2026 European Conference on Computing in Construction
Corfu, Greece
July 12-15, 2026

MULTI-CLASS PREDICTION OF OCCUPATIONAL ACCIDENT TYPES USING XGBOOST IN PIPELINE CONSTRUCTION

Ahmet Eren Ünlü[1], Recep Özkan[2], and Onur Behzat Tokdemir[2]
[1]Istanbul Technical University, İstanbul, Türkiye
[2]Turkish-German University, İstanbul, Türkiye
[3]Artificial Intelligence and Data Science Research Center, Istanbul Technical University, İstanbul, Türkiye

Abstract

In this study, the XGBoost algorithm was chosen to predict accident types using data from incidents that occurred in pipeline construction projects, and a multi-class classification approach was presented. A model was trained on 1,184 real accident records obtained from pipeline construction projects. This data includes information such as the time of the accident and the type of treatment required. Classification results showed an overall accuracy of 83%, a macro-average F1 score of 0.82, and a weighted F1 score of 0.81. Some classes showed high recall rates, while recall was low in categories with limited data. The results demonstrate that XGBoost can be a suitable model for predicting accident types using historical incident datasets and can contribute to risk analysis in construction projects through artificial intelligence.

Introduction

Construction sites have the highest number of fatal work injuries of any private industry sector, accounting for approximately one in five occupational deaths (BLS, 2024). Therefore, the construction industry poses a significant risk to both human safety and property (Mostofi & Toğan, 2022). According to the International Labour Organization (ILO, 2023), construction workers in economically developed countries are three to four times more likely to experience fatal workplace accidents compared to those in other sectors. Globally, an estimated 2.93 million workers die annually from occupational accidents and work-related diseases, while an additional 395 million suffer from non-fatal injuries (ILO, 2023). The economic impact of these incidents is substantial, estimated at approximately 4% of global GDP, which is equivalent to nearly $3 trillion in direct and indirect costs (ILO, 2023). While construction work accounts for only about 7% of global employment, it is responsible for a disproportionate 30–40% of fatal occupational injuries (ILO, 2023). Beyond macroeconomic costs, workplace accidents directly impair corporate performance by reducing productivity, operating profit, and sales growth at the firm level (Dong Koo Kim & Park, 2021). In recent years, the growing global demand for transporting energy, water, and chemicals has accelerated the development of pipeline construction projects. However, this expansion has been accompanied by a significant rise in workplace accidents within the sector (Mammadov et al., 2023). Pipeline construction activities are categorized among the highest-risk subsectors, characterized by unique challenges such as exposure to pressurized systems, flammable substances, and geographically remote working environments (Mammadov et al., 2023). Despite these risks, data-driven approaches to accident prediction in pipeline construction remain limited, highlighting the critical need for targeted research in this domain.

Most current studies focus on estimating accident severity or risk levels rather than identifying specific accident types that may occur. However, accident type prediction offers a proactive approach to safety management, as it allows for the prediction of potential events that may occur on-site and the implementation of preventive measures before accidents happen. Therefore, accident type prediction can support more effective safety planning, training, and resource allocation, especially in high-risk environments.

Pipeline construction activities are categorized among the highest-risk subsectors, characterized by unique challenges such as exposure to pressurized systems, flammable substances, and geographically remote working environments (Mammadov et al., 2023). These types of hazards differentiate pipeline construction from other constructions. Despite these distinctive risks, data-driven approaches specifically focused on accident type prediction in pipeline construction remain limited. Therefore, the motivation for this study is to fill this gap by applying a multi-class classification model using XGBoost for accident types occurring in pipeline construction.

Literature Review

In the construction industry, the use of machine learning (ML) and artificial intelligence (AI)-based models as estimation and prediction methods is increasing (Doğan et al., 2022; Koc et al., 2024). Given the persistently high accident rates in the sector, the application of data-driven predictive models has become a practical necessity not only to reduce accident frequency but also to uncover complex causal relationships that traditional statistical methods fail to capture (Alkaissy et al., 2023).

Machine learning can be particularly beneficial because it excels with complex and large datasets (Wang et al., 2022). As Sarkar and Maiti (2020) mention, tree-based ML models are increasingly used to improve accident prediction and support safety planning. XGBoost is a popular ensemble learning algorithm known for its speed, high predictive accuracy, and ability to handle large or complex datasets (Asselman et al., 2023). For example, Pang et al. (2024) developed an early-warning model for traffic accidents using XGBoost and achieved significant prediction accuracy.

