

A Scoping Literature Review of Natural Language Processing Application to Safety Occurrence Reports

Jon Ricketts , David Barry , Weisi Guo and Jonathan Pelham 

School of Aerospace, Transport & Manufacturing, Cranfield University, Cranfield MK43 0AL, UK

* Correspondence: j.ricketts@cranfield.ac.uk

Abstract: Safety occurrence reports can contain valuable information on how incidents occur, revealing knowledge that can assist safety practitioners. This paper presents and discusses a literature review exploring how Natural Language Processing (NLP) has been applied to occurrence reports within safety-critical industries, informing further research on the topic and highlighting common challenges. Some of the uses of NLP include the ability for occurrence reports to be automatically classified against categories, and entities such as causes and consequences to be extracted from the text as well as the semantic searching of occurrence databases. The review revealed that machine learning models form the dominant method when applying NLP, although rule-based algorithms still provide a viable option for some entity extraction tasks. Recent advances in deep learning models such as Bidirectional Transformers for Language Understanding are now achieving a high accuracy while eliminating the need to substantially pre-process text. The construction of safety-themed datasets would be of benefit for the application of NLP to occurrence reporting, as this would allow the fine-tuning of current language models to safety tasks. An interesting approach is the use of topic modelling, which represents a shift away from the prescriptive classification taxonomies, splitting data into “topics”. Where many papers focus on the computational accuracy of models, they would also benefit from real-world trials to further inform usefulness. It is anticipated that NLP will soon become a mainstream tool used by safety practitioners to efficiently process and gain knowledge from safety-related text.

Keywords: natural language processing; occurrence reporting; incident reporting; safety monitoring; safety management system



Citation: Ricketts, J.; Barry, D.; Guo, W.; Pelham, J. A Scoping Literature Review of Natural Language Processing Application to Safety Occurrence Reports. *Safety* **2023**, *9*, 22. <https://doi.org/10.3390/safety9020022>

Academic Editor: Raphael Grzebieta

Received: 13 February 2023

Revised: 21 March 2023

Accepted: 27 March 2023

Published: 5 April 2023



Copyright: © 2023 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Safety occurrence reporting systems used within safety-critical industries are capable of producing large quantities of textual data. In a typical sociotechnical system, these data will often contain a variety of information from technical issues through to organisational and cultural problems, assisting in the prevention of accidents. Presently, a lot of these data are reviewed by human beings to classify and identify relevant trends to improve safety. The advent of Natural Language Processing (NLP) has allowed machines to undertake this task, be able to automatically classify information and possibly extract knowledge from the reports [1–3].

NLP is a field of research overlapping computer science and artificial intelligence concerned with the ability to process natural languages; this generally consists of translating the natural language into data that a computer can use [4]. Present day computations on natural language are being undertaken using deep learning and machine learning techniques [5]. Machine learning involves the use of algorithms to parse data and learn from it, before making predictions and providing an output for a given task. Hence, the machine is “trained” on large amounts of data and algorithms that give it the ability to “learn”. Deep learning is considered a subset of machine learning, based upon neural networks. Neural networks consist of one or more layers of neurons, connected by weighted links to

take input data and produce an output. The “deep” term of deep learning is essentially taking these neural networks and increasing the layers and neurons. This is to create rich hierarchical representations by training neural networks with many hidden layers [6].

Early applications of NLP to safety occurrence and incident reports began with expression matching to highlight human factor concerns [7] through to classification, automatically identifying safety issues via a Support Vector Machine technique [8,9]. More recent papers recognise the specialist language used in many areas, deploying both topic modelling [10,11] and the state of the art machine learning and deep learning models [12–14].

There is currently an absence of comprehensive reviews on the application of NLP to occurrence reporting in safety. The aim of this paper is to explore the existing literature covering the application of NLP within safety occurrence reporting across multiple industries, identifying the computational methods deployed and associated challenges and limitations, informing future research on the application of NLP to safety occurrence data (in the context of this review, occurrence reporting is inclusive of incident reporting).

The main contribution of this paper is to present how and why NLP has been applied to safety occurrence reporting. The findings from this paper can assist safety practitioners to understand what approaches are available alongside their performance limits and challenges.

2. Method

This paper utilises a systematic review method [15] to identify and discuss academic papers that relate to the use of NLP within safety occurrence reporting.

In order to locate relevant papers, both the search terms/strings and databases need to be carefully selected. Safety occurrence reporting covers multiple industries (e.g., transport, medical and construction); therefore, the search encompasses all these industries for a full appreciation of how NLP may have been applied.

The databases selected for the search were: ScienceDirect, Scopus and Web of Science. These databases contain full, peer-reviewed papers while covering journals relevant to this literature review.

The search term (“NLP” OR “Natural Language Processing”) AND (“Report” OR “Occurrence”) AND “Safety” was used across the title, abstract and keywords. The addition of “Safety” to the search term dramatically reduced the quantity of search results, ensuring the analysis of the results was more manageable and relevant to the field of research. A further search string was created where NLP was replaced by “Text Mining”, from which, although returning duplicate results, several new and relevant articles were discovered. The results of the search strings are shown in Table 1.

Table 1. Search results from literature databases as of December 2022.

Database	(“NLP” or “Natural Language Processing”) and (“Report” or “Occurrence”) and “Safety”	“Text Mining” and (“Report” or “Occurrence”) and “Safety”
ScienceDirect	60	56
Web of Science	92	78
Scopus	306	223

After the removal of duplicates, the paper titles and abstracts were manually screened against an inclusion criteria that clearly bound this review, where the papers must match the following attributes:

- Original work.
- Full text is available.
- Written in English.
- NLP is specifically applied to safety occurrence reports.
- Published between 2012–2022.

As a further analysis to enable pearling research, the authors and citations of these papers were loaded into VOSviewer software [16] creating an interactive network map based on the bibliographic coupling of each document. The larger the size of the author node, the greater importance of that publication. The proximity of the nodes indicates the strength of the bibliographic content, while the links show the citations between the various papers. After application of the criteria and pearling research, 61 papers were left for review. The overall process is depicted in Figure 1.

