sharma[2] and Ashis Bhattacherjee[3]

[1]*School of Advanced Engineering, University of Petroleum and Energy Studies, Dehradun – 248007, Uttarakhand, India; aksaini4@gmail.com*
[2]*Department of Mining Engineering, National Institute of Technology, Surathkal - 575025, Karnataka, India*
[3]*Department of Mining Engineering, Indian Institute of Technology Kharagpur - 721302, West Bengal, India*

Abstract

Accident investigations in high-risk industries are intended to support organisational learning by identifying the systemic conditions that contribute to adverse events. Meaningful accident investigation requires moving beyond isolated procedural failures to examine how systemic structures and constraints shape unsafe outcomes. This study examines whether regulatory accident investigations reflect such a multi-level understanding or whether they create an illusion of causality by emphasising a narrow set of apparent sources, particularly violations of Standard Operating Procedures (SOP). Using official fatal accident narratives published by the Directorate General of Mines Safety (DGMS), India, a Bidirectional Encoder Representations from Transformers (BERT)- based topic modelling approach was applied to two textual components: (i) the original fatal accident narratives and (ii) the corresponding recommendations or counterfactual statements issued by the regulator. Topic modelling of the accident narratives revealed 13 distinct themes, including transportation-related accidents, roof collapses, explosions, equipment failures, and drowning. In contrast, the analysis of DGMS recommendations identified only three dominant topics, two of which were heavily centred on regulatory compliance with the Coal Mine Regulations and the Central Electricity Authority Regulations, together accounting for 91.25% of all cases. These findings demonstrate a clear mismatch between the multi-level complexity of accident causation and the compliance-focused framing adopted in regulatory investigations. The reliance on SOP violations as primary causal explanations may risk obscuring deeper systemic vulnerabilities. An explicit system-based approach is therefore essential to avoid the illusion of causality created by oversimplified explanations and support more effective learning and prevention in high-risk socio-technical systems.

Keywords: *Accident Investigation, BERT Topic Modelling, DGMS, SOP Deviation, System Safety Causality Illusion*

Received: 24.02.2026; **Revised:** 12.08.2026; **Accepted:** 13.08.2026

1. Introduction

Among high-risk industries, mining remains one of the most hazardous industries worldwide due to its inherent exposure to numerous hazards and complex human-machine interactions[1]. Despite the implementation of occupational safety regulations and ongoing efforts to improve mining safety, fatal and serious accidents continue to occur in both developed and developing mining economies[2,3]. A comprehensive scoping review confirmed that mining accounts for approximately 8% of all global fatal occupational injuries, with injury rates among miners consistently ranging between 4.4% and 10.6% across studied populations[2,4,5]. In India, coal remains the backbone of the nation's energy infrastructure, accounting for roughly 70% of total power generation and expected to remain unaltered[6]. This reliance on coal places enormous operational pressure on the mining sector and introduces implicit safety implications. The DGMS, a statutory authority

*Author for correspondence

in India responsible for mine inspections, accident investigations, enforcement of safety standards, and regulatory guidance, reports a total of 226 fatalities in coal and lignite mines across the country between 2020 and 2024[7]. These figures highlight that regulatory control led to relative safety improvements; however, fatal accidents in Indian mines continue at a rate that demands a fundamental re-examination of how the industry understands, investigates, and learns from adverse events. The persistent recurrence of fatalities across diverse operational contexts suggests that the root of the problem may lie not merely in the SOP violation, but in deeper systemic issues.

Accidents in socio-technical systems are emergent outcomes produced by the interaction of technical, organisational, human, and environmental elements; they cannot be reduced to a single "root cause". Evidence from systems-oriented safety research demonstrates that failures often result from gradual shifts and interactions across system levels rather than isolated operator error[8]. Treating deviations from SOPs as the root cause is a reductionist approach that obscures upstream influences, management decisions, resource constraints, production pressures, design of controls, and latent organisational conditions that make non-compliance likely or even rational from workers' perspectives. Systemic accident causation models explicitly link frontline workers' behaviours to these upstream controls and show how organisational structures and feedback loops generate conditions for deviation[9]. When investigations adopt a narrow causal view that centres on SOP deviation, corrective actions typically emphasise retraining, discipline, or re-issuing procedures. Such responses address only proximal behaviours and leave deeper systemic issues intact. Literature on normalisation of deviance and "drift into failure" shows how organisational practices and incentives can normalise risky adaptations over time, making procedural fixes alone insufficient[10].

The reliance on SOP non-compliance as the primary cause of accidents is increasingly recognised as inadequate for understanding accidents in modern socio-technical systems. Rules and procedures are abstractions that cannot fully anticipate the variability of real-world operations. As argued by Dekker, work-as-imagined by designers and regulators often diverges from work-as-done by practitioners operating under dynamic conditions[11]. When accidents occur, investigations that take work-as-imagined as a reference point tend to judge deviations as failures, rather than as adaptive responses to systemic constraints. Moreover, compliance-based explanations can inadvertently reinforce blame-oriented cultures, discouraging open reporting and learning. By attributing accidents to individual or local procedural failures, organisations may avoid confronting uncomfortable questions about resource allocation, production incentives, staffing levels, or regulatory effectiveness. This phenomenon has been documented across multiple high-risk domains, including aviation, healthcare, and offshore energy, where formal compliance has coexisted with persistent safety problems[12,13].

Despite the potential offered by advanced analytical methods, the current framework for accident investigation and response in India's mining sector remains firmly rooted in a compliance-based paradigm. Under the Mines Act, 1952, and the Coal Mines Regulations, 2017, DGMS is mandated to conduct enquiries into fatal and serious accidents, identify the factors that led to the incident, and recommend corrective measures to prevent recurrence. However, the nature of these enquiries and the recommendations that follow tend to centre narrowly on whether established SOPs and statutory regulations were adhered to at the time of the accident. When violations are identified, whether related to the Coal Mine Regulations, the Central Electricity Authority Regulations, or other prescribed safety standards, they are treated as the primary causal explanation for the adverse event. When violations are identified, whether related to the Coal Mine Regulations, the Central Electricity Authority Regulations, or other prescribed safety standards, they are treated in the broader safety-science literature as symptomatic of a compliance-driven investigative logic that risks reducing complex, multi-dimensional accidents to instances of non-compliance at the level of on-site conduct, rather than examining the broader organisational and regulatory system within which that conduct occurs[11,14]. This compliance-driven approach mirrors a broader global issue acknowledged

by occupational safety regulators themselves. According to the U.S. Occupational Safety and Health Administration (OSHA), "It is far too easy, and often misleading, to conclude that carelessness or failure to follow a procedure alone was the cause of an incident. To do so fails to discover the underlying or root causes of the incident, and therefore fails to identify the systemic changes and measures needed to prevent future incidents[15]."

The limitations of compliance-centred accident investigation are well documented in the broader safety science literature, which has long argued for a paradigmatic shift from linear, event-based causal reasoning to a systemic understanding of how accidents emerge from the interaction of factors distributed across multiple levels of a sociotechnical system. The theoretical foundation for this shift was laid by Jens Rasmussen (1997), who proposed a risk management framework in which higher-level decisions, made by legislators, regulators, and senior organisational actors, cascade downward through the system hierarchy to shape the conditions experienced by front-line workers[16]. In this view, an accident is not simply the result of an individual's failure to follow a rule; it is the emergent outcome of a system in which the rules themselves may be poorly designed, inadequately enforced, insufficiently communicated, or structurally incompatible with the realities of operational practice. This perspective has been reinforced by empirical studies conducted across a range of industries[17].