Similarly, Mammadov et al. (2023) predicted accident outcomes in pipeline construction using 11 different ML algorithms across 12 sub-datasets created by considering different event types, outcomes, timings, and causes, demonstrating the value of ML in pipeline safety management. Unlike the present study, that work focused on predicting accident outcomes rather than accident types, with deep learning outperforming other algorithms in most configurations. The study proposed a prescriptive prevention model and provides a methodological basis relevant to the current work. Davoudi Kakhki et al. (2019) compared the performance of SVMs, boosted trees, and Naïve Bayes on over 33,000 compensation claims in the agribusiness sector. Their models achieved accuracy rates between 92% and 98% in predicting injury severity, demonstrating ML's effectiveness in modeling occupational injury outcomes. Similarly, Ayhan and Tokdemir (2019) used a hybrid model integrating artificial neural networks and fuzzy logic to predict construction accident outcomes with up to 84% accuracy. Although their focus was on outcome severity rather than accident types, their work bears methodological relevance to the present study. Recent innovations have turned toward graph-based ML models. Mostofi and Togan (2023) proposed a multi-head graph attention network (GAT) for construction safety prediction, achieving up to 87.1% accuracy. Building on earlier work in which a graph convolutional network achieved 94% accuracy in predicting accident severity outcomes (Mostofi et al., 2022), their approach improves the interpretability of accidents by connecting them, thereby providing a more detailed understanding of workplace accident risks.

In another large-scale study, Choi et al. (2020) analysed over 137,000 accident records from Korea, finding that Random Forest outperformed other algorithms, particularly when class imbalance was addressed with random oversampling, an approach conceptually aligned with the SMOTE (Synthetic Minority Oversampling Technique) technique used in the present study. Luo et al. (2023) achieved 82% accuracy in predicting accident severity using a Random Forest model trained with entropy-based text mining and SMOTE, identifying key risk factors including lack of supervision and inadequate training. More recently, Alsulamy et al. (2026) demonstrated that by including the nature of the incident (NOI) as an explanatory variable, the XGBoost model achieved 89% accuracy in predicting accident severity (SOI).

Yoo et al. (2024) developed a multi-class classification model for nine accident types in Korean construction sites using SMOTE and machine learning algorithms, achieving strong predictive performance and demonstrating that accident type classification is a viable and practically relevant prediction target. Alkaissy et al. (2023) further showed that ensemble methods including XGBoost can effectively classify injury types across large-scale construction accident datasets, highlighting the suitability of boosting algorithms for multi-class safety prediction tasks. Beyond conventional ML approaches, Park et al. (2026) compared ML and large language model (LLM) approaches for accident prediction, concluding that ML-based models remain highly effective for structured classification tasks even as newer language-based methods emerge.

While most of the mentioned studies focused on predicting the severity or consequences of accidents, this research aims to predict accident types using the XGBoost algorithm on a dataset of 1,184 pipeline construction accident records. Pipeline construction projects involve distinctive, dynamic, and complex working conditions that differ significantly from other construction types. Activities involving oil, gas, and chemicals introduce critical OHS risks such as fire, explosion, and chemical exposure, while factors such as limited worker experience, extensive geographic scope, and the high number of stakeholders further elevate accident risks (Mammadov et al., 2023). Despite this, pipeline construction remains a less-researched domain in the safety prediction literature, and this study aims to address this gap. Furthermore, this research incorporates a synthetic oversampling technique to mitigate class imbalance and its negative impacts on prediction performance.

Methodology

XGBoost was chosen in this study due to its scalability, high predictive accuracy, and ability to handle large and complex datasets. It employs a novel sparsity-aware algorithm for sparse data and directly handles missing values, making it robust in real-world applications. The algorithm builds an ensemble of decision trees in an additive manner, incorporating L1 and L2 regularization terms to prevent overfitting and enhance generalization performance. Additionally, XGBoost provides built-in feature importance scores, enabling interpretable insights into the key predictors of accident types (Chen & Guestrin, 2016). It builds models sequentially. Each new tree attempts to learn from the previous tree's errors to improve its results. The general form of the objective function in the t-iteration is defined as follows:

