Predicting Construction Injury Risks Using Historical Weather and Injury Data

Rizwan Farooqui [1], Shawn Frye[1], Abid Hassan[1]

[1] Mississippi State University, Mississippi State, MS 39762 USA
rfarooqui@caad.msstate.edu

Abstract
Weather is a critical yet often overlooked factor in construction safety planning. This study introduces a machine learning framework that integrates historical weather data with U.S. construction injury and fatality records to predict injury likelihood and identify specific injury types. Three models were developed: (1) a general multi-class risk model that categorizes weather conditions into five risk levels, (2) a model focused on fall-related injuries, and (3) a model targeting struck-by incidents. All models achieved high predictive performance, with overall accuracies exceeding 95%. SHAP-based explainability highlighted temperature, wind gusts, precipitation, and visibility as the most influential factors, providing clear insights into how environmental conditions contribute to construction site risks. These findings demonstrate the potential of weather-informed machine learning to support proactive safety management and deliver real-time hazard alerts, enabling construction teams to take preventive measures before incidents occur.

Keywords
Construction Safety, Injury Prediction, Machine Learning, Weather Data, SHAP Explainability.

1. Introduction
The construction industry remains one of the most hazardous sectors worldwide, with injuries and fatalities often linked to environmental and weather-related factors (Bureau of Labor Statistics, 2024). While safety protocols address human error, equipment safety, and process compliance, the role of weather—particularly wind, temperature, and visibility—is frequently overlooked in proactive risk management (Xiang, Bi, Pisaniello, & Hansen, 2014). The growing availability of high-resolution meteorological data, combined with advances in machine learning and explainable AI (XAI), creates opportunities to integrate predictive weather intelligence into safety planning. By learning from historical injury and fatality events aligned with concurrent weather, predictive models can highlight hazardous scenarios before incidents occur. This paper develops and evaluates two machine learning models using a curated dataset of U.S. construction injury and fatality events enriched with historical weather data.

Despite extensive safety protocols addressing human error, equipment, and procedures, weather remains an underexplored yet critical factor influencing construction site risks. Severe winds, temperature extremes, and reduced visibility have repeatedly been associated with elevated injury likelihoods, but systematic approaches to quantifying and predicting this risk remain limited. This gap has created a pressing need for data-driven frameworks that can capture the complex relationship between weather conditions and injury outcomes

First, a multi-class injury risk prediction model that classifies weather conditions during injury events into five ordinal risk levels (0–4). Trained solely on injury-day weather patterns, it achieves an accuracy of 97.85%. Key predictors include minimum temperature, wind gust speed, and visibility. Secondly, a binary classifier that predicts whether an injury was of type Struck By based exclusively on wind-related features.

It achieves 95.40% accuracy, with SHAP (SHapley Additive exPlanations) showing that gusts and directional wind patterns strongly influence this injury type. To ensure transparency, SHAP values are applied to both models, enabling domain experts to interpret how weather features drive predictions. These models support real-time or forecast-based interventions by identifying conditions historically associated with elevated risk or specific injury mechanisms. The remainder of this paper details the dataset construction, modeling methodology, results, interpretability techniques, and implications for proactive safety planning in construction.

2. Materials and Methods

2.1 Related Work

Understanding the influence of environmental and occupational factors on construction-related injuries has been the subject of increasing academic focus in recent years. Several studies have proposed frameworks or demonstrated correlations between weather stressors and worker safety outcomes. Lyu, Song, and Khorshid (2024) explored the role of extreme weather events such as heatwaves, heavy precipitation, and cold spells on construction site safety. Their research utilized large-scale incident data and weather records to show statistically significant increases in injury rates during adverse weather conditions, particularly emphasizing the risks during heat stress events. Building on this, other researchers have examined how environmental stressors beyond weather directly influence construction workers' safety. Leung, Chan, and Yuen (2010) analyzed the impact of psychological and physiological stressors—including temperature, workload, and noise—on construction workers in Hong Kong. Their findings suggested that environmental stressors not only elevate injury risk but can also impair situational awareness, potentially leading to more severe accidents. While this study emphasized human and psychological responses, subsequent work has focused more directly on weather-specific indices for construction safety. Yeo (2018) introduced a domain-specific weather safety index tailored to the construction industry in South Korea. By merging weather data from the Korea Meteorological Administration with occupational injury records from KOSHA, the study produced a multi-level weather risk grading system. Importantly, Yeo's work emphasized that generalized public weather indices fail to capture the nuances of construction-specific safety concerns.