While the systemic accident causation models are well established in the safety science literature, empirical evidence examining whether regulatory investigation practice actually reflects this systemic understanding remains scarce, particularly in the context of developing economies such as India. Researchers have repeatedly noted a persistent gap between the advancement of systems-oriented frameworks in academic discourse and their adoption within formal regulatory inquiry processes[12,18]. For the Indian mining context specifically, no study has systematically examined whether DGMS investigation reports and their associated recommendations reflect the multi-level, systemic framing advocated by contemporary safety science, or whether they default to the compliance-centred explanatory logic that safety researchers have long argued is insufficient[19]. The framing of causality in official investigation outputs has downstream implications for the nature and scope of corrective actions adopted at both the organisational and regulatory levels. Where accident explanations remain fixed to proximal, compliance-related factors, the resulting recommendations may not fully capture the range of systemic factors that contributed to the accident.

Accident investigation has long been recognised as a primary mechanism for organisational learning in safety-critical industries. Traditional investigation approaches rely on expert judgment, interviews, and document analysis to reconstruct sequences of events and identify contributing factors. While these methods can provide rich qualitative insights, they are often limited by cognitive biases, institutional norms, and the selective framing of evidence. In large regulatory systems, accident reports accumulate over decades, forming extensive textual corpora that remain underutilised beyond case-by-case analysis. Recent advances in Natural Language Processing (NLP) and machine learning offer new opportunities to extract latent patterns from large collections of accident narratives[20]. Techniques such as topic modelling, word embeddings, and transformer-based language models enable systematic analysis of unstructured text at a scale that was previously infeasible. In safety research, NLP has been increasingly applied to analyse accident reports, near-miss databases, and safety incident narratives[21,22].

By analysing accident narratives collectively rather than individually, NLP-based approaches can reveal dominant themes and systemic issues that may not be apparent to investigators focused on single events. More importantly, such methods can be used to compare different components of accident documentation, such as descriptions of events versus official conclusions or recommendations. This comparative perspective is particularly valuable for examining how causality is framed and whether explanatory emphasis aligns with the actual complexity of accidents. Despite the availability of rich accident data, regulatory and organisational approaches to accident investigation in

India remain largely rooted in compliance-oriented thinking. The present study employed a BERT topic modelling approach to identify and compare the themes in the textual description of accidents and counterfactuals/recommendations provided by DGMS.

Against this backdrop, the present study addresses a significant gap in the understanding of how accident investigations are conducted and communicated within the Indian mining regulatory context. While the theoretical case for systemic accident analysis is well established in the international safety science literature, there is a conspicuous absence of empirical evidence examining whether the actual outputs of regulatory investigations in Indian mines reflect this systemic understanding or whether they default to narrow, compliance-based explanations. This study addresses that gap by applying a BERT-based topic modelling technique to a corpus of official fatal accident narratives published by DGMS, as well as to the corresponding regulatory recommendations or counterfactual statements issued alongside those narratives. This computational analysis allows for an objective assessment of the thematic scope and distribution of topics in both the accident descriptions and the regulatory responses, revealing whether the counterfactuals track the diversity of hazards and contexts documented in the narratives or whether they converge on a narrow set of compliance-related prescriptions. The findings are intended to contribute not only to the literature on Indian mining safety but also to the broader methodological conversation about how regulatory bodies can integrate systems thinking into their investigative and learning processes, thereby moving beyond the illusion of causality created by oversimplified, compliance-centred explanations.

2. Methodology

The present study employed BERT-based topic modelling to identify latent themes in fatal accident descriptions and the corresponding counterfactuals/recommended actions documented in accident investigation reports published by the DGMS between 2016 and 2022. The initial dataset comprised 338 DGMS-reported fatal accidents. After excluding records with missing narrative text and incomplete counterfactual/recommendation fields, the final dataset consisted of 336 accident reports corresponding to 400 individual fatalities. Each fatality was treated as a separate data point. Figure 1 shows the methodology flow chart.

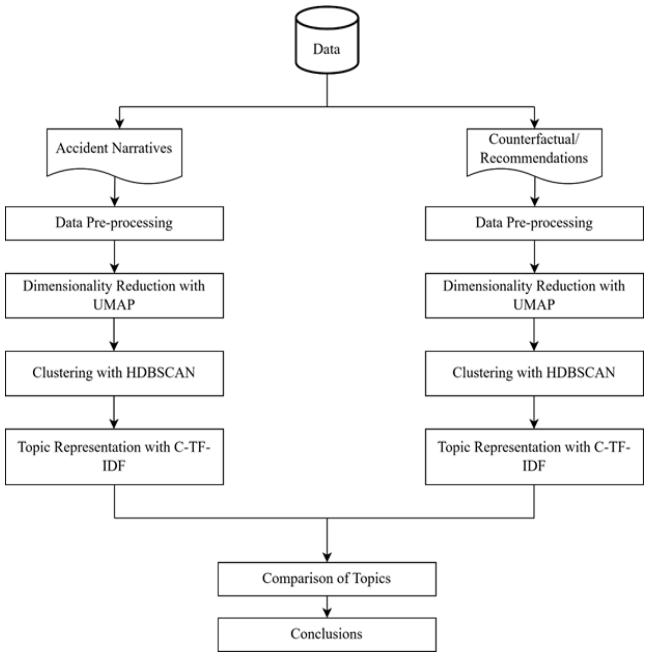


Figure 1. Methodology flow chart.

As shown in Figure 1, the analytical workflow proceeded in four stages. First, DGMS fatal accident narratives and their corresponding counterfactual/recommendation statements were extracted and compiled into a structured dataset. Second, both text corpora were cleaned and normalised. Third, BERT-based topic modelling (BERTopic) was applied independently to the pre-processed narrative corpus and the pre-processed counterfactual corpus, producing two separate sets of topics. Fourth, the resulting topics were validated and interpreted using coherence, Normalised Pointwise Mutual Information (NPMI), inter-topic similarity, and topic diversity metrics, combined with qualitative review of representative excerpts.

2.1 Data Pre-processing

Before topic modelling, both the accident-narrative corpus and the counterfactual/recommendation

corpus were pre-processed to reduce lexical noise while preserving domain-relevant terminology, consistent with BERTopic's recommendation of minimal, targeted pre-processing given its use of contextual embeddings[23]. Pre-processing involved: (i) conversion of text to lowercase; (ii) removal of punctuation and numerals; (iii) removal of standard English stop words; and (iv) lemmatisation of remaining tokens to their base form.

Example of an accident narrative before and after data pre-processing:

Before: "While an empty load haul dump machine, taking the right turn at a junction on the way to lift coal from faces in a development district of an underground coal mine, a dresser standing at the turning beside the coal pillar, was hit by the side of the bucket of the machine, inflicting serious bodily injuries which proved fatal after three hours."

After: "empty load haul dump machine right turn junction lift coal face development district underground coal mine dresser turning coal pillar hit side bucket machine inflict serious bodily injury fatal hour."

2.2 Topic Modelling

Topic modelling is a broad class of unsupervised machine learning techniques designed to discover latent semantic structures, referred to as topics, within large collections of unstructured textual data[24]. Topic modelling has become a foundational tool in NLP, information retrieval, and text mining, enabling researchers and practitioners to automatically extract meaningful thematic patterns from corpora without prior labelling[25]. The evolution of topic modelling methods has progressed through three distinct paradigms: Probabilistic generative models, matrix factorisation approaches, and, most recently, neural embedding-based methods.

Under the probabilistic paradigm of topic modelling, Latent Dirichlet Allocation (LDA) was introduced[24]. LDA operates under the generative assumption that each document in a corpus is a probabilistic mixture of a fixed number of latent topics, and each topic, in turn, is modelled as a probability distribution over the entire vocabulary[24]. Additionally, Non-negative Matrix Factorisation (NMF) represented an alternative approach that decomposes a term-document matrix into two lower-rank non-negative matrices, effectively separating the corpus into latent factors interpretable as topics[26].