$$\sum_{i=1}^{n} l(y_i, \hat{y}_i^{(t-1)} + f_t(x_i)) + \Omega(f_i) \tag{1}$$

where l is the loss function measuring the difference between the predicted and true labels, ft is the function (decision tree) added at iteration t, and Ω(ft) is the regularization term that penalizes model complexity to avoid overfitting. Unlike traditional gradient boosting methods, XGBoost incorporates second-order derivatives of the loss function to approximate gradients more accurately and directly incorporates regularization into the objective, which enhances generalization. Additionally, techniques such as column subsampling, sparsity-aware learning, and optimized handling of missing values contribute to its superior performance and computational efficiency. The workflow of the study is shown in Figure 1.

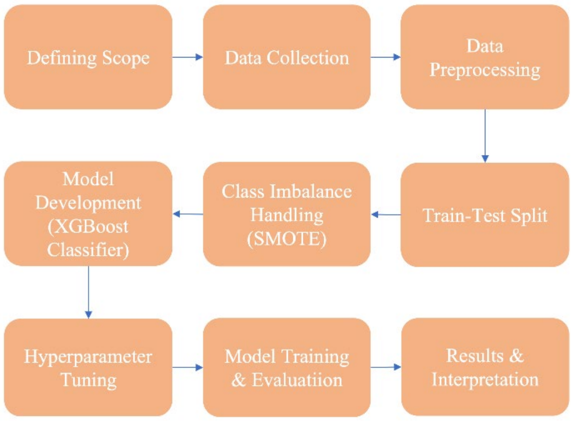

Figure 1: The workflow of the study.

Defining Scope and Data Collection

This study aims to develop a model to predict the types of accidents that may occur during pipeline construction. While results in the literature are generally predicted based on the causes of the accident, this study directly considers the type of accident as the target variable. To create the model, the necessary data were first determined, and then the data collection process was initiated. The data was collected by anonymizing it from a pipeline construction project spanning Asia to Europe, involving numerous countries and stakeholders. The organization of over 5000 unstructured accident reports resulted in 1184 structured accident records. The present study reuses this dataset with a distinct focus, rather than predicting accident outcomes, the current work utilizes root causes as input features to perform multi-class classification for predicting accident types using XGBoost. Each record is labelled with one of six predefined accident types (classes 0-5) and includes various categorical and temporal characteristics such as the accident outcome, primary causes, immediate causes, and time and date of occurrence. The six accident types used as classification targets are derived from event/exposure categories coded according to the OIICS (Occupational Injury and Illness Classification System) (United States. Bureau of Labor Statistics. (2007). The categories reflect the distribution of accident types recorded in pipeline construction projects. The classes and descriptions of the accident types are shown in Table 1.

Table 1: Class Types and Description.

Class	Description
0	Bodily reaction and exertion (Injuries resulting from physical stress, overexertion, heavy lifting, or repetitive motions)
1	Contacts with objects and equipment (Incidents where a worker is struck by, caught in, or crushed by moving machinery, equipment, or falling materials)
2	Exposure to harmful substances or environments (Incidents involving exposure to toxic chemicals)
3	Falls (Slips and trips on the same level, or falls from a higher level)
4	Fires and explosions (Accidents caused by the ignition of fiery gases)
5	Transportation accidents (Accidents involving vehicles in a construction zone)

Data Preprocessing

The data used included time information such as the hour (0–23), month (1–12), and day of the week (0: Monday – 6: Sunday). In addition, the accident's outcome (e.g., injury status), root causes (RC), and immediate causes (ICF) were also included in the model. These causes were coded according to the World Health Organization (WHO, 2001) classification system. All data was converted into numerical forms. Categorical variables and textual expressions were numerically encoded to facilitate processing by the algorithm. Missing data was removed to prevent any adverse impact on the model's results. The detailed descriptions and data types of the input features used in the proposed model are summarized in Table 2.

Table 2: Input Features.