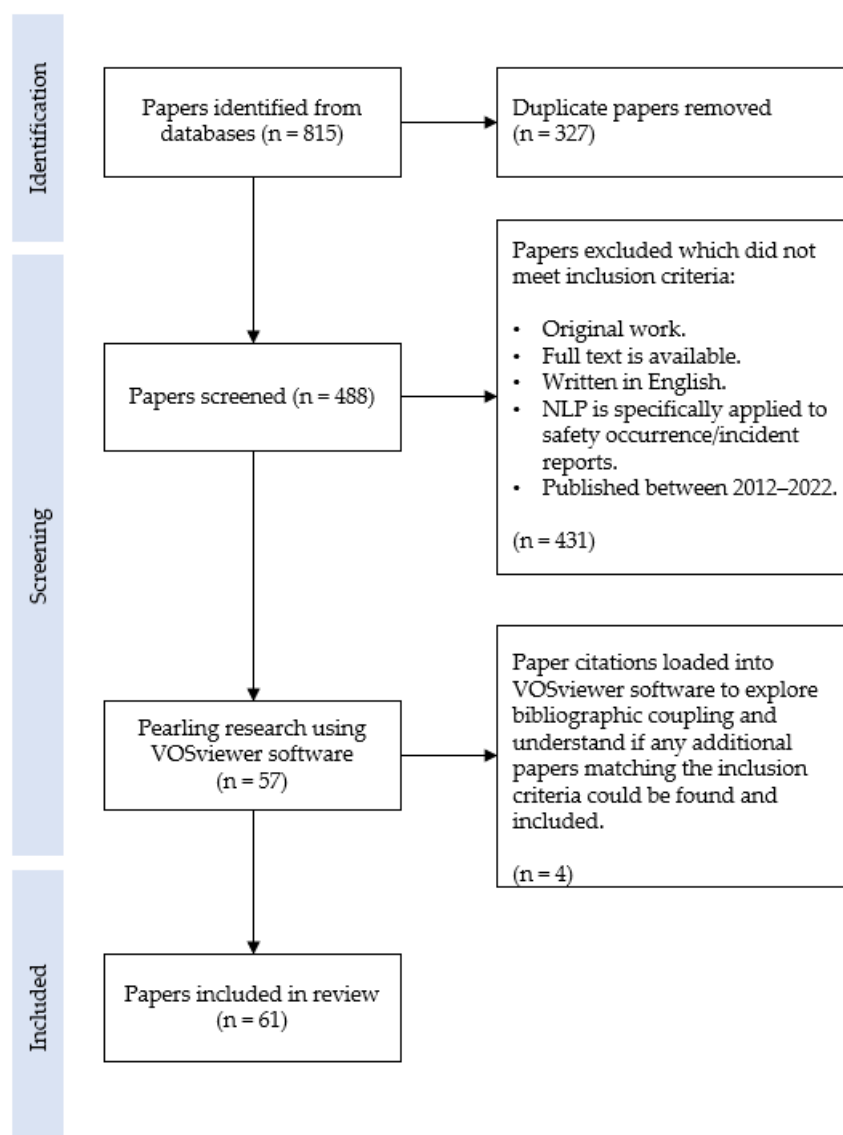


Figure 1. Overview of consecutive stages and results from literature review.

The papers were categorised against industry, general aim and computational method(s) used. Table 2 provides the definitions used to categorise the papers against an aim (typical NLP task) and computational method. If a paper featured multiple aims or methods, then these were recorded.


Table 2. Definitions for paper aims and computational methods used in this review.

Paper Categories		Definition
Paper Aim	Classification	Methods that seek to predict a category or class (e.g., assigning occurrence reports to given categories).
	Clustering	The partitioning of data into similar groups.
	Entity Extraction	The extraction of given entities from the text such as hazards, causes, consequences, etc.
	Injury Prediction	Forecasting injury based upon available data.
	Reveal Knowledge	Methods that focus on revealing knowledge from the data such as production of knowledge graphs or case-based reasoning methods.
	Risk Variables	Methods that explicitly highlight risks from the data and demonstrate risk relationships.
	Semantic Search	Ability to semantically search the data rather than traditional lexical searches.
	Text Summarisation	The summarisation of a larger body of text into a smaller, concise version.
	Topic Modelling	Topic modelling methods that seek to generate a number of topics from the data, providing an alternative method of analysis.
	Accident Prediction	Forecasting given accidents based on available data.
	Question and Answering	Methods that allow for specific questions to be answered from the data.
Computational Method	Database Cleansing	Methods used to improve database quality.
	Machine Learning	Any paper utilising machine learning methods that are defined as computational techniques enabling systems to learn from data or experience. Employing a set of statistical methods to find patterns in existing data and to then use patterns to make predictions [6].
	Deep Learning	Papers explicitly stating a deep learning method. Deep learning is a subset of machine learning creating rich hierarchical representations through the training of neural networks with many hidden layers [6].
	Rule-based algorithm	Methods that do not use machine learning but rather programmed rules to parse text and provide results.
	Ontology	Methods that explicitly state the development of an ontology. Ontologies generally describe taxonomic relationships [17].

3. Results

This section summarises the findings of the literature review with particular focus on the categorised aims of the individual papers and notable methods used.

Figure 2 shows the number of papers published each year and computational method used. It is clear there is an increasing trend of publications over time. Since 2020 there has been a shift from machine learning methods to deep learning methods.

**Figure 2.** Number of papers published each year featuring NLP and occurrence reporting.

A further insight was to understand what industries the papers covered (Figure 3). Half of the papers featured the aerospace and construction industries, while a quarter were formed of the medical and rail industries.




Figure 3. Number of papers per industry.

The aforementioned VOSviewer software was used to understand the citations between industries (Figure 4), and it was shown that papers featuring the construction industry were most heavily cited, followed by the aviation industry.

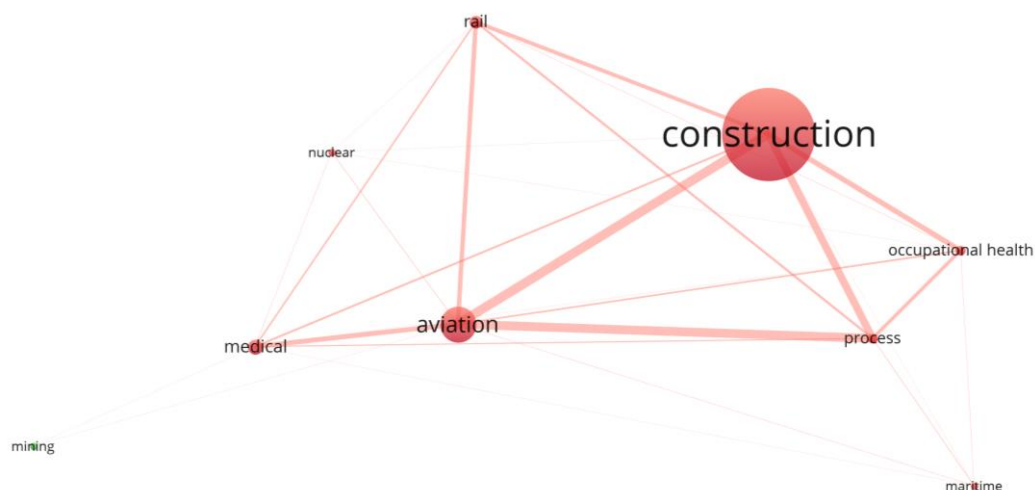


Figure 4. Network visualisation of papers grouped by industry and weighted by citations.

Each paper was assessed to understand its aim, the results of which are shown in Figure 5. Popular aims were to classify reports or extract entities such as causes and consequences. A few paper aims did not naturally fit to the categories defined in Table 2; therefore, the aims of these papers were recorded as “Visualise Safety Risk” [18] and “Similar Case Retrieval” [19]. The aim to reveal knowledge was broken down into “Knowledge Graph” [13] and “Knowledge Database” [20].




Figure 5. General aims of the papers.