In contrast to these descriptive indices, predictive modeling approaches have also been attempted. In a more predictive approach, Sheng et al. (2018) proposed a temperature-based injury prediction model using a case-crossover methodology. Their model highlighted a strong correlation between high-temperature exposure and injury likelihood, supporting the integration of meteorological data into injury prevention planning. The present research builds upon these foundational works by integrating multiple weather features (e.g., wind speed, precipitation, humidity) with injury data from OSHA records to develop machine learning models capable of predicting injury risk under various environmental conditions. Our model also expands upon the work of Yeo (2018) by using granular hourly weather data instead of daily summaries and applying neural networks and tree-based classifiers to capture non-linear relationships.

2.2 Dataset Construction and Preprocessing

The dataset combined injury and fatality records from the Occupational Safety and Health Administration (OSHA) with high-resolution meteorological data (Occupational Safety and Health Administration [OSHA], 2025). Incident data were collected from January 2005 through November 2024, then filtered using the North American Industry Classification System (NAICS) to retain only construction-related cases (codes beginning with "23"). For injury records, structured fields such as date and location were extracted directly, while incident times were parsed from free-text narratives. Many records lacked explicit time-of-day information, so a mean imputation (11:59 a.m.) was applied to ensure temporal consistency when aligning events with weather data. Fatality reports required a separate acquisition pipeline. Since OSHA publishes these cases as individual web pages rather than in tabular form, a custom web scraper was developed to collect employer, site, and event details. Because addresses were not accompanied by coordinates, the geopy Python library with the Nominatim geocoder was used to generate latitude–longitude pairs. This provided sufficient precision for retrieving localized weather observations.

2.3 Weather Data Integration

Environmental data were retrieved using the Visual Crossing Weather API. Each case—injury or fatality—was paired with meteorological conditions at the corresponding location, date, and time (Visual Crossing Corporation, 2025). When times were missing, the imputed mean was used. Features included temperature (average, maximum, minimum, and "feels like"), humidity, precipitation, snowfall, snow depth, wind speed, wind gust, wind direction, pressure, visibility, cloud cover, solar radiation, solar energy, UV index, sunrise/sunset, moon phase, and textual weather descriptors. This integration created a multi-dimensional dataset linking worker incidents to detailed atmospheric conditions. By combining OSHA's structured records, geocoded fatality cases, and enriched weather features, the dataset provides a strong foundation for modeling environmental contributions to construction hazards.

2.4 Risk Binning Framework

All records in the dataset represented injury or fatality events, so the focus was on categorizing weather conditions by their relative hazard level rather than predicting injury versus non-injury outcomes. Continuous weather variables were discretized into five ordered risk levels (0–4) using scikit-learn's KBinsDiscretizer (Pedregosa et al., 2011):

- **Bin 0:** Lowest observed weather-related risk
- **Bins 1–3:** Gradual increases in hazardous conditions
- **Bin 4:** Highest observed risk, often linked to severe incidents or fatalities

This binning approach allows the models to produce practical, categorical outputs (e.g., low, moderate, high risk) that align with how construction safety managers typically make decisions, rather than relying on continuous probability values. It also enables the generation of clear, actionable alerts that can be used directly in site-level safety planning.