The introduction of transformer-based pre-trained language models, most notably BERT, marked a paradigm shift in how text is represented computationally. Unlike Bag-of-Words (BoW) or term Frequency-Inverse Document Frequency (TF-IDF) vectors, contextual embeddings produced by models such as BERT encode rich semantic and syntactic information, capturing word meaning as it varies depending on surrounding context. This representational power has been used for topic modelling through two primary architectures: Top2Vec[27] and BERTopic[23]. Top2Vec trains document and word embeddings jointly using a neural network and identifies topics as dense clusters in the joint embedding space. BERTopic, by contrast, employs a modular pipeline that separates the embedding, dimensionality reduction, clustering, and topic representation stages, offering significantly greater flexibility and reproducibility.

The superiority of BERT-based topic modelling over traditional approaches rests on several advantages. First, contextual embeddings capture polysemy and semantic nuance that BoW-based models entirely miss. Second, BERTopic produces significantly more stable topic assignments across runs compared to LDA, as the clustering is performed on fixed pre-trained embeddings rather than on stochastic sampling procedures[23]. Third, BERT-based models handle short text far more effectively, as the pre-trained embeddings supply contextual information even when individual documents contain very few tokens, a critical advantage for analysing social media, reviews, and survey responses[28].

Furthermore, a comparative analysis conducted in the context of healthcare application reviews demonstrated that BERTopic generated more cohesive and focused topics that were easier to interpret compared to LDA, with the labelling process being substantially faster due to fewer but more discriminative

Table 1. Comparison of topic modelling methods

Feature	LDA	NMF	BERTopic
Text Representation [23-26]	Bag-of-Words	TF-IDF Matrix	Contextual Embeddings (BERT)
Semantic Capture [23, 28]	Word co-occurrence only	Word frequency patterns	Deep contextual semantics
Short Text Performance [28,31]	Poor	Moderate	Excellent
Number of Topics [23,24,28]	Must be pre-specified	Must be pre-specified	Determined automatically
Run-to-Run Stability [28]	Low	Moderate	High
Interpretability [28,31]	Moderate	Moderate	High (focused keywords)
Preprocessing Required [23,28]	Extensive	Moderate	Minimal

words per topic[29]. A biomedical study integrating AI modules further confirmed that clusters produced by BERTopic were more compact and well-separated on t-SNE visualisations, indicating improved coherence and distinctiveness between topics compared to traditional LDA[30]. Collectively, these findings establish BERT-based approaches as the methodological choice for contemporary topic modelling. Table 1 shows the comparative analysis of LDA, NMF, and BERTopic.

2.2.1 BERTopic

BERTopic, as formalised by Grootendorst, constitutes a four-stage neural topic modelling pipeline that transforms raw text into coherent, interpretable topic representations[23]. The architecture is designed to be modular, meaning that each stage can be independently configured or replaced with alternative components without compromising the integrity of the overall pipeline. The BERTopic pipeline was implemented in Python using the open-source 'bertopic' package[23], with document embeddings generated using the 'sentence-transformers' library and the pre-trained 'all-MiniLM-L6-v2' model, dimensionality reduction performed using Uniform Manifold Approximation and Projection (UMAP) ('umap-learn'), and clustering performed using Hierarchical Density Based Spatial Clustering of Applications with Noise (HDBSCAN) ('hdbscan'). The description of the BERTopic pipeline is given below:

Step 1: Document Embedding: Raw text is converted into dense semantic vectors using a pre-trained Sentence-Transformer model (e.g., all-MiniLM-L6-v2), producing an n × d embedding matrix.

Step 2: Dimensionality Reduction: UMAP reduces the high-dimensional embeddings to a lower-dimensional space while preserving local and global semantic structure.

Step 3: Clustering: HDBSCAN identifies dense clusters in the reduced space, automatically determining the number of topics and assigning noise documents to topic -1.

Step 4: Topic Representation: C-TF-IDF extracts the most discriminative keywords for each cluster, generating concise and interpretable topic labels.

In the present study, topic quality was assessed using topic coherence, NPMI, inter-topic similarity, and topic diversity. These metrics are widely used in the topic modelling literature to evaluate both within-topic consistency and between-topic distinctiveness[32].

Topic coherence (C_v) measures the degree to which the most important words of a topic are semantically

related. It has been shown to correlate most strongly with human judgments of topic interpretability and is therefore widely adopted[32]. NPMI measures the extent to which pairs of topic words co-occur more frequently than expected by chance. Unlike C_v, NPMI relies strictly on lexical co-occurrence and is therefore a more conservative metric. NPMI values range from -1 to +1. Positive values indicate stronger-than-chance co-occurrence, zero indicates independence, and negative values indicate less-than-chance co-occurrence[33,34]. Negative NPMI values are common in embedding-based topic models and do not necessarily indicate poor topic quality, particularly when accompanied by high topic coherence scores.

Inter-topic similarity measures the semantic proximity between topics and is used to assess redundancy or hierarchical relationships[35]. In embedding-based models, this is typically computed using cosine similarity. Topic diversity evaluates the extent to which topics share keywords. High topic diversity indicates that topics are lexically distinct and that the model avoids redundancy[36]. This metric is global rather than topic-wise and is designed to capture distinctiveness across the entire set of topics by quantifying the extent to which topics share common keywords. Lower values indicate substantial overlap between topics, whereas higher values reflect clearer separation.

3. Results and Discussion

The results section is divided into two parts. The first part presents the topic modelling results derived from accident narrative texts, while the second part presents the topic modelling results from counterfactual/recommendation texts.

3.1 Accident Narratives

Table 2 presents the results of BERT-based topic modelling applied to 400 coal mine fatality accident narratives.

The topic modelling analysis identified thirteen distinct themes, each characterised by associated keywords and document distribution. Together, these topics capture a broad spectrum of operational,

Table 2. Topic distribution and characteristics of coal mine fatality accident narratives

Topic ID	Topic Label	Keywords	Count (%)
Topic-1	Unclassified/General Mining Injuries	'persons', 'person', 'injuries', 'coal', 'machine', 'succumbed', 'bodily', 'road', 'drill', 'spot'	75 (18.75%)
Topic 0	Surface Mining Vehicle Operations Related Accidents	'truck', 'opencast', 'opencast coal', 'grader', 'injuries', 'road', 'coal', 'inflicting', 'motor', 'bodily'	43 (10.75%)
Topic 1	Tipper/Haul Truck Accidents	'tipper', 'right', 'succumbed', 'injuries succumbed', 'injuries', 'bodily injuries', 'road', 'left', 'loaded', 'bodily'	42 (10.5%)
Topic 2	Roof Fall and Support Failures	'roof', 'persons', 'crew', 'measuring', 'height', 'escaped', 'drilling', 'fell height', 'support', 'width'	30 (7.5%)
Topic 3	Coal Rib/Sidewall Collapse	'height', 'rib', 'measuring', 'coal', 'men', 'injuries', 'suddenly', 'inflicting', 'fell', 'person'	29 (7.25%)
Topic 4	Heavy Mobile Equipment Collisions	'dumper', 'operator', 'tanker', 'road', 'haul road', 'haul', 'water tanker', 'water', 'dump', 'vehicle'	29 (7.25%)
Topic 5	Blasting and Explosive Operations Related Accident	'blasting', 'explosive', 'roof', 'measuring', 'holes', 'dresser', 'injuries', 'parted', 'depillaring', 'panel'	25 (6.25%)