Table 3 displays the papers associated with each categorical aim and computational method.

Table 3. Categories of each paper per aim and computational method.

Paper Categories		Papers
Paper Aim	Classification	[2,3,21–32]
	Clustering	[33–37]
	Entity Extraction	[1,12,14,17,32,38–51]
	Injury Prediction	[52]
	Reveal Knowledge	[13,20]
	Risk Variables	[53]
	Semantic Search	[54]
	Text Summarisation	[55]
	Topic Modelling	[10,11,56–58]
	Accident Prediction	[59,60]
	Question and Answering	[61]
	Visualise Safety Risk	[18]
	Similar Case Retrieval	[19]
	Database Cleansing	[62,63]
Computational Method	Deep Learning	[13,14,22,27,32,42,44,59,61,64,65]
	Machine Learning	[2,10–12,21,23,24,26,28–31,33–39,41,43,45–53,55–58,60,62,63,66,67]
	Rule-based algorithm	[1,3,19,20,38,40]
	Ontology	[17]

3.1. Classification

Classification of text is a common NLP application that can be applied to safety reporting/occurrence systems in that reports are typically classified against a given taxonomy for further analysis and reporting.

Tixier et al. [1] are major contributors to the research having been cited numerous times. Their study sought to automatically classify construction injury reports against a standard taxonomy (energy source, injury type, body part, injury severity). The method was based on hand-crafted rules and a keywords dictionary to extract outcomes and precursors from unstructured injury reports with over 95% accuracy.

A selection of papers [2,25,28,31,68] sought to classify occurrences against current system taxonomies (e.g., air safety reporting system), which is of benefit to current business and regulatory needs being that NLP would be able to quickly parse reports, while the alternative manual option would be too time-consuming.

The literature indicates that the machine learning Random Forest (RF) algorithm is a proven, high accuracy model for occurrence reporting classification. RF builds multiple decision trees and merges them together to gain an accurate prediction, which has been shown to achieve an accuracy of 80–93% when categorising aviation occurrences [2].

Although limited data are often an issue, in a study by Tanguy et al. [31], it did not prove problematic as runway excursion could be reliably classified while, on average, forming a small percentage of the overall occurrences. It was proposed that reports being classified with a precision of 95% or higher could be processed without human verification [31].

A deep neural network for classification using Universal Language Model Finetuning (ULMFiT) [28] comprising of a recurrent neural network and a classifier using a pretrained Wikipedia texts language model was fine-tuned with safety record narratives. It was predicted that with the increasing accessibility of NLP tools, they will soon form part of the safety analyst's standard toolset [28].

Closely linked to system taxonomy classification is the classification of occurrences to specific elements of the accident sequence, such as cause, type of incident and resultant effects. Bidirectional Encoder Representations from Transformers (BERT) has been used to automatically classify near-miss information [27]. BERT improves upon single word embedding models by taking into account the number of occurrences of a given word, for example, providing a different contextual embedding for homographs such as “a bat was used” and “a bat flew in”. In this instance, the BERT approach was able to achieve an accuracy of 86.9%. Recent papers feature hybrid approaches leveraging several computational methods to improve performance [13,22].

3.2. Entity Extraction

Entity extraction is where the NLP method can extract given entities (terms) from passages of text, for example, geographical places or people's names. In terms of safety reporting, the requirement could be to extract safety events, hazards, causes, etc. These could then be analysed in a more convenient form.

A training dataset underpins entity recognition models, where many safety activities would require the identification of bespoke entities. Fortunately, a number of software tools exist to create these datasets. One example is “APLenty” developed by the University of Manchester, which has been used to annotate hazards, consequences and mitigation strategies for construction safety [69]. The same methodology could be extended to other industries.

The extraction of pertinent information from occurrences was another theme identified within the papers. A natural language framework for automatic information extraction modelling, identifying features such as accident type, date, etc., has been proposed [44]. Risk factors have been extracted from accident reports with good results [45], while identifying causal relationships from reports has been shown to reduce manual workload [70]. In the medical industry, identifying harm events in patient care and categorising the harm event types based on their severity level has been undertaken [49].

A combined approach of rule-based gazetteers and machine learning has been conducted where occurrence reports were scanned for causes, consequences and hazards to validate a hazard identification artefact [38]. An entity recognition model trained on identifying causes and consequences then returned any new hazards not identified by the gazetteer.

A recent deep learning approach utilises a Long Short-Term Memory model to extract causal factors, being more accurate and adaptable than traditional machine learning methods [42].

3.3. Topic Modelling

Topic modelling is a collective term for a number of unsupervised machine learning models that capture meaning from a selection of documents. The development of topic modelling can be traced back to Latent Semantic Analysis (LSA) [71], which developed into Probabilistic LSA [72]. A further development was Latent Dirichlet Allocation (LDA), a generative probabilistic model operating via a three level Bayesian hierarchical model [73].

Topic modelling offers a different view to occurrence report analysis where the entire collection of reports can be divided into a chosen number of topics. This offers a more flexible alternative to the traditional classification taxonomies in use and enables emerging themes to be noticed. However, this does not mean that classification taxonomies are immediately redundant as supplemented with topic modelling, more insight can be gained from the data.

Topic modelling has been trialled on safety reports where it was determined that topic modelling was suitable for the data, with the majority of topics being relevant and independent from the metadata attributes [31]. Such a technique would be useful for data without a thorough classification scheme. Further work was carried out applying LDA to a fourteen-year sample of the ASRS database for temporal analysis [10]. This generated 200 topics, which were further verified by a panel of safety experts who declared that topic modelling would be useful for safety.

A different form of topic modelling has been conducted whereby feature word vectors of narrative text were obtained via Word2Vec training [56]. An LDA model was then used to map the latent semantic space, forming the document topic feature vectors of narrative text in a report. The approach yielded a marginally higher coherence score than LDA alone across a number of topics ranging from 1–20 [56].

Structural Topic Modelling (STM) intends to go further than basic topic modelling by highlighting links between aspects and certain conditions. For example, linking a particular failure to an aircraft type. STM was able to identify known issues and uncover previously unreported issues; however, it lacked the specific detail to direct action for which a human analyst is still required [11].

3.4. Semantic Search, Database Cleansing and Visualisation

A further promising area is the ability to create systems capable of performing semantic searches on databases. Semantic search refers to conducting a search that accounts for meaning and context, unlike classic lexical searches for literal term matches. The ability to semantically search records is of real use within safety engineering to understand past occurrences, learn from experience and provide question/answer-type responses in hazard identification activities.

Report similarity had been researched by two papers [19,20] showing promise from a safety analyst's perspective where a dataset could be interrogated to understand if an occurrence had previously happened without specific taxonomy labelling. The French Direction Generale de l'Aviation Civile (DGAC) trialled "timeplot" similarity software in an operational context as a temporal representation of similarity over time [31].

NLP has been used to improve search characteristics, including faceted search offering an intuitive retrieval of critical incident reports [74]. The combination of a keyword-based search and a semantic search resulted in good recall values.