Feature Name	Description	Data Type
Time_block	Time of the accident categorized into 4 blocks (Morning, Afternoon, Evening, Night)	Categorical
Season	Season of the accident (Winter, Spring, Summer, Autumn)	Categorical
Weekday	Day of the week (0: Monday - 6: Sunday)	Categorical
F/Injury /ADPD/NM	The immediate outcome/severity of the accident	Categorical
ICF1, ICF2	Immediate causes of the accident (based on WHO classification)	Categorical
RC1, RC2	Root causes of the accident (based on WHO classification)	Categorical

Handling Class Imbalance with SMOTE

During the analysis, it was observed that the dataset was unbalanced, meaning the available accident data was not evenly distributed, and therefore some types were underrepresented. To prevent data leakage, SMOTE (Synthetic Minority Oversampling Technique) was applied exclusively to the training subset (80%) after the train-test split. Specifically, the minority classes, such as Fires and Explosions (Class 4) and Transportation Accidents (Class 5), were oversampled to match the instance count of the majority class. Since the minority classes contained very few instances, k_neighbors were set to 1 to ensure SMOTE could be applied without error, as a larger neighborhood size would exceed the available sample count in these classes, thereby ensuring the XGBoost model learns unbiased decision boundaries.

Hyperparameter Tuning

To obtain the results, the hyperparameters were first adjusted. One key parameter is the maximum depth, which determines how deep each decision tree can go. Deeper trees can learn more complex patterns, but they can also lead to overfitting. The learning rate determines how much each tree contributes to the overall prediction. A lower learning rate increases learning time but generally yields more consistent and accurate results. The N_estimators determine how many trees the model will build. Generally, using more trees improves performance but also increases training time. Finally, the subsampling parameter determines the proportion of training data used to build each tree. Using less than the full data adds randomness to the training process and reduces overfitting. Another important hyperparameter is `colsample_bytree`, which specifies the proportion of randomly selected features per tree. Using only a portion of the features allows the model to generalize better and reduces the likelihood of overfitting.

Several key XGBoost hyperparameters were selected to improve model performance. The maximum tree depth was adjusted to balance complexity and overfitting. A low learning rate was used to ensure reliable model learning. The number of estimators was appropriately chosen to increase accuracy without significantly increasing training time.

In addition, subsample and colsample_bytree settings were used to introduce randomness and improve generalization.

Model training was initially performed with the default hyperparameters, followed by a limited grid search to optimize them. The best model performance was obtained with the following parameter configuration: maximum tree depth 6, learning rate 0.1, N_estimators 200, subsampling rate 0.8, and column sampling rate per tree 0.8. The hyperparameter configuration established for the XGBoost model to balance complexity and prevent overfitting is presented in Table 3.

Table 3: Hyperparameters.

Hyperparameter	Value	Description
Max_depth	6	Maximum depth of a tree. Balances model complexity and overfitting.
Learning_rate	0.1	Step size shrinkage used in update to prevent overfitting.
N_estimators	200	Number of gradient boosted trees.
Subsample	0.8	Subsample ratio of the training instances.
Colsample_bytree	0.8	Subsample ratio of columns when constructing each tree.
Objective	Multi: softmax	Learning objective for multi-class classification.

Model Training and Evaluation

The final XGBoost model was trained on the preprocessed dataset following SMOTE-based augmentation. The model was trained with an 80%-20% training-to-test split. The model's performance was evaluated using metrics such as accuracy, recall, F1 score, and precision, which indicate how well the model predicted each type of accident. All analyses were conducted in Google Colab using Python. The following libraries were used: pandas and numpy for data processing, scikit-learn for model evaluation and train-test splitting, imbalanced-learn for SMOTE implementation, and xgboost for model training and prediction. Visualizations were generated using matplotlib and seaborn.

Results and Discussion

Precision, recall, F1-score, and overall accuracy metrics were used to obtain and evaluate the results. Precision measures how accurate the model's positive predictions are. Recall indicates how well the model captures all real samples of a class. The F1 score is the harmonic mean of precision and recall. It provides a balanced measure when both false positives and false negatives are significant. It is particularly useful when a single measure is needed that considers both precision and recall.

In accident type prediction using the XGBoost model, True Positive (TP) refers to correctly predicted accident types, False Positive (FP) refers to accident types that did not occur but were incorrectly predicted, and False Negative (FN) refers to actual accident types that the model could not predict; these terms are necessary to calculate recall, precision, and the F1 score to evaluate classification performance. The equations for calculating these metrics are given in Equations 2, 3 and 4.

$$\text{Recall} = \frac{TP}{TP+FN} \tag{2}$$

$$\text{Precision} = \frac{TP}{TP+FP} \tag{3}$$

$$\text{F1-Score} = \frac{2 \text{ x Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}} \tag{4}$$

This study categorized 1184 pipeline construction accidents into 6 main types. We measured the XGBoost Classification algorithm’s performance, which reached 83% overall accuracy. The macro average F1 score was 0.82, and the weighted F1 score was 0.81. The classification performance of the model across all classes is graphically represented in Figure 2.