Topic 6	Conveyor Belt Systems Related Accident	'belt', 'conveyor', 'conveyor belt', 'drum', 'bunker', 'fell', 'head', 'coal', 'plant', 'succumbed'	25 (6.25%)
Topic 7	Overburden and Bench Instability-Related Accident	'overburden', 'benches', 'dumped', 'workers', 'cracks', 'hole blasting', 'deep hole', 'machinery', 'failure', 'form'	23 (5.75%)
Topic 8	Mobile Equipment Machinery Accidents	'operator', 'engine', 'bodily', 'chain', 'ramp', 'started', 'injuries', 'track chain', 'track', 'yard'	19 (4.75%)
Topic 9	Electrical Contact/ Electrocution	'electric', 'pole', 'line', 'overhead', 'received', 'overhead line', 'contact', 'came contact', 'came', 'got'	18 (4.5%)
Topic 10	Underground Haulage Systems Related Accident	'tubs', 'haulage', 'set', 'loaded', 'roadway', 'coal tubs', 'rope', 'loaded coal', 'endless', 'track'	17 (4.25%)
Topic 11	Water Accumulation and Drowning	'water', 'sump', 'pump', 'dip', 'main', 'drowned', 'shelter', 'operators', 'got drowned', 'away'	13 (3.25%)
Topic 12	Fire and Thermal Burns	'hot', 'burn', 'bench', 'ash', 'burn injuries', 'hospital', 'succumbed hospital', 'causing', 'causing burn', 'days'	12 (3.0%)

geotechnical, mechanical, and environmental hazards across surface and underground mining. The distribution of documents across topics is uneven, reflecting both the heterogeneity of mining operations and the concentration of accident occurrences around specific high-risk activities. The largest category, labelled Unclassified/General Mining Injuries (18.75%), comprises incident descriptions that lack sufficient detail to be clearly associated with a specific topic/hazard. Among the classified topics, Surface Mining Vehicle Operations-related accidents (10.75%) and Tipper/Haul Truck Accidents (10.5%) emerge as the most prevalent operational themes. These topics predominantly involve heavy vehicle movement on haul roads, interactions between mobile equipment, and loss-of-control scenarios. The close prevalence of these two topics suggests that vehicle-related hazards remain one of the dominant sources of injury in surface mines.

Ground control-related accidents collectively represent a substantial proportion of the dataset. Roof Fall and Support Failures-related accidents (7.5%) and Coal Rib/Sidewall Collapse (7.25%) capture incidents associated with the instability of underground excavations and coal faces. The keywords within these topics emphasise measurements, heights, and sudden collapses, suggesting that failures often occur during drilling or production activities. The comparable prevalence of these two topics indicates that both roof and rib instabilities remain persistent and distinct safety challenges, despite long-standing regulatory controls and support design guidelines. Similarly, Overburden and Bench Instability (5.75%) reflects geotechnical failures during surface operations, including bench cracking and overburden collapse during or following blasting. Although less frequent than vehicle-related incidents, such events typically carry a high severity potential due to the scale of material movement involved[1].

Several topics highlight hazards associated with fixed and mobile equipment systems. Heavy Mobile Equipment Collisions (7.25%) and Mobile Equipment Machinery Incidents (4.75%) encompass accidents involving dumpers, tankers, engines, and track systems. While both relate to machinery, the separation of these topics suggests differing risk profiles, one dominated

Table 3. Representative accident narratives for each topic category

Topic ID	Representative Document Excerpt
Topic 0	While the truck driver was directing the movement of the truck from behind near the coal stockpile, the truck suddenly moved backwards after being operated by the *khalasi* (helper), causing severe bodily injuries to the driver. The driver later succumbed to the injuries while being transported to the hospital.
Topic 1	While a tipper was negotiating a left-hand curve in a railway siding of mines, a tipper helper passing behind a coal heap was hit by the front left wheel of the tipper, thus sustaining serious lacerated injuries on his stomach, which proved fatal on the way to hospital within half an hour.
Topic 2	While a crew of seven support personnel were erecting steel chock at goaf edge of a split gallery in depillaring panel in a belowground coal mine, a piece of sand stone measuring 1.4m (length) x 1.0m (width) x 0.20m-0.40m (thick) parted off from a height of 4.1m and fell on one of them, inflicting serious bodily injuries, which proved fatal instantaneously and others escaped unhurt.
Topic 3	While a Sirdar and Shot-firer were inspecting the face after blasting in a depillaring district of a belowground coal mine, the roof measuring about 12m in length x 1-4.0m in width x 0-3.0m in thickness fell from a height of 3.7m, due to extension of fall from the goaf into the gallery, resulting in fatal injuries to Sirdar, whereas the Shot-firer escaped unhurt.
Topic 4	While a 120 Tonne carrying capacity dumper loaded with Overburden (OB) was on the way to the OB dump yard on a near-level and 19.80m wide haul road, it collided nose to tail with a stationary 120 Tonne dumper parked on the left side of the haul road, seriously. injuring the operator and entrapping him in the damaged cabin, to which he succumbed almost instantaneously
Topic 5	While contractual workers and machinery were deployed to form benches in dumped overburden, in-situ overburden and coal in an opencast mine containing known geological disturbances; cracks developed due to the dead weight of dumped overburden, deep hole blasting and repeated plying of Heavy Earth Moving Machinery (HEMM), resulting in violent failure of in-situ overburden and coal benches, followed by instant sliding of dump in an area of 720m x 366m burying 23 workers along with 18 HEMM.
Topic 6	While two workers were approaching the drive head of the apron feeder through the dribble conveyor for examining/cleaning clogged material after stoppage of apron feeder/ dribble (spillage) conveyor in a coal crushing plant of a Mine, suddenly the dribble conveyor came into operation, thereby trapping both the workers; however, one of the worker got entangled in between the channel of the dribble conveyor and the base, and he received serious bodily injuries which proved fatal soon after.
Topic 7	While a person was dressing the roof of a junction in a development district of an underground coal mine, a mass of shaley sandstone roof, measuring about 3.6m x 2.3m x 0.3m to 1.0m(thick), fell from a height of about 2.2m, inflicting fatal injuries to him and serious bodily injuries to another person standing nearby.
Topic 8	While a person inadvertently came in contact with the pole of a faulty electric line which received power from switchgear with defunct protective systems, received severe electric shock and died.
Topic 9	While a dozer was passing through a blasted area over an old developed waterlogged fiery gallery in an opencast coal mine, due to sudden collapse of parting, the dozer submerged into the waterlogged and fiery gallery, resulting in the instantaneous death of the dozer operator because of blunt and burn injuries.

Topic 10	While a pump operator was ascending alone in a cage, jumped to the shaft bottom. from the cage by lifting the fencing oneself, from a height of 70m, inflicting injuries on his head, to which he succumbed almost instantly.
Topic 11	While a person was connecting a cable to the SCEH-6 kiosk unit, which was installed near the edge of the M-28 sump, fell into the sump from a height of 7m and drowned.
Topic 12	While the previously dumped fly ash over an overburden bench side was being side cast by deployment of an excavator at the bench side in an opencast coal mine; portion of the bench side consisting of fly ash, measuring about 35m (L) x 11.5m (H) x 5m (D), failed, resulting burying of the person engaged there to supervise the operations/work place, who was declared dead at the hospital after one hour and fifteen minutes after the accident.

by vehicle-to-vehicle or vehicle-to-infrastructure interactions, and the other by mechanical failures or operator-machine interface issues. Process-specific hazards are captured in topics such as Blasting and Explosive Operations (6.25%) and Conveyor Belt Systems (6.25%). Blasting-related incidents reflect the inherently high-risk nature of explosive use. Conveyor belt incidents, on the other hand, frequently involve entanglement, falling materials, or contact with moving components, indicating ongoing challenges in guarding, maintenance, and safe access around continuous material-handling systems.

Electrical hazards, represented by Electrical Contact/Electrocution (4.5%), remain a significant though less frequent source of accidents. The keywords suggest contact with overhead lines and poles, pointing to risks during equipment movement or maintenance activities. Despite lower prevalence, such incidents are typically severe or fatal, emphasising the importance of visibility, clearance management, and isolation procedures. Underground transport risks are reflected in Underground Haulage Systems (4.25%), which involve tubs, ropes, tracks, and loaded haulage systems. These incidents highlight the continued relevance of traditional haulage hazards in underground mines, particularly in constrained roadways where human-machine interactions are unavoidable. Environmental and secondary hazards are captured by Water Accumulation and Drowning (3.25%) and Fire and Thermal Burns (3.0%). Although these topics account for smaller proportions of fatal accidents, they often involve delayed detection, inadequate monitoring, or challenges with emergency response. Their presence as distinct topics indicates that such hazards, while less frequent, remain systemic issues rather than isolated anomalies.