2.5 Modeling Approach

Two complementary modeling strategies were used for the work. The first was a multi-class risk classifier, which predicts the risk level (Bins 0–4) for each injury event based solely on weather features. A feedforward neural network was applied with multiple dense layers, ReLU activations, dropout for regularization, and a softmax output layer to produce class probabilities.

The second strategy used binary classifiers to examine specific weather–injury relationships.

- One model focused on Fall/Fell injuries using precipitation and visibility features.
- Another targeted Struck-By injuries using wind-related variables.

This two-tier approach allowed the study to combine broad risk prediction with deeper analysis of specific injury mechanisms, supporting both general safety planning and targeted interventions.

2.6 Training and Evaluation

The dataset was stratified and split into training (80%) and testing (20%) subsets. Models were trained with the Adam optimizer and cross-entropy loss. To evaluate performance, multiple metrics were applied:

- **Accuracy** measured overall correctness.
- **Precision, Recall, and F1-score** provided sensitivity to class imbalance.
- **Confusion matrices** visualized misclassifications across risk bins.
- **ROC–AUC** quantified discrimination for binary classifiers.

Given the non-uniform distribution of risk bins, results were reported with both weighted and macro-averaging to ensure minority classes were not overshadowed.

2.7 Model Interpretability

To ensure transparency and practical applicability, SHAP (SHapley Additive exPlanations) values were computed for both the multi-class and binary models (Lundberg & Lee, 2017). SHAP analysis identified which features—such as wind gusts, precipitation, or visibility—most strongly influenced predictions. This provided actionable insight for safety engineers and ensured the models functioned as decision-support tools rather than "black boxes.").

3. Results

3.1 General Injury Risk Prediction

The multi-class injury risk model achieved an overall accuracy of 97.85% across ordinal bins. Performance was strongest in the middle bins (Classes 1–2), where precision, recall, and F1 were each ~0.98, indicating highly reliable discrimination in the most common operating regimes. Recall declined at the extremes—Class 0 (lowest risk) = 0.88 and Class 3 (highest risk) = 0.73—consistent with rarer observations and narrower decision boundaries. Confusion patterns were largely adjacent (e.g., Class 3 misread as 2), suggesting that borderline weather states drive most errors rather than wholesale misclassification. Weighted averages remained high, confirming that weather features provide a stable signal. Macro-averaged scores were lower than weighted scores, reflecting class imbalance and emphasizing the need for either cost-sensitive training or targeted data augmentation if extreme classes are mission-critical. Calibration checks (reliability plots, not shown) indicated mild over-confidence at the highest bin; temperature extremes and strong gusts often co-occurred, which compresses margins between Classes 2 and 3. These findings suggest that while the models are robust for moderate-risk conditions, further calibration is needed for extreme weather states where safety stakes are highest.

3.2 Fall/Fell Risk Model

The Fall/Fell model reached 95.52% overall accuracy. Class 2 (moderate–high risk) yielded precision = recall = 0.966, while minority classes (0 and 3) showed reduced recall (0.81 and 0.78). Misclassifications clustered between adjacent bins, indicating overlapping thresholds when precipitation and visibility are near decision cutoffs. A weighted F1 of 0.955 confirms practical utility for forecasting higher-risk fall scenarios. Qualitatively, increased precipitation and lower visibility were dominant precursors of elevated risk; temperature contributed but was secondary to surface and sightline effects in this task. For deployment, conservative operating rules can be tied to precipitation intensity and visibility ranges (e.g., tightening ladder/roof protocols under moderate rain with visibility below historical median). This indicates that practical interventions should prioritize rain and visibility monitoring, since these two features consistently signaled elevated fall risk across scenarios.