In addition to the overall evaluation, the results obtained for each accident class were examined separately. The class in which the model was most successful in classification was Class 4, which refers to Fires and Explosions. In this class, the model achieved a precision of 0.96, a recall of 1.00, and an F1 score of 0.98. Fires and explosions are among the high-risk categories in pipeline

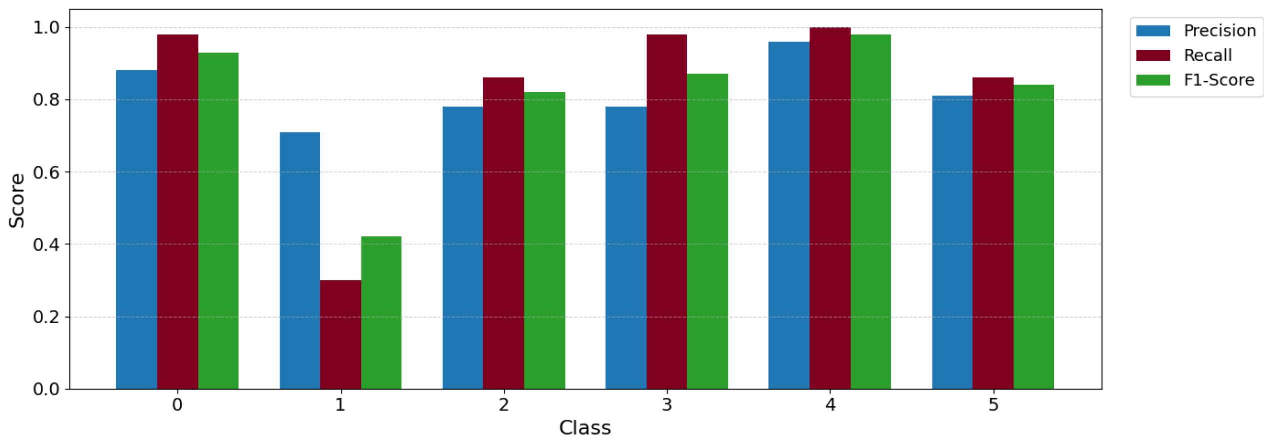

Figure 2: Class-wise classification performance of the XGBoost model

construction, and the model's successful predictive performance in this class is a noteworthy contribution to site safety.

Similarly, the recall value of 0.98 obtained for Class 3 (Falls), one of the most common and dangerous types of accidents in the construction industry, demonstrates that the model has an extremely low accident-risk skipping rate and offers a reliable monitoring mechanism. In this class, the precision was 0.78, and the F1 score was 0.87. Compared to binary classifications, the increase in the number of classes the model can choose in multi-class classifications brings with it the risk of misclassification. Therefore, achieving these metric values among 6 classes is a remarkable achievement.

When examining the results for Class 0 (Bodily reaction and exertion), the precision value was 0.88, while the recall and F1 score metrics were 0.98 and 0.93, respectively. The results show that the model can offer a reliable monitoring mechanism in this classification as well. The recall value of 0.86 and the F1 score of 0.84 obtained for Class 5 (Transportation accidents) show that the model can also distinguish these types of accidents with high accuracy.

Despite these successful results, the model's most significant weakness is its recall of 0.30 and F1 score of 0.42 in the Class 1 (Contact with Objects and Equipment) category. An examination of the predictions revealed that actual Class 1 instances were frequently misclassified as Class 3 (Falls). This misclassification largely stems from the fact that these accident types often share similar underlying causes or occur simultaneously on-site, which blurs the decision boundaries for the algorithm. Furthermore, since the SMOTE technique was primarily targeted at extreme minority categories (such as Classes 4 and 5) to match the majority class, the inherent data limitations of Class 1 remained. Combined with its lack of distinct predictive features, this prevented the model from effectively learning the unique patterns required to separate Class 1 from more dominant, similar classes like Class 3.