Table 3 presents sample accident narratives for each identified topic, illustrating the characteristic patterns and circumstances associated with different categories of coal mine fatalities. Each topic is represented by one exemplar narrative extracted from the dataset with stop words removed, demonstrating the typical scenarios within that category.

Table 4 shows pairwise semantic similarity between topics identified, computed as cosine similarity between topic embeddings. The inter-topic similarity matrix provides insight into the semantic relationships among the identified accident themes. Rather than viewing topics as isolated categories, the similarity structure highlights clusters of related activities and shared underlying conditions that shape accident occurrence.

Several topic pairs exhibit moderate to high semantic similarity, indicating closely related operational contexts. For example, Topics 0 and 1 show a relatively high similarity (0.59), reflecting the strong conceptual linkage between general surface vehicle operations and tipper or haul-truck-specific incidents. This relationship suggests that these topics are not independent risk domains, but rather represent different manifestations of a shared traffic and mobile equipment system, influenced by common factors such as haul-road design, visibility, operator behaviour, traffic control, and management practices. Similarly, the moderate similarity between Topics 0, 4, and 6 indicates interaction between vehicle operations, heavy

mobile equipment, and associated machinery systems, reinforcing the idea that mobile equipment hazards arise from interacting system elements rather than single-point failures. A second notable cluster emerges among Topics 3, 4, and 6, where similarity values exceed 0.49 in several pairwise comparisons. These topics relate to ground control failures and equipment interactions, suggesting that geotechnical instability and machinery operations are often coupled in practice. Such coupling is characteristic of tightly coordinated mining activities, where changes in ground conditions, equipment positioning, and work sequencing can collectively influence accident risk.

In contrast, many topic pairs display low similarity values (below 0.30), indicating that the topic model has successfully distinguished between semantically different hazards. This separation supports the validity of the extracted topics and aligns with the high topic diversity observed in the model. Importantly, Topic 8 shows consistently low similarity with all other topics, with values largely below 0.10. This pattern suggests that Topic 8 represents a highly distinct accident mechanism, likely associated with specialised operational conditions or less frequent but clearly defined events. Such topics, while contributing a smaller proportion of incidents, may correspond to

Table 4. Inter-topic similarity matrix of accident narrative topics

	Topic 0	Topic 1	Topic 2	Topic 3	Topic 4	Topic 5	Topic 6	Topic 7	Topic 8	Topic 9	Topic 10	Topic 11	Topic 12
Topic 0	1.00	0.59	0.44	0.45	0.52	0.39	0.44	0.42	0.08	0.42	0.27	0.30	0.24
Topic 1	0.59	1.00	0.53	0.25	0.34	0.48	0.31	0.29	0.08	0.45	0.21	0.25	0.15
Topic 2	0.44	0.53	1.00	0.15	0.25	0.43	0.19	0.25	0.07	0.41	0.16	0.23	0.11
Topic 3	0.45	0.25	0.15	1.00	0.57	0.15	0.56	0.23	0.04	0.20	0.13	0.16	0.19
Topic 4	0.52	0.34	0.25	0.57	1.00	0.19	0.49	0.33	0.08	0.31	0.19	0.25	0.20
Topic 5	0.39	0.48	0.43	0.15	0.19	1.00	0.17	0.19	0.08	0.37	0.15	0.17	0.21
Topic 6	0.44	0.31	0.19	0.56	0.49	0.17	1.00	0.26	0.11	0.26	0.17	0.18	0.14
Topic 7	0.42	0.29	0.25	0.23	0.33	0.19	0.26	1.00	0.06	0.30	0.23	0.23	0.15
Topic 8	0.08	0.08	0.07	0.04	0.08	0.08	0.11	0.06	1.00	0.10	0.04	0.02	0.06
Topic 9	0.42	0.45	0.41	0.20	0.31	0.37	0.26	0.30	0.10	1.00	0.20	0.22	0.16
Topic 10	0.27	0.21	0.16	0.13	0.19	0.15	0.17	0.23	0.04	0.20	1.00	0.14	0.11
Topic 11	0.30	0.25	0.23	0.16	0.25	0.17	0.18	0.23	0.02	0.22	0.14	1.00	0.09
Topic 12	0.24	0.15	0.11	0.19	0.20	0.21	0.14	0.15	0.06	0.16	0.11	0.09	1.00

latent systemic vulnerabilities that are not captured by more dominant hazard categories.

The overall structure of the similarity matrix reflects a system characterised by both clustering and differentiation. Clustering indicates areas where multiple accident types share common operational and organisational influences, while differentiation highlights domains where hazards are context-specific and less interconnected. This dual structure is typical of complex socio-technical systems, where local adaptations and variations in work practice coexist with standardised processes and shared infrastructure. From an organisational learning perspective, the inter-topic similarity analysis provides evidence that, for accident categories exhibiting moderate-to-high inter-topic similarity, safety interventions should not be designed solely around individual accident categories. At the same time, the presence of highly distinct topics, such as Topic 8, underscores the need for targeted, category-specific controls and monitoring strategies for specialised risks that may otherwise be overlooked in aggregated safety analyses.

Table 5 reports the topic-wise coherence score and NPMI for each topic identified using BERTopic. Topic diversity is reported at the model level to indicate overall distinctiveness across topics. The topic-wise evaluation using coherence (C_v), NPMI, and topic diversity provides a comprehensive assessment of the quality, interpretability, and robustness of the extracted topics. Together, these metrics allow validation of the topic model not only in terms of internal semantic consistency, but also in relation to lexical behaviour and inter-topic distinctiveness.

Overall, the majority of topics exhibit moderate to very strong semantic coherence, with C_v values exceeding 0.60 for most topics. Several topics demonstrate particularly high coherence, including Topics 7, 9, and 10, with C_v values above 0.89. Such high coherence indicates that, within each of these topics, the top keywords are strongly semantically related to one another, and that the corresponding accident narratives describe well-defined and internally consistent hazards. A smaller number of topics, such as topics 0 and 8, exhibit comparatively lower coherence values (below 0.50). Rather than indicating poor model performance, such variability may reflect the inherent complexity and context-dependence of certain operational domains.

The NPMI values display greater variation across topics, ranging from moderately positive to negative. Topics with higher NPMI values, particularly Topics 9, 10, and 7, indicate stronger lexical co-occurrence among topic keywords, suggesting that the associated terms frequently appear together within individual accident reports. These topics likely correspond to operational scenarios where specific technical elements or processes are repeatedly described using consistent terminology. In contrast, several topics exhibit low or negative NPMI values, including Topics 3, 5, and 8. Negative NPMI indicates that while the keywords are semantically related, they do not frequently co-occur within the same documents. This pattern is characteristic of embedding-based topic models applied to short and heterogeneous texts, such as accident narratives. Importantly, the coexistence of high C_v values with low or negative NPMI in several topics suggests that semantic coherence is not driven solely by lexical repetition. Instead, it reflects deeper contextual similarity captured by the embedding-based representation: BERT-derived embeddings encode semantic relationships between words and phrases that are used in similar situational contexts across narratives (for example, synonymous or paraphrased descriptions of the same underlying hazard mechanism, such as 'roof fell' and 'strata collapsed'), even when such terms do not co-occur within the same individual document. NPMI, by contrast, is a purely lexical co-occurrence measure and cannot capture this form of semantic equivalence.

Topic diversity, reported at the model level, is high (0.84), indicating strong overall distinctiveness among topics. This high diversity confirms that the model avoids excessive overlap in keyword usage and that each topic captures a relatively unique aspect of the accidents. In combination with the inter-topic similarity analysis, this supports the conclusion that the topic set is both well separated and meaningfully structured, rather than fragmented or redundant.