Distilled BERT (40% fewer parameters than original BERT) has been used to provide answers from free text narratives to set questions [61]. A total of 70% of the questions were answered correctly, while further work was identified with training the model with safety expert feedback and investigating the use of more advanced models.

NLP has been used to improve and clean a safety report database where previously a significant effort would usually be required to identify, address, clean and repair data errors and inconsistencies [63].

The generation of visual analytics from close call reports, where words are shown as nodes with their relationships as links within a network, has been proposed [18]. The

technique was found to be useful for identifying risks in the small test set; however, the language differences from different groups of people in a larger dataset would be problematic. The study also recognized that significant contextual safety knowledge would be required by the analyst using this method and that the human is a vital part of the process.

4. Discussion

4.1. Key Challenge of Applying NLP to Safety Occurrence Reports

It can be argued that the biggest challenge facing the application of NLP to safety occurrence reporting is the textual data characteristics. It can be expected that a given occurrence report will feature a free text field for the reporter to enter information about the occurrence. This is often the valuable data for safety investigations (and NLP) where the incident and its surrounding circumstances are described, revealing causes, hazards and other factors that can be used to continually improve safety. The free text field can then be further enhanced by additional data such as times, dates, temperature, location, etc. How these data are analysed depends very much on the industry and task at hand.

As an example, an extract from an occurrence held on the Aviation Safety Reporting System (ASRS) is shown below:

“DURING FINAL APCH TO LNDG ZONE, R-HAND ENG COWLING EXITED ACFT STRIKING MAIN ROTOR BLADE AND REAR CTR MAIN VERT STABILIZER. THE SHATTERED COWLING DROPPED TO GND IN PIECES APPROX 4 BLOCKS NNE OF THE LNDG ZONE CAUSING NO INJURIES OR PROPERTY DAMAGE.” [75]

The extract describes an engine cowling detaching from a helicopter, hitting the main rotor blades and vertical stabiliser. While the description will make sense to those in the aviation industry, those who are unfamiliar may struggle with the terse language and number of abbreviations scattered throughout; APCH—approach, ACFT—aircraft and CTR—centre. This goes to show that not only does the safety practitioner dissecting these reports need a grounding in safety theory, but they also need to have a good understanding of the industry, and its operations and technical terminology. Likewise, NLP needs to reflect this. Although the above example focuses on an aviation occurrence report, the same issues are present within other industries.

Freely available NLP tools and models are usually trained on vast amounts of text such as Wikipedia pages, and therefore have not encountered industry-specific terminology. This ensures that the processing of industry/safety-specific text (such as the example above) to provide useful responses can be inaccurate. For many safety activities, accuracy is vital, as the results can influence safety-related decision making.

In order to overcome the aforementioned challenge with the data, a couple of options for NLP machine learning models are:

1. Fine-tune the model. The “standard” model is further trained on a specific dataset (e.g., collection of safety assessment reports) [42,76].
2. Train model from scratch. The model is trained on the safety-specific data, although this is where the second challenge is presented: quantity. If we take BERT as an example, this was trained on 3300 M words [77]. Unless the organisation has an equally large repository of information or is able to accumulate data from a number of regulators, then it is unlikely to match a similar level of data input.

4.2. Common Issues When Applying NLP to Safety Occurrence Reports

From reviewing the papers discussed above, it is possible to draw out some of the main challenges when applying NLP to safety occurrence reporting. These are listed in Table 4 (the list is not exhaustive but can be used as a starting point for NLP projects).

Table 4. Common challenges when applying NLP to safety occurrence reporting.

Challenge	Potential Solution
Use of language/semantics including the use of acronyms and spelling errors, which can confuse algorithms and require extensive effort to normalise text prior to machine learning.	Use the data to train a model from scratch or fine-tune a model, enabling the model to “learn” the new terminology. Alternatively, standardise the text by parsing it through a bespoke dictionary of acronyms and domain specific terms.
Language can differ across a single organisation/domain.	As above, standardisation rules can also be applied to reduce the text into common terms.
Contextual safety knowledge is often required to understand if the results are useful.	Incorporate the knowledge of domain safety experts through review or workshops. Construct bespoke datasets to capture context and feed into machine learning models.
Data cleaning itself can require significant effort at the start of a project.	Allocate enough time to clean and normalise data at the start of the project. As above, handwritten rules can be used to speed up this process, organising the text into appropriate formats for onward processing.
Model overfitting leading to erroneous results.	Depends upon the language model; however, one aim is to reduce the amount of “noise” in the data and ensuring training data are appropriate.
Classification errors occur between labels that share similar expressions.	Analyse model output samples and fine-tune parameters of the model.
Data may fit multiple categories, adding complexity to the machine learning model.	Consider using a multi-classifier machine learning model.
Component failure can be difficult to recognize, given that it can form part of a wider event leading to surplus information that detracts the classifier from the actual cause.	As suggested by Tanguy et al. [31] “build a relationship with the data” taking time to understand what is required and adapt the model accordingly.
Care needs to be taken to avoid bias and properly train/maintain models.	Algorithmic bias is unavoidable; however, it can be reduced by targeted sampling or re-weighting. The “UnBias: Emancipating Users Against Algorithmic Biases for a Trusted Digital Economy” project offers solutions and tools to reduce bias [78].
The results are only as good as the training data. Therefore it is important to ensure the training data are accurate.	Invest time and resources at the start of the project to cleanse and check the training data are suitable for the task.
Incident reports typically only detail “what went wrong”. For a balanced view, knowledge of what went well is required.	Dependent upon the safety system in use and if data on “successes” are recorded. Alternatively, data could be gathered from employees as to what safety mitigations work well.
Distrust in model outputs or unable to achieve high levels of accuracy.	Dependent upon the intended use of the output, due to the nature of safety engineering, the model may be used initially to enhance a safety practitioners’ role as a support tool. Evaluation of the completed model can be carried out via case studies using experienced safety practitioners.
The data might be imbalanced where one of the classes (minority class) contains a much smaller number of examples than the remaining class (majority class).	Undersampling can be used to remove some instances of the majority class or oversampling to create new instances of the minority class.

4.3. Limitations

Having covered the application of NLP described in academic papers, this section highlights some of the main limitations, and what is required if NLP is to become a regular tool in safety occurrence reporting.

4.3.1. Lack of Safety-Focused Datasets

A core element of many language models is a defined dataset on which it has been trained. This can simply be a large corpus of text or it may be a bespoke dataset created

to solve a particular task. In the case of entity recognition, this would be passages of text with annotated entities. These datasets are resource extensive to create. For example, to create a safety-specific entity recognition dataset, several safety practitioners with the requisite knowledge would be required to annotate the text, a time-consuming and therefore expensive task. If interest continues to grow in deploying NLP to safety activities, then there may be a shift to creating such datasets, although these would likely be industry-specific due to the differing use of language. Other factors such as model drift will play a role: if a language model has previously been based upon a crewed aviation theme, over time, uncrewed aircraft have become more prevalent. The model in this case would not be ideally adjusted for the new terminology and language that comes with uncrewed aircraft, most likely leading to poor results. Therefore, model drift would need to be monitored and models re-baselined with current data.