These results demonstrate that models like XGBoost can be used as a powerful monitoring tool for pipeline construction safety. Still, their performance may be inadequate for accidents with characteristics similar to other classes, such as Class 1. In future studies, working with more and diverse datasets in pipeline construction and testing the model for construction accidents across different areas will be meaningful for evaluating the model's success.

Compared to traditional safety analysis methods that typically rely on descriptive statistics or severity-based assessments, the proposed approach allows safety managers to understand accident patterns more proactively by identifying the likely type of incident before it occurs, helping project owners complete the project with less risk.

One of the most important strengths of this study is its focus on accident type prediction rather than accident severity prediction. While severity prediction is useful for post-incident assessment, predicting accident types provides more actionable information for safety planning, training programs, and preventive strategies. Specifically, identifying the most likely accident categories allows safety managers to allocate resources more effectively and design targeted interventions for high-risk activities.

Alongside these strengths, the study also has some limitations. The use of SMOTE to address class imbalance adds synthetic samples to the training data. While SMOTE was applied only to the training subset to prevent data leakage, the generated samples may not fully represent real-world accident scenarios and could introduce bias to the model. Although the study includes accident types from different construction sites, it focuses on a defined set of accident categories. While these categories reflect real-world classifications, future research could examine a broader or more detailed set of accident types to enhance the model's comprehensiveness.

Finally, although the model demonstrates strong classification performance, its practical application in real construction environments requires integration with real-time data collection systems and validation using live project data. Future work should focus on testing the model in operational environments and integrating it with digital construction technologies such as BIM and IoT-based monitoring systems.

Overall, this work lays a foundation for the development of proactive, data-driven safety management tools in pipeline construction and highlights the potential of machine learning approaches to improve occupational safety outcomes.

Conclusion

This study investigated the success of the XGBoost model, trained on 1184 real-world accident data from pipeline construction sites, in predicting accident types using a multi-class classification method. Given the potential for both material and non-material damage from accidents in the construction sector, the study focused on predicting these types of accidents. Accident types were numerically translated from text and categorized into six classes. The model was trained on the accident data. Despite the relatively small dataset, the model achieved good prediction performance with 83% accuracy and a macro F1 score of 0.82. These findings validate the potential of machine learning approaches, such as XGBoost, for occupational safety analysis in the construction industry. The proposed model can be particularly beneficial for safety managers, site supervisors, and project owners by enabling automated classification of incident reports and helping uncover complex relationships between root causes and specific accident types, thereby strengthening post-incident analysis and future safety planning. Future studies could validate the model using larger, more diverse datasets in the pipeline field. It is also possible to make predictions on similar datasets using different ML algorithms. Furthermore, applying this model across different sub-sectors of the construction industry could yield interesting results. The current results provide a reliable reference point for ongoing developments in AI-assisted accident prevention.

Acknowledgments

Recep Özkan gratefully acknowledges the support received from the Scientific and Technological Research Council of Turkey (TÜBİTAK) through the 2211-[A] National PhD Scholarship Program.

References

Alkaissy, Maryam & Arashpour, Associate Prof. Mehrdad & Golafshani, Emad & Hosseini, M. Reza & Khanmohammadi, Sadegh & Bai, Yu & Feng, Haibo. (2023). Enhancing construction safety: Machine learning-based classification of injury types. Safety Science. 162. https://doi.org/106102. 10.1016/j.ssci.2023.106102.

Alsulamy, S., Alshayeb, M., Inam, I., & Ahmed, A. (2026). Machine learning applications for predicting safety incidents in construction industry. Scientific Reports. https://doi.org/10.1038/s41598-025-34763-0

Asselman, A., Khaldi, M., & Aammou, S. (2023). Enhancing the prediction of student performance based on the machine learning XGBoost algorithm. Interactive Learning Environments, 31(6), 3360-3379. https://doi.org/10.1080/10494820.2021.1928235.

Ayhan, B. U., & Tokdemir, O. B. (2019). Predicting the outcome of construction incidents. Safety science, 113, 91-104. https://doi.org/10.1016/j.ssci.2018.11.001.

Chen, T., & Guestrin, C. (2016). Xgboost: A scalable tree boosting system. In Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining (pp. 785-794). https://doi.org/10.1145/2939672.2939785.

Choi, J., Gu, B., Chin, S., & Lee, J. S. (2020). Machine learning predictive model based on national data for fatal accidents of construction workers. Automation in Construction, 110, 102974. https://doi.org/10.1016/j.autcon.2019.102974.