Table 5. Topic-wise coherence and lexical association metrics for the accident narrative topics

Topic Id	Coherence Value	NPMI	Topic Diversity
Topic 0	0.496	-0.0251	0.84
Topic 1	0.657	0.0558	
Topic 2	0.878	0.0271	
Topic 3	0.604	-0.0611	
Topic 4	0.669	0.0113	
Topic 5	0.640	-0.0711	
Topic 6	0.614	0.0215	
Topic 7	0.994	0.1770	
Topic 8	0.478	-0.1924	
Topic 9	0.894	0.3236	
Topic 10	0.943	0.3521	
Topic 11	0.721	0.1092	
Topic 12	0.807	-0.0351	

Taken together, these metrics provide an indication that the topics are interpretable, distinct, and useful representations of accident themes.

3.2 Counterfactual/ Recommendation Narratives

The analysis of counterfactual statements issued by the regulatory authority reveals a strong and consistent emphasis on regulatory compliance and prescribed preventive measures. The topic modelling results show that an overwhelming majority of counterfactuals (86.75%) fall under Regulatory Compliance and Preventive Measures, with dominant references to coal mine regulations, statutory requirements, and assertions that accidents could have been averted through adherence to existing rules. This pattern indicates that regulatory accident analysis is primarily framed around deviations from SOPs and formal regulatory provisions, rather than a broader examination of systemic conditions. Such an emphasis is consistent with what has been described in the safety literature as a normative or rule-based model of accident causation, where accidents are retrospectively explained by identifying violations of established procedures or human errors[14, 16, 37].

The smaller proportion of counterfactuals related to Ground Stability and Dump Management (6.25%) and Electrical Safety Standards (4.5%) further reinforces this observation. Although these topics refer to technically substantive hazard domains, the counterfactual logic remains largely compliance-oriented, focusing on statutory limits, approvals, and adherence to specific regulatory clauses. The analysis thus suggests that even when technical hazards are addressed, they are framed predominantly in terms of whether predefined controls were followed, rather than how those controls functioned or failed to function within real operational contexts[38].

This finding aligns with prior critiques of regulatory accident investigations, which have noted a tendency to focus on rule violations rather than examining how organisational pressures, resource constraints, and adaptive work practices shape unsafe conditions[11,12]. Systems safety research emphasises that safety is an emergent property of interactions among system components, and that compliance with procedures does

not necessarily guarantee safe outcomes, particularly in dynamic and high-risk industries such as mining[39].

Moreover, the dominance of a single compliance-focused topic suggests a relatively narrow construction of counterfactuals, where alternative accident trajectories are framed almost exclusively around regulatory adherence. In contrast, systems-oriented approaches view counterfactual analysis as an opportunity to explore how different configurations of system functions, decision-making processes, and control structures might have altered outcomes[16,38]. From this perspective, regulatory counterfactuals appear to function primarily as post hoc validations of existing rules, rather than as mechanisms for identifying deeper systemic vulnerabilities or opportunities for adaptive improvement. The implications for organisational learning are significant. When accident investigations emphasise SOP deviations as the primary explanatory mechanism, learning tends to be limited to reinforcing existing rules or issuing additional procedures. Such outcomes have been shown to produce diminishing safety returns over time, as they fail to address the underlying conditions that shape everyday work and risk-taking[8]. In the context of mining safety, this may contribute to the sector-wide recurrence of similar accident-type patterns across different mines and years, despite repeated regulatory interventions, rather than a reduction in the diversity of hazard types documented over time.

Overall, the topic modelling results suggest that regulatory accident analysis is strongly oriented toward compliance-based explanations, with limited engagement with the broader socio-technical dynamics of mining systems. While this approach is effective for legal enforcement and standardisation, it constrains deeper system-level learning. Integrating systems safety concepts, such as control structure analysis, functional interactions, and performance variability, into regulatory counterfactual analysis could enhance its explanatory depth and support more resilient safety management practices[16,38,40].

Table 6 presents the topic structure derived from BERTopic modelling of the regulatory counterfactual/recommendation statements issued by DGMS following fatal accident investigations. Three topics were identified: Topic A (Regulatory Compliance and Preventive Measures) overwhelmingly dominates the corpus at 86.75%, characterised by generic references to coal mine regulations and statements that the accident could have been averted through compliance. Topic B (Ground Stability and Dump Management, 6.25%) and Topic C (Electrical Safety Standards under the Central Electricity Authority Regulation, 4.5%) form small, more specialised clusters addressing strata/dump-related and electrical hazard-specific recommendations, respectively.

Table 7 shows representative excerpts that demonstrate how regulatory counterfactuals are predominantly framed in terms of adherence to formal regulations and prescribed operating procedures. Across all topics, accident prevention is articulated

Table 6. Distribution of topics derived from regulatory counterfactual analyses

Topic ID	Topic Label	Keywords	Count (%)
Topic A	Regulatory Compliance and Preventive Measures (Coal Mine Regulation)	'regulation', 'mines', 'coal', 'coal mines', 'regulations', 'required', 'mines regulations', 'regulation coal', 'averted', 'accident averted'	347 (86.75%)
Topic B	Ground Stability and Dump Management	'dump', 'benches', 'statute', 'prevent danger', 'cracks', 'dgms', 'exceeding', 'prior', 'danger', 'given'	25 (6.25%)
Topic C	Electrical Safety Standards (Central Electricity Authority Regulation)	'supply', 'electric', 'measures', 'central', 'authority', 'relating', 'safety', 'regulations', 'required', 'regulation'	18 (4.5%)

Table 7. Representative accident narratives for each topic category of counterfactuals

Topic ID	Representative Document Excerpt
Topic A	The support personnel been not deployed for work under the roof of a split gallery, which was not kept adequately supported by full column grouted roof bolts in rows at a maximum interval of 1.2m between bolts in the same row and 1.2m between rows of bolts and also with cogs at interval of not exceeding 2.4m as the galley was wider than 4.8m, as required by condition No.4(a) and 5(c) respectively, of the Systematic support rules of the depillaring panel modified vide Directorate's letter No. JA/010491/ Permission/ 95/2015/536 dated 11.02.2016, under Regulation 108 of the Coal Mines Regulations, 1957; ii) effective compliance with the provisions of the Systematic support rules of depillaring panel modified vide Directorate's letter no. JA/010491/Permission/95/2015/536 dated 11.02.2016, been ensured and the mine was not worked in contravention thereof as required by Regulation 108(5) of the Coal Mines Regulations, 1957; this accident could have been averted.
Topic B	Geotechnical study been done after occurrence of roof fall in the area and accordingly Strata Control and Monitoring Plan(SCAMP) been framed with due regard to the engineering classification of strata, local geological condition, past experience and accordingly support plan for securing the sides been framed and implemented before allowing any person to the work in the area as required under Reg.123(1) (2) of Coal Mines Regulations, 2017; monitoring of support performance, measurement of strata behavior including movement of the sides and re-setting of supports, and clearing of falls of ground been done properly according to support plan as required under Reg.123(6)of the Coal Mines Regulations, 2017; every ledge and every prominent crack or slip been supported with at least a pair of cogs and suitably lagged, as required under Reg.124(10) of Coal Mines Regulations, 2017; the overhanging sides been dressed down as required under Reg.124(11) of Coal Mines Regulations, 2017; every place wherein fall of roof or sides has occurred, no work of cleaning the dislodged coal or the fall or any part thereof shall be undertaken nor shall any persons be allowed to pass until the newly exposed roof and sides in the vicinity thereof have been examined and made safe if necessary by temporary supports as required under Reg.124(15) of Coal Mines Regulations, 2017; the roof and sides of all working places and travelling roadways including airways and travelling roadways to second outlets been made and kept secured as required under Reg.115(1) of Coal Mines Regulations, 2017, this accident could have been averted.
Topic C	The audio visual prestart warning alarm not been defunct in contravention of Regulation 97(5) of the Coal Mines Regulations, 2017. ii) the cleaning of belt conveyor not been carried out while in motion in contravention with Code of Practice Clause No. A(5) framed in compliance with Regulation 97(6) of the Coal Mines Regulations, 2017. iii) Fail safe system existed, so that belt conveyor could not be started or put in motion while spillage cleaning is in progress as required under Regulation 97(6) of the Coal Mines Regulations, 2017 iv) Sufficient lighting has been provided at the tail end drum of the conveyor belt as required under Code of Practice Clause No. A(9) framed in compliance with Regulation 97(6) read with DG Tech Cir 07/2016 and Regulation 178(b) of the Coal Mines Regulations, 2017. v) The effective means of telecommunication, along with broadcasting facilities been provided at the drive head of the conveyor belt as required under Regulation 97(4) of the Coal Mines Regulations, 2017; this accident could have been averted.