Closely linked with the creation of datasets is the quality of the data. The raw text data may need to be pre-processed prior to a language model phase, in which case the output needs to be accurate and repeatable. Likewise, bespoke datasets need to be fit for purpose. Where vast generic datasets are created through the use of Amazon Turk workers [79] or equivalent, this is not possible within safety engineering as the workers require knowledge of safety and the intended industry.

4.3.2. Model Evaluation beyond Metrics

Depending on how extensive the use of NLP will be, the performance needs to be assessed beyond typical machine learning metrics such as accuracy, precision and F1 measures. Feedback from safety practitioners is invaluable to truly assess usefulness. For example, a question and answering machine learning model may exhibit poor performance in terms of computational metrics, such as not extracting exact (computer-anticipated) answers. However, when used in an operational context by safety practitioners the generated answers may capture all that is required to be useful.

4.3.3. Trustworthiness and Model Interpretability

A further consideration is around trustworthiness and the integration of NLP tools with safety practitioners. “Trust” is typically placed in systems that demonstrate repeatable behaviour and performance to deliver successful outcomes; a single failure can start to erode this trust. From the wider perspective, machine learning technologies often fail to live up to our expectations through being inaccurate, unreliable and discriminatory [80]. Within safety reporting, the results of a given NLP system might influence safety decisions or risk to life, therefore requiring an element of trust.

Machine learning models can be treated as a “black box” where the internal workings are not fully known [81]. This is likely to be unacceptable for many safety tasks where the rationale behind outputs made by the model need to be clear, especially if risk to life is involved or the output needs to be traceable. In this case, the system could be limited to a “decision support role”, or if the setup allows, to substantiate outputs with evidence. For example, supplying the document reference and passage of text that a generated answer has come from.

4.3.4. Data Protection

The final limitation considered within this paper is that of data protection. Machine learning models are often data-hungry and require vast amounts of occurrence reports to improve performance. Where some databases are publicly available (e.g., ASRS), others are not, and due to data privacy policies, any models trained on these data may not be made public. It would be encouraging to see larger repositories of occurrence data available for public use in the future; this could be provided by regulatory bodies in a format cleansed of personal data.

5. Conclusions

This paper introduces the topic of NLP within the occurrence reporting context while providing a review of the research to date.

The latest deep learning developments such as BERT are starting to be introduced in the most recent papers with promising results. Where the majority of papers discuss language models trained upon large datasets such as Wikipedia pages, there appear to be few specific safety-focused datasets available. This is likely due to the unique field and the fact that the majority of safety databases/repositories are unlikely to match datasets based upon Wikipedia in terms of sheer size and variety. The construction of safety-themed datasets going forward would be of benefit to the application of NLP to occurrence reporting, as this will allow the fine-tuning of current language models to safety tasks.

Semantic search appears to be an area of development where a small portion of the reviewed papers addressed this [54,61]. At the time of writing, OpenAI's ChatGPT [82] has received extensive media coverage with its ability to understand and answer natural questions and solving tasks such as writing computer code based on prompts. However, it is limited by only being trained on a knowledge base up until 2021, being susceptible to hallucination (where it produces an incorrect but plausible response) and producing verbose responses [83]. ChatGPT or other Generative Pretrained Transformer (GPT) models could fulfil a semantic search capability for safety documentation (e.g., a collection of safety assessments or an incident database).

Going forward, it would be satisfying to see regulatory bodies starting to encourage the use of NLP, although caution should be taken that it is not a case of "one size fits all" and the application should be tailored to the given task. There could be a temptation to fully automate typical safety processes (e.g., monitoring trends or indicators) through NLP. To some extent this is possible; however, the quality of the encompassing safety management system is important, alongside organisational understanding and motivation as these aspects will drive improvements [84].

The authors do not envisage the use of NLP making the safety practitioner redundant but rather offering an insightful tool to aid their work and increase efficiency. It is anticipated that standard NLP tools and methods will soon be used to assist in many safety activities from the generation of artefacts to ongoing safety monitoring.

Author Contributions: J.R., D.B., W.G. and J.P. contributed to the conception of the literature review. J.R. retrieved and screened the papers. J.R. wrote the manuscript. D.B., W.G. and J.P. provided supervision and reviewed the manuscript, providing feedback and corrections. All authors have read and agreed to the published version of the manuscript.

Funding: The APC was funded by UKRI.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: No new data were created or analysed in this study. Data sharing is not applicable to this article.

Acknowledgments: J Ricketts thanks the contribution of the IMechE Whitworth Senior Scholarship Award and sponsorship of BAE Systems.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Tixier, A.J.P.; Hollowell, M.R.; Rajagopalan, B.; Bowman, D. Automated Content Analysis for Construction Safety: A Natural Language Processing System to Extract Precursors and Outcomes from Unstructured Injury Reports. *Autom. Constr.* **2016**, *62*, 45–56. [CrossRef]
2. De Vries, V. Classification of Aviation Safety Reports Using Machine Learning. In Proceedings of the 2020 International Conference on Artificial Intelligence and Data Analytics for Air Transportation, AIDA-AT 2020, Singapore, 3–4 February 2020; IEEE: Piscataway, NJ, USA, 2020.

3. Hughes, P.; Shipp, D.; Figueres-Esteban, M.; van Gulijk, C. From Free-Text to Structured Safety Management: Introduction of a Semi-Automated Classification Method of Railway Hazard Reports to Elements on a Bow-Tie Diagram. *Saf. Sci.* **2018**, *110*, 11–19. [\[CrossRef\]](#)
4. Lane, H.; Howard, C.; Hapke, H. *Natural Language Processing in Action*; Manning Publications Co.: Shelter Island, NY, USA, 2019; ISBN 9781617294631.
5. Ghosh, S.; Gunning, D. *Natural Language Processing Fundamentals*; Packt Publishing: Birmingham, UK, 2019.
6. ISO 22989:2022(E); Information Technology—Artificial Intelligence—Artificial Intelligence Concepts and Terminology. International Organization for Standardization: Geneva, Switzerland, 2022.
7. Posse, C.; Matzke, B.; Anderson, C.; Brothers, A.; Matzke, M.; Ferryman, T. Extracting Information from Narratives: An Application to Aviation Safety Reports. In Proceedings of the IEEE Aerospace Conference, Big Sky, MT, USA, 5–12 March 2005; IEEE: Piscataway, NJ, USA, 2005; pp. 3678–3690.
8. Oza, N.; Castle, J.P.; Stutz, J. Classification of Aeronautics System Health and Safety Documents. *IEEE Trans. Syst. Man Cybern. Part C Appl. Rev.* **2009**, *39*, 670–680. [\[CrossRef\]](#)
9. Wolfe, S. Wordplay: An Examination of Semantic Approaches to Classify Safety Reports. In Proceedings of the AIAA Infotech@Aerospace 2007 Conference and Exhibit, Rohnert Park, CA, USA, 7–10 May 2007.
10. Robinson, S.D. Temporal Topic Modeling Applied to Aviation Safety Reports: A Subject Matter Expert Review. *Saf. Sci.* **2019**, *116*, 275–286. [\[CrossRef\]](#)
11. Kuhn, K.D. Using Structural Topic Modeling to Identify Latent Topics and Trends in Aviation Incident Reports. *Transp. Res. Part C Emerg. Technol.* **2018**, *87*, 105–122. [\[CrossRef\]](#)
12. Baker, H.; Hallowell, M.R.; Tixier, A.J.P. Automatically Learning Construction Injury Precursors from Text. *Autom. Constr.* **2020**, *118*, 103145. [\[CrossRef\]](#)
13. Liu, C.; Yang, S. Using Text Mining to Establish Knowledge Graph from Accident/Incident Reports in Risk Assessment. *Expert Syst. Appl.* **2022**, *207*, 117991. [\[CrossRef\]](#)
14. Rybak, N.; Hassall, M. Deep Learning Unsupervised Text-Based Detection of Anomalies in U.S. Chemical Safety and Hazard Investigation Board Reports. In Proceedings of the International Conference on Electrical, Computer, Communications and Mechatronics Engineering, ICECCME 2021, Mauritius, 7–8 October 2021; IEEE: Piscataway, NJ, USA, 2021; pp. 7–8.
15. Denyer, D.; Tranfield, D. Producing a Systematic Review. In *The Sage Handbook of Organizational Research Methods*; Sage Publications Ltd.: Thousand Oaks, CA, USA, 2009; pp. 671–689. ISBN 978-1-4129-3118-2.
16. Perianes-Rodriguez, A.; Waltman, L.; van Eck, N.J. Constructing Bibliometric Networks: A Comparison between Full and Fractional Counting. *J. Informetr.* **2016**, *10*, 1178–1195. [\[CrossRef\]](#)
17. Hughes, P.; Robinson, R.; Figueres-Esteban, M.; van Gulijk, C. Extracting Safety Information from Multi-Lingual Accident Reports Using an Ontology-Based Approach. *Saf. Sci.* **2019**, *118*, 288–297. [\[CrossRef\]](#)
18. Figueres-Esteban, M.; Hughes, P.; van Gulijk, C. Visual Analytics for Text-Based Railway Incident Reports. *Saf. Sci.* **2016**, *89*, 72–76. [\[CrossRef\]](#)
19. Fan, H.; Li, H. Retrieving Similar Cases for Alternative Dispute Resolution in Construction Accidents Using Text Mining Techniques. *Autom. Constr.* **2013**, *34*, 85–91. [\[CrossRef\]](#)
20. Wu, H.; Zhong, B.; Medjdoub, B.; Xing, X.; Jiao, L. An Ontological Metro Accident Case Retrieval Using CBR and NLP. *Appl. Sci.* **2020**, *10*, 5298. [\[CrossRef\]](#)
21. Hou, Q.; Wang, L.; Yuan, T. Research on Automatic Classifying Method for Incident Reports with Runway Incursion. In Proceedings of the 4th International Conference on Information Science, Electrical, and Automation Engineering (ISEAE 2022), Guangzhou, China, 1 August 2022; p. 122573T. [\[CrossRef\]](#)
22. Zhang, F. A Hybrid Structured Deep Neural Network with Word2Vec for Construction Accident Causes Classification. *Int. J. Constr. Manag.* **2022**, *22*, 1120–1140. [\[CrossRef\]](#)
23. Madeira, T.; Melício, R.; Valério, D.; Santos, L. Machine Learning and Natural Language Processing for Prediction of Human Factors in Aviation Incident Reports. *Aerospace* **2021**, *8*, 47. [\[CrossRef\]](#)
24. Evans, H.P.; Anastasiou, A.; Edwards, A.; Hibbert, P.; Makeham, M.; Luz, S.; Sheikh, A.; Donaldson, L.; Carson-Stevens, A. Automated Classification of Primary Care Patient Safety Incident Report Content and Severity Using Supervised Machine Learning (ML) Approaches. *Health Inform. J.* **2020**, *26*, 3123–3139. [\[CrossRef\]](#) [\[PubMed\]](#)
25. Goodrum, H.; Roberts, K.; Bernstam, E.V. Automatic Classification of Scanned Electronic Health Record Documents. *Int. J. Med. Inform.* **2020**, *144*, 104302. [\[CrossRef\]](#) [\[PubMed\]](#)
26. Cheng, M.-Y.; Kusoemo, D.; Gosno, R.A. Text Mining-Based Construction Site Accident Classification Using Hybrid Supervised Machine Learning. *Autom. Constr.* **2020**, *118*, 103265. [\[CrossRef\]](#)
27. Fang, W.; Luo, H.; Xu, S.; Love, P.E.D.; Lu, Z.; Ye, C. Automated Text Classification of Near-Misses from Safety Reports: An Improved Deep Learning Approach. *Adv. Eng. Inform.* **2020**, *44*, 101060. [\[CrossRef\]](#)
28. Marev, K.; Georgiev, K. Automated Aviation Occurrences Categorization. In Proceedings of the ICMT 2019—7th International Conference on Military Technologies, Brno, Czech Republic, 30–31 May 2019; IEEE: Piscataway, NJ, USA, 2019; pp. 1–5.
29. Zhang, F.; Fleyeh, H.; Wang, X.; Lu, M. Construction Site Accident Analysis Using Text Mining and Natural Language Processing Techniques. *Autom. Constr.* **2019**, *99*, 238–248. [\[CrossRef\]](#)

30. Heidarysafa, M.; Kowsari, K.; Barnes, L.; Brown, D. Analysis of Railway Accidents' Narratives Using Deep Learning. In Proceedings of the 17th IEEE International Conference on Machine Learning and Applications, ICMLA 2018, Orlando, FL, USA, 17–20 December 2018; IEEE: Piscataway, NJ, USA, 2019; pp. 1446–1453.
31. Tanguy, L.; Tulechki, N.; Urieli, A.; Hermann, E.; Raynal, C. Natural Language Processing for Aviation Safety Reports: From Classification to Interactive Analysis. *Comput. Ind.* **2016**, *78*, 80–95. [\[CrossRef\]](#)
32. Jidkov, V.; Abielmona, R.; Teske, A. PE Enabling Maritime Risk Assessment Using Natural Language Processing-Based Deep Learning Techniques. In Proceedings of the 2020 IEEE Symposium Series on Computational Intelligence, SSCI 2020, Canberra, Australia, 1–4 December 2020; pp. 2469–2476.
33. Miyamoto, A.; Bendarkar, M.V.; Mavris, D.N. Natural Language Processing of Aviation Safety Reports to Identify Inefficient Operational Patterns. *Aerospace* **2022**, *9*, 450. [\[CrossRef\]](#)
34. Rose, R.L.; Puranik, T.G.; Mavris, D.N. Natural Language Processing Based Method for Clustering and Analysis of Aviation Safety Narratives. *Aerospace* **2020**, *7*, 143. [\[CrossRef\]](#)
35. Liu, J.; Wong, Z.S.Y.; Tsui, K.L.; So, H.Y.; Kwok, A. Exploring Hidden In-Hospital Fall Clusters from Incident Reports Using Text Analytics. *Stud. Health Technol. Inform.* **2019**, *264*, 1526–1527. [\[CrossRef\]](#) [\[PubMed\]](#)
36. Chokor, A.; Naganathan, H.; Chong, W.K.; Asmar, M. El Analyzing Arizona OSHA Injury Reports Using Unsupervised Machine Learning. *Procedia Eng.* **2016**, *145*, 1588–1593. [\[CrossRef\]](#)
37. Tirunagari, S.; Hanninen, M.; Stahlberg, K.; Kujala, P. Mining Causal Relations and Concepts in Maritime. In Proceedings of the TechSamudra 2012, International Conference cum Exhibition on Technology of the Sea, Visakhapatnam, India, 6–8 December 2012; Volume 1, pp. 548–566.
38. Ricketts, J.; Pelham, J.; Barry, D.; Guo, W. An NLP Framework for Extracting Causes, Consequences, and Hazards from Occurrence Reports to Validate a HAZOP Study. In Proceedings of the 2022 IEEE/AIAA 41st Digital Avionics Systems Conference (DASC), Portsmouth, VA, USA, 18–22 September 2022; IEEE: Portsmouth, VA, USA, 2022; pp. 1–8.
39. Liu, G.; Boyd, M.; Yu, M.; Halim, S.Z.; Quddus, N. Identifying Causality and Contributory Factors of Pipeline Incidents by Employing Natural Language Processing and Text Mining Techniques. *Process Saf. Environ. Prot.* **2021**, *152*, 37–46. [\[CrossRef\]](#)
40. Shekhar, H.; Agarwal, S. Automated Analysis through Natural Language Processing of DGMS Fatality Reports on Indian Non-Coal Mines. In Proceedings of the 5th International Conference on Information Systems and Computer Networks, ISCON 2021, Mathura, India, 22–23 October 2021; IEEE: Piscataway, NJ, USA, 2021; pp. 1–6.
41. Valcamonico, D.; Baraldi, P.; Zio, E. Natural Language Processing and Bayesian Networks for the Analysis of Process Safety Events. In Proceedings of the 2021 5th International Conference on System Reliability and Safety, ICSRS 2021, Palermo, Italy, 24–26 November 2021; IEEE: Piscataway, NJ, USA, 2021; pp. 216–221.
42. Dong, T.; Yang, Q.; Ebadi, N.; Luo, X.R.; Rad, P. Identifying Incident Causal Factors to Improve Aviation Transportation Safety: Proposing a Deep Learning Approach. *J. Adv. Transp.* **2021**, *2021*, 5540046. [\[CrossRef\]](#)
43. Wang, G.; Liu, M.; Cao, D.; Tan, D. Identifying High-Frequency–Low-Severity Construction Safety Risks: An Empirical Study Based on Official Supervision Reports in Shanghai. *Eng. Constr. Archit. Manag.* **2021**, *29*, 940–960. [\[CrossRef\]](#)
44. Feng, D.; Chen, H. A Small Samples Training Framework for Deep Learning-Based Automatic Information Extraction: Case Study of Construction Accident News Reports Analysis. *Adv. Eng. Inform.* **2021**, *47*, 101256. [\[CrossRef\]](#)
45. Hua, L.; Zheng, W.; Gao, S. Extraction and Analysis of Risk Factors from Chinese Railway Accident Reports. In Proceedings of the 2019 IEEE Intelligent Transportation Systems Conference, ITSC 2019, Auckland, New Zealand, 27–30 October 2019; IEEE: Piscataway, NJ, USA, 2019; pp. 869–874.
46. Zhao, Y.; Diao, X.; Huang, J.; Smidts, C. Automated Identification of Causal Relationships in Nuclear Power Plant Event Reports. *Nucl. Technol.* **2019**, *205*, 1021–1034. [\[CrossRef\]](#)
47. Song, B.; Suh, Y. Narrative Texts-Based Anomaly Detection Using Accident Report Documents: The Case of Chemical Process Safety. *J. Loss Prev. Process Ind.* **2019**, *57*, 47–54. [\[CrossRef\]](#)
48. Zhao, Y.; Diao, X.; Smidts, C. Preliminary Study of Automated Analysis of Nuclear Power Plant Event Reports Based on Natural Language Processing Techniques. In Proceedings of the Probabilistic Safety Assessment and Management PSAM 14, Los Angeles, CA, USA, 16–21 September 2018.
49. Cohan, A.; Ratwani, R.; Fong, A.; Goharian, N. Identifying Harm Events in Clinical Care through Medical Narratives. In Proceedings of the 8th ACM International Conference on Bioinformatics, Computational Biology and Health Informatics, Boston, MA, USA, 20–23 August 2017; pp. 52–59. [\[CrossRef\]](#)
50. Fong, A.; Harriott, N.; Walters, D.M.; Foley, H.; Morrissey, R.; Ratwani, R.R. Integrating Natural Language Processing Expertise with Patient Safety Event Review Committees to Improve the Analysis of Medication Events. *Int. J. Med. Inform.* **2017**, *104*, 120–125. [\[CrossRef\]](#)
51. Tixier, A.J.P.; Hollowell, M.R.; Rajagopalan, B.; Bowman, D. Construction Safety Clash Detection: Identifying Safety Incompatibilities among Fundamental Attributes Using Data Mining. *Autom. Constr.* **2017**, *74*, 39–54. [\[CrossRef\]](#)
52. Tixier, A.J.P.; Hollowell, M.R.; Rajagopalan, B.; Bowman, D. Application of Machine Learning to Construction Injury Prediction. *Autom. Constr.* **2016**, *69*, 102–114. [\[CrossRef\]](#)

53. Wang, Z.; Yin, J. Risk Assessment of Inland Waterborne Transportation Using Data Mining. *Marit. Policy Manag.* **2020**, *47*, 633–648. [CrossRef]
54. Denecke, K. Concept-Based Retrieval from Critical Incident Reports. *Stud. Health Technol. Inform.* **2017**, *236*, 1–7. [CrossRef] [PubMed]
55. Zhao, Z.; Yang, Y.; Wang, Y.; Zhang, J.; Wang, D.; Luo, X. Summarization of Coal Mine Accident Reports: A Natural-Language-Processing-Based Approach. *Commun. Comput. Inf. Sci.* **2020**, *1329*, 103–115. [CrossRef]
56. Luo, Y.; Shi, H. Using Lda2vec Topic Modeling to Identify Latent Topics in Aviation Safety Reports. In Proceedings of the 2019 IEEE/ACIS 18th International Conference on Computer and Information Science (ICIS), Beijing, China, 17–19 June 2019; pp. 518–523. [CrossRef]
57. Kuhn, K.D. Topics and Trends in Incident Reports Using Structural Topic Modeling to Explore Aviation Safety Reporting System Data. In Proceedings of the 12th USA/EUROPE Air Traffic Management R&D Seminar, Seattle, WA, USA, 27–30 June 2017.
58. Robinson, S.D. Visual Representation of Safety Narratives. *Saf. Sci.* **2016**, *88*, 123–128. [CrossRef]
59. Zhang, X.; Srinivasan, P.; Mahadevan, S. Sequential Deep Learning from NTSB Reports for Aviation Safety Prognosis. *Saf. Sci.* **2021**, *142*, 105390. [CrossRef]
60. Baker, H.; Hallowell, M.R.; Tixier, A.J.P. AI-Based Prediction of Independent Construction Safety Outcomes from Universal Attributes. *Autom. Constr.* **2020**, *118*, 103146. [CrossRef]
61. Kierszbaum, S.; Lapasset, L. Applying Distilled BERT for Question Answering on ASRS Reports. In Proceedings of the 2020 New Trends in Civil Aviation (NTCA), Prague, Czech Republic, 23–24 November 2020; pp. 33–38. [CrossRef]
62. Macedo, J.B.; Ramos, P.M.S.; Maior, C.B.S.; Moura, M.J.C.; Lins, I.D.; Vilela, R.F.T. Identifying Low-Quality Patterns in Accident Reports from Textual Data. *Int. J. Occup. Saf. Ergon.* **2022**. [CrossRef]
63. Dorsey, L.C.; Wang, B.; Grabowski, M.; Merrick, J.; Harrald, J.R. Self Healing Databases for Predictive Risk Analytics in Safety-Critical Systems. *J. Loss Prev. Process Ind.* **2020**, *63*, 104014. [CrossRef]
64. Ramos, P.; Macêdo, J.B.; Maior, C.B.S.; Moura, M.C.; Lins, I.D. Combining BERT with Numerical Features to Classify Injury Leave Based on Accident Description. *Proc. Inst. Mech. Eng. Part O J. Risk Reliab.* **2022**, 1–12. [CrossRef]
65. Kierszbaum, S.; Klein, T.; Lapasset, L. ASRS-CMFS vs. RoBERTa: Comparing Two Pre-Trained Language Models to Predict Anomalies in Aviation Occurrence Reports with a Low Volume of In-Domain Data Available. *Aerospace* **2022**, *9*, 591. [CrossRef]
66. Jiao, Y.; Dong, J.; Han, J.; Sun, H. Classification and Causes Identification of Chinese Civil Aviation Incident Reports. *Appl. Sci.* **2022**, *12*, 10765. [CrossRef]
67. Gillespie, A.; Reader, T.W. Online Patient Feedback as a Safety Valve: An Automated Language Analysis of Unnoticed and Unresolved Safety Incidents. *Risk Anal.* **2022**, 1–15. [CrossRef] [PubMed]
68. Wong, Z.S.Y.; So, H.Y.; Kwok, B.S.C.; Lai, M.W.S.; Sun, D.T.F. Medication-Rights Detection Using Incident Reports: A Natural Language Processing and Deep Neural Network Approach. *Health Inform. J.* **2020**, *26*, 1777–1794. [CrossRef]
69. Thompson, P.; Yates, T.; Inan, E.; Ananiadou, S. Semantic Annotation for Improved Safety in Construction Work. In Proceedings of the Twelfth Language Resources and Evaluation Conference, Marseille, France, 11–16 May 2020; pp. 1990–1999.
70. Han, L.; Ball, R.; Pamer, C.A.; Altman, R.B.; Proestel, S. Development of an Automated Assessment Tool for MedWatch Reports in the FDA Adverse Event Reporting System. *J. Am. Med. Inform. Assoc.* **2017**, *24*, 913–920. [CrossRef]
71. Deerwester, S.; Furnas, G.W.; Landauer, T.K.; Harshman, R. Indexing by Latent Semantic Analysis Scott. *J. Am. Soc. Inf. Sci.* **1990**, *41*, 391–407. [CrossRef]
72. Hofmann, T. Unsupervised Learning by Probabilistic Latent Semantic Analysis. *Mach. Learn.* **2001**, *42*, 177–196. [CrossRef]
73. Blei, D.M.; Ng, A.Y.; Jordan, M.I. Latent Dirichlet Allocation. *J. Mach. Learn. Res.* **2003**, *3*, 993–1022.
74. Denecke, K. Automatic Analysis of Critical Incident Reports: Requirements and Use Cases. *Stud. Health Technol. Inform.* **2016**, *223*, 85–92. [CrossRef] [PubMed]
75. ASRS Report ACN 353289; ASRS: Kitty Hawk, NC, USA, 1996.
76. Macêdo, J.B.; das Chagas Moura, M.; Aichele, D.; Lins, I.D. Identification of Risk Features Using Text Mining and BERT-Based Models: Application to an Oil Refinery. *Process Saf. Environ. Prot.* **2022**, *158*, 382–399. [CrossRef]
77. Devlin, J.; Chang, M.W.; Lee, K.; Toutanova, K. BERT: Pre-Training of Deep Bidirectional Transformers for Language Understanding. In Proceedings of the NAACL HLT 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Minneapolis, MN, USA, 2–7 June 2019; Volume 1, pp. 4171–4186.
78. Unbias Project. Available online: <https://unbias.wp.horizon.ac.uk/> (accessed on 14 September 2020).
79. Saeidi, M.; Bartolo, M.; Lewis, P.; Singh, S.; Rocktäschel, T.; Sheldon, M.; Bouchard, G.; Riedel, S. Interpretation of Natural Language Rules in Conversational Machine Reading. In Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, EMNLP 2018, Brussels, Belgium, 31 October–4 November 2018; Volume 1, pp. 2087–2097.
80. Newman, J. *A Taxonomy of Trustworthiness for Artificial Intelligence*; CLTC: North Charleston, SC, USA, 2023.
81. Rudin, C. Stop Explaining Black Box Machine Learning Models for High Stakes Decisions and Use Interpretable Models Instead. *Nat. Mach. Intell.* **2019**, *1*, 206–215. [CrossRef] [PubMed]
82. OpenAI ChatGPT: Optimizing Language Models for Dialogue. Available online: <https://openai.com/blog/chatgpt/> (accessed on 10 February 2023).

-
83. Chatterjee, J.; Dethlefs, N. This New Conversational AI Model Can Be Your Friend, Philosopher, and Guide. and Even Your Worst Enemy. *Patterns* **2023**, *4*, 1–3. [[CrossRef](#)] [[PubMed](#)]
 84. Wreathall, J. Leading? Lagging? Whatever! *Saf. Sci.* **2009**, *47*, 493–494. [[CrossRef](#)]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.

A scoping literature review of natural language processing application to safety occurrence reports

Ricketts, Jon

2023-04-05

Attribution 4.0 International

Ricketts J, Barry D, Guo W, Pelham J. (2023) A scoping literature review of natural language processing application to safety occurrence reports. *Safety*, Volume 9, Issue 2, April 2023, Article number 22

<https://doi.org/10.3390/safety9020022>

Downloaded from CERES Research Repository, Cranfield University