Doğan, N. B., Ayhan, B. U., Kazar, G., Saygili, M., Ayözen, Y. E., & Tokdemir, O. B. (2022). Predicting the cost outcome of construction quality problems using case-based reasoning (CBR). Buildings, 12(11), 1946. https://doi.org/10.3390/buildings12111946.

International Labour Organization. (2023). *Nearly 3 million people die of work-related accidents and diseases.* https://www.ilo.org/resource/news/nearly-3-million-people-die-work-related-accidents-and-diseases

Kakhki, F. D., Freeman, S. A., & Mosher, G. A. (2019). Evaluating machine learning performance in predicting injury severity in agribusiness industries. Safety science, 117, 257-262. https://doi.org/10.1016/j.ssci.2019.04.026.

Kim, D. K., & Park, S. (2021). An analysis of the effects of occupational accidents on corporate management performance. Safety science, 138, 105228. https://doi.org/10.1016/j.ssci.2021.105228.

Koc, K., Budayan, C., Ekmekcioğlu, Ö., & Tokdemir, O. B. (2024). Predicting cost impacts of nonconformances in construction projects using interpretable machine learning. Journal of Construction Engineering and Management, 150(1), 04023143. https://doi.org/10.1061/JCEMD4.COENG-13857.

Luo, X., Li, X., Goh, Y. M., Song, X., & Liu, Q. (2023). Application of machine learning technology for occupational accident severity prediction in the case of construction collapse accidents. Safety science, 163, 106138. https://doi.org/10.1016/j.ssci.2023.106138.

Mammadov, A., Kazar, G., Koc, K., & Tokdemir, O. B. (2023). Predicting accident outcomes in cross-border pipeline construction projects using machine learning algorithms. Arabian Journal for Science and Engineering, 48(10), 13771-13789. https://doi.org/10.1007/s13369-023-07964-w

Mostofi, F., & Toğan, V. (2023). Construction safety predictions with multi-head attention graph and sparse accident networks. Automation in Construction, 156, 105102. https://doi.org/10.1016/j.autcon.2023.105102

Mostofi, F., Toğan, V., Ayözen, Y. E., & Tokdemir, O. B. (2022). Construction safety risk model with construction accident network: A graph convolutional network approach. Sustainability, 14(23), 15906. https://doi.org/10.3390/su142315906

Pang, Y., Wang, X., Tang, Y., Guobin, G., Chen, J., Yang, B., ... & Wen, X. (2024, November). A traffic accident early warning model based on XGBOOST. In Fourth International Conference on Advanced Algorithms and Neural Networks (AANN 2024) (Vol. 13416, pp. 494-499). SPIE. https://doi.org/10.1117/12.3049832.

Park, S., Jeong, K., Kim, J., Kim, H., Yoon, S., Kim, J., & Jeon, J. (2026). Comparative evaluation of machine learning and LLM-based approaches for construction accident prediction. Journal of Asian Architecture and Building Engineering, 1-16. https://doi.org/10.1080/13467581.2026.2622236

Sarkar, S., & Maiti, J. (2020). Machine learning in occupational accident analysis: A review using science mapping approach with citation network analysis. Safety science, 131, 104900. https://doi.org/10.1016/j.ssci.2020.104900.

United States. Bureau of Labor Statistics. (2007). *Occupational injury and illness classification manual*. Bureau of Labor Statistics.

U.S. Bureau of Labor Statistics. (2024). *Census of Fatal Occupational Injuries Summary, 2023.* U.S. Department of Labor. https://www.bls.gov/news.release/cfoi.nr0.htm

Wang, C. C., Kuo, P. H., & Chen, G. Y. (2022). Machine learning prediction of turning precision using optimized XGBoost model. Applied Sciences, 12(15), 7739. https://doi.org/10.3390/app12157739.

World Health Organization. (2001). *International Classification of Functioning, Disability and Health (ICF)*. World Health Organization. https://iris.who.int/handle/10665/42407

Yoo, J. W., Park, J., & Park, H. (2024). Enhancing safety of construction workers in Korea: An integrated text mining and machine learning framework for predicting accident types. International journal of injury control and safety promotion, 31(2), 203-215. https://doi.org/10.1080/17457300.2023.2300424