through explicit references to statutory clauses, notifications, and regulatory mandates, with repeated assertions that the accident "could have been averted" through compliance. The excerpts show limited engagement with the broader operational context or systemic interactions preceding the accident, instead emphasising procedural correctness, supervisory compliance, and conformity with regulatory provisions.

Table 8 shows an inter-topic similarity matrix that indicates the regulatory counterfactual themes are largely distinct, with generally low similarity values across topics. The strongest relationship is observed between Topic A (Regulatory Compliance and Preventive Measures) and Topic C (Electrical Safety Standards), with a similarity value of 0.36. This moderate overlap reflects a shared compliance-oriented logic, where accident prevention is articulated through detailed references to statutory provisions and prescribed safety measures. Electrical safety counterfactuals thus appear to be framed as a specialised extension of broader regulatory compliance rather than as an emergent or systemic risk domain. In contrast, Topic B (Ground Stability and Dump Management) shows low similarity with both Topic A (0.21) and Topic C (0.062), indicating that geotechnical hazards are treated as a relatively separate regulatory category.

The topic quality metrics indicate that all three regulatory counterfactual topics exhibit high semantic coherence, with identical C_v values of 0.79. This suggests that the keywords and representative narratives within each topic are internally consistent and that the regulatory authority employs a stable and well-defined vocabulary when framing accident prevention. From an analytical perspective, this high coherence reflects a standardised mode of reasoning, where counterfactual explanations are constructed around clearly articulated regulatory provisions and procedural expectations.

The NPMI values, however, show greater variation across topics. Topic C exhibits the highest NPMI (0.34), indicating strong lexical co-occurrence and a highly standardised regulatory language in the context of electrical safety. Topic A shows moderate positive NPMI (0.24), consistent with frequent co-occurrence of compliance-related terms across counterfactual statements. In contrast, Topic B displays a negative NPMI (-0.22), suggesting that while the topic is semantically coherent, the associated terms do not consistently co-occur within individual documents. This pattern likely reflects the context-dependent and situational nature of ground stability narratives, which are framed using varied vocabulary depending on local conditions. Topic diversity, reported at the model level, is very high (0.90), indicating strong distinctiveness among the three topics and minimal lexical overlap.

Table 8. Inter-topic similarity matrix of counterfactual topics

	Topic A	**Topic B**	**Topic C**
Topic A	1	0.21	0.36
Topic B	0.21	1	0.062
Topic C	0.36	0.062	1

Table 9. Topic-wise coherence and lexical association metrics for the counterfactual topics

Topic Id	**Coherence Value**	**NPMI**	**Topic Diversity**
Topic A	0.79	0.24	0.9
Topic B	0.79	-0.22	
Topic C	0.79	0.34	

4. Policy Implication

The findings of the present study carry important implications for how regulatory bodies can evolve their investigative frameworks to better reflect the systemic complexity of fatal accidents in Indian coal mines. The most direct policy implication concerns the methodological tools used to structure accident investigations. Regulatory bodies currently employ enquiry-based investigation practices oriented around regulatory non-compliance, but a growing body of evidence supports the adoption of systems-thinking methods such as AcciMap, the Systems Theoretic Accident Model and Processes (STAMP), and its companion analytical technique, the Causal Analysis using Systems Theory (CAST). These methods explicitly map contributory factors across multiple levels of the sociotechnical system, from government policy and regulatory bodies down to physical processes and equipment, rather than terminating analysis at the point of a procedural violation. Rasmussen's AcciMap approach provides a particularly suitable template for regulatory investigation in high-risk industries because it recognises that decisions made at higher organisational and governmental levels cascade downward and shape the conditions encountered by frontline workers[16]. Applying AcciMap-based analysis to investigation reports would allow investigators to systematically document contributory factors at each level of the system, producing a richer causal picture that more accurately reflects the distributed origins of fatal accidents[16]. STAMP and CAST offer a complementary approach: rather than tracing a chain of events backwards to a proximate cause, CAST examines how control structures across organisational hierarchies failed to enforce safety constraints, making it well-suited for analysing repeated accident patterns across large corpora of investigation reports[41]. The adoption of systemic methods is unlikely to be effective without accompanying investments in investigator training and competency development. Regulatory bodies could consider developing structured training programs for their officers on the application of at least one recognised systemic investigation method, supported by standardised templates to guide the documentation of contributory factors across system levels.

5. Conclusion

This study examined whether regulatory accident investigations in Indian coal mines reflect the multi-level complexity of accident causation or whether they tend toward a narrow, compliance-driven framing that may not fully capture the systemic conditions underlying adverse events. By applying BERT-based topic modelling to both fatal accident narratives and the corresponding regulatory counterfactuals issued by the regulatory body, the study provides empirical evidence of a clear mismatch between the complexity of accidents as described in narratives and the simplified causal logic reflected in regulatory recommendations. Several limitations of the study should be acknowledged. Firstly, the analysis relies on officially published narratives and counterfactuals, which may themselves be shaped by institutional norms and reporting conventions. Second, topic modelling captures dominant semantic patterns but does not establish causal relationships. Third, the study focuses on coal mine fatalities in India, which may limit generalisability to other sectors or regulatory regimes. Fourth, the analysis used a combined dataset because mine-type-specific subsets were insufficient for stable BERTopic modelling. Future research should examine opencast and underground mines separately using larger datasets.

Notwithstanding these limitations, the study makes a meaningful contribution by providing the first empirical assessment of thematic alignment between accident descriptions and regulatory responses in the Indian mining context. Longitudinal comparative studies across regulatory jurisdictions could further illuminate how investigative culture, institutional mandate, and resource availability influence the depth and scope of accident learning. Ultimately, embedding systems safety principles more explicitly into the investigative and reporting practices of regulatory bodies remains a critical step toward moving beyond proximal explanations and building the kind of organisational learning that can

meaningfully reduce fatal accident rates in Indian coal mining.