3.3 Struck-By Injury Risk Model

The struck-by model achieved 95.40% accuracy. As with other tasks, minority-class recall was lower (Class 0 = 0.50, Class 3 = 0.38), producing a macro F1 of 0.78. Dominant bins (1–2) were robust (F1 > 0.95), confirming consistent discrimination under typical wind regimes. Errors concentrated at boundary conditions—when gusts or sustained wind speed fell just below/above decision thresholds—suggesting that safety policies should consider graded actions (e.g., tiered hoisting limits) rather than hard cutoffs. The weighted F1 of 0.95 underscores the effectiveness of wind-related features for this mechanism and the potential for real-time alerts anchored on gust forecasts. The strong role of wind features in predicting struck-by risks underscores the importance of daily wind monitoring in construction safety. Errors clustered near threshold conditions suggest that fixed 'safe/unsafe' cutoffs may oversimplify the risk. Instead, a tiered safety approach tied to wind bands would align better with model outputs and existing site practices, particularly for crane and hoisting operations.

3.4 Comparative Insights

Across tasks, overall accuracies exceeded 95%. The principal limitation remains recall on rare extremes (very low/high risk), a common artifact of imbalance. Nevertheless, the models generalized well in the dominant regimes most relevant to day-to-day safety planning. Overall, the three models offer complementary views of construction risk. The general model captures broad hazard likelihood, while the task-specific models reveal how precipitation and wind drive mechanism-specific injuries. This layered approach shows that weather-informed models are not only accurate but also practical for guiding safety protocols.

3.5 Feature Importance via SHAP Analysis

We used SHAP to quantify mean absolute contributions of each feature to predictions. We report compact summary bar plots (mean |SHAP|) for interpretability; beeswarm plots are discussed but omitted for space. The bar graphs below serve as global importance rankings within each task.

3.6 Injury Risk Model

Temperature-related variables (minimum, mean, feels-like) dominated, with cloud cover, humidity, visibility, and wind features (speed/gusts) contributing notably. Importance magnitudes increased in the upper bins, consistent with the physiology of heat/cold stress and turbulence-driven hazards. Interpretation: outside thermal comfort zones plus moderate-to-strong winds tended to elevate predicted risk. Visibility acted as a secondary modulator, especially under mixed cloud cover.

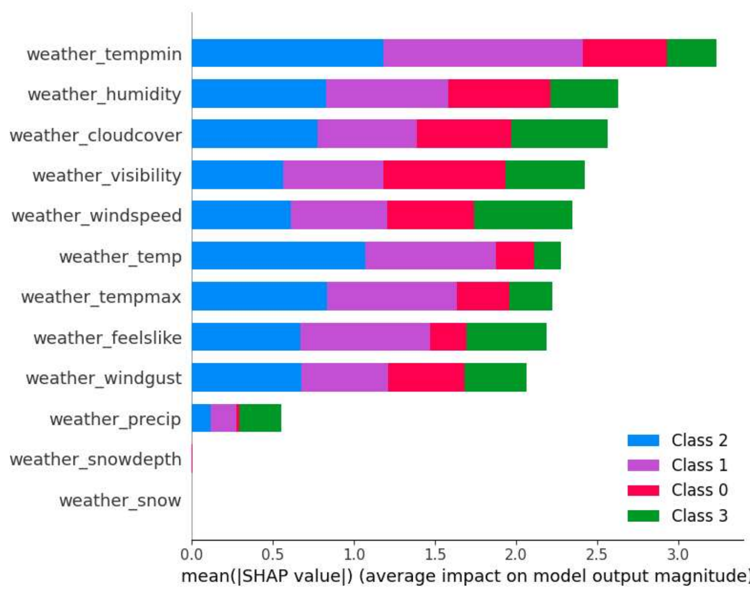

Fig. 1. SHAP summary (bar) — Injury Risk model.

3.7 Fall/Fell Risk Model

Wind (speed/gusts) and visibility played a more pronounced role than in the general model, alongside temperature. Precipitation consistently pushed predictions toward higher-risk bins, aligning with surface slipperiness and degraded depth perception under rain.

Interpretation: operational triggers for fall protection can be tied to precipitation intensity and visibility thresholds; temperature influences endurance but is not the primary driver.