6. References

1. Kumar A, Upadhyay R, Samanta B, Bhattacherjee A, editors. A novel approach to mining risk assessment: Fuzzy opinion aggregation and fuzzy inference system. International Conference on Sustainable and Innovative Mining Practices; 2023: Springer. https://doi.org/10.1007/978-3-031-76614-5_3
2. Cruz-Ausejo L, Cama-Ttito NA, Solano PF, Copez-Lonzoy A, Vera-Ponce VJ. Occupational accidents in mining workers: Scoping review of studies published in the last 13 years. BMJ Open. 2024; 14(10):e080572. https://doi.org/10.1136/bmjopen-2023-080572 PMid:39395823 PMCid:PMC11474850.
3. Sharma A, Kumar A, Saini S. Associations of occupational and demographic factors with cause of fatalities in the Indian coal mining sector. J Mines Met Fuels. 2025; 73(2).
4. Gomero Cuadra R, Vargas Zegarra R, Roldán Silva A, Mejía CR. Asociación entre parámetros bioquímicos y los accidentes de trabajo de mineros peruanos entre los años 2009 y 2014. Rev Asoc Esp Espec Med Trab. 2018; 27(3):150-156.
5. ILO. Mining: A hazardous work: International Labour Organisation; 2015. Available from: https://www.ilo.org/resource/mining-hazardous-work.
6. Tiwari V, Garg A, Kapshe M, Deshpande A, Vishwanathan S. Assessing possibilities for coal continuance in India under climate constraints. Int J Greenh Gas Control. 2023; 122:103811. https://doi.org/10.1016/j.ijggc.2022.103811
7. DGMS. Annual Report 2024; 2025.
8. Dekker S. Drift into failure: From hunting broken components to understanding complex systems. CRC Press; 2016. https://doi.org/10.1201/9781315257396
9. Zhang Y, Dong C, Guo W, Dai J, Zhao Z. Systems theoretic accident model and process (STAMP): A literature review. Safety Sci. 2022; 152:105596. https://doi.org/10.1016/j.ssci.2021.105596 PMCid:PMC10718422.
10. Sedlar N, Irwin A, Martin D, Roberts R. A qualitative systematic review on the application of the normalization of deviance phenomenon within high-risk industries. J Saf Res. 2023; 84:290-305. https://doi.org/10.1016/j.jsr.2022.11.005 PMid:36868658.
11. Dekker S. The field guide to understanding 'human error'. CRC Press; 2017. https://doi.org/10.1201/9781315239675
12. Leveson NG. Engineering a safer world: Systems thinking applied to safety. The MIT Press; 2016.
13. Leveson NG. An introduction to system safety engineering. MIT Press; 2023.
14. Hopkins A. Lessons from Esso's gas plant explosion at Longford. Warwick Pearse, Clare Gallagher and Liz Bluff. 2001; 1.
15. OSHA. Incident [Accident] Investigations: A Guide for Employers; 2015.
16. Rasmussen J. Risk management in a dynamic society: A modelling problem. Safety Sci. 1997; 27(2-3):183-213. https://doi.org/10.1016/S0925-7535(97)00052-0 PMCid:PMC10718422.
17. Stemn E, Hassall ME, Bofinger C. Systemic constraints to effective learning from incidents in the Ghanaian mining industry: A correspondence analysis and AcciMap approach. Safety Sci. 2020; 123:104565. https://doi.org/10.1016/j.ssci.2019.104565
18. Stemn E, Bofinger C, Cliff D, Hassall ME. Investigating the maturity of incident investigations of the Ghanaian mining industry and its effect on safety performance. Safety. 2019; 5(1):3. https://doi.org/10.3390/safety5010003
19. Underwood P, Waterson P. Systems thinking, the Swiss Cheese model and accident analysis: A comparative systemic analysis of the Grayrigg train derailment using the ATSB, AcciMap and STAMP models. Accid Anal Prev. 2014; 68:75-94. https://doi.org/10.1016/j.aap.2013.07.027 PMid:23973170.
20. Kumar A, Senapati A, Upadhyay R, Chatterjee S, Bhattacherjee A, Samanta B. An uncertainty-aware decision support system: Integrating text narratives and conformal prediction for trustworthy accident code classification. Process Saf Environ Prot. 2025: 108134. https://doi.org/10.1016/j.psep.2025.108134
21. Ganguli R, Miller P, Pothina R. Effectiveness of natural language processing-based machine learning in analyzing incident narratives at a mine. Minerals. 2021; 11(7):776. https://doi.org/10.3390/min11070776

22. Sharma A, Kumar A, Vardhan H, Mangalpady A, Mandal BB, Senapati A, *et al*. Human-in-the-loop data analytics for classifying fatal mining accident causes using natural language processing and machine learning techniques. Min Metall Explor. 2025; 42(6):4155-4167. https://doi.org/10.1007/s42461-025-01351-9

23. Grootendorst M. BERTopic: Neural topic modelling with a class-based TF-IDF procedure. arXiv preprint arXiv:220305794. 2022.

24. Blei DM, Ng AY, Jordan MI. Latent Dirichlet allocation. J Mach Learn Res. 2003; 3(Jan):993-1022.

25. George LE, Birla L, editors. A study of topic modelling methods. 2018 second international conference on intelligent computing and control systems (iciccs); 2018: IEEE. https://doi.org/10.1109/ICCONS.2018.8663152

26. Lee DD, Seung HS. Learning the parts of objects by non-negative matrix factorization. Nature. 1999; 401(6755):788-791. https://doi.org/10.1038/44565 PMid: 10548103.

27. Angelov D. Top2vec: Distributed representations of topics. arXiv preprint arXiv:200809470. 2020.

28. Egger R, Yu J. A topic modelling comparison between LDA, NMF, Top2Vec, and bertopic to demystify Twitter posts. Front Sociol. 2022; 7:886498. https://doi.org/10.3389/fsoc.2022.886498 PMid:35602001 PMCid:PMC9120935.

29. El-Gayar O, Al-Ramahi M, Wahbeh A, Nasralah T, Elnoshokaty A. A comparative analysis of the interpretability of LDA and LLM for topic modelling: The case of healthcare apps. 2024.

30. Ma L, Chen R, Ge W, Rogers P, Lyn-Cook B, Hong H, *et al*. AI-powered topic modelling: Comparing LDA and BERTopic in analyzing opioid-related cardiovascular risks in women. Exp Biol Med. 2025; 250:10389. https://doi.org/10.3389/ebm.2025.10389 PMid:40093658 PMCid: PMC11906279.

31. Albalawi R, Yeap TH, Benyoucef M. Using topic modelling methods for short-text data: A comparative analysis. Frontiers in artificial intelligence. 2020; 3:42. https://doi.org/10.3389/frai.2020.00042 PMid:33733159 PMCID: PMC7861298.

32. Röder M, Both A, Hinneburg A, editors. Exploring the space of topic coherence measures. Proceedings of the eighth ACM International Conference on Web Search and Data Mining; 2015. https://doi.org/10.1145/2684822.2685324

33. Bouma G. Normalized (pointwise) mutual information in collocation extraction. Proceedings of GSCL. 2009; 30:31-40.

34. Lau JH, Newman D, Baldwin T, editors. Machine reading tea leaves: Automatically evaluating topic coherence and topic model quality. Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics; 2014. https://doi.org/10.3115/v1/E14-1056

35. Mimno D, Wallach H, Talley E, Leenders M, McCallum A, editors. Optimizing semantic coherence in topic models. Proceedings of the 2011 conference on empirical methods in natural language processing; 2011.

36. Dieng AB, Ruiz FJ, Blei DM. Topic modelling in embedding spaces. Trans Assoc Comput Linguist. 2020; 8:439-453. https://doi.org/10.1162/tacl_a_00325 PMCid: PMC9607740.

37. Mitropoulos P, Abdelhamid TS, Howell GA. Systems model of construction accident causation. J Constr Eng Manage. 2005; 131(7):816-825. https://doi.org/10.1061/(ASCE)0733-9364(2005)131:7(816)

38. Leveson N. A new accident model for engineering safer systems. Safety Sci. 2004; 42(4):237-270. https://doi.org/10.1016/S0925-7535(03)00047-X PMCid: PMC10718422.

39. Hollnagel E. Safety-I and safety-II: The past and future of safety management: CRC Press; 2018. https://doi.org/10.1201/9781315607511

40. Hollnagel E. FRAM: the functional resonance analysis method: Modelling complex socio-technical systems. CRC Press; 2017.

41. Patriarca R, Chatzimichailidou M, Karanikas N, Di Gravio G. The past and present of System-Theoretic Accident Model and Processes (STAMP) and its associated techniques: A scoping review. Safety Sci. 2022; 146:105566. https://doi.org/10.1016/j.ssci.2021.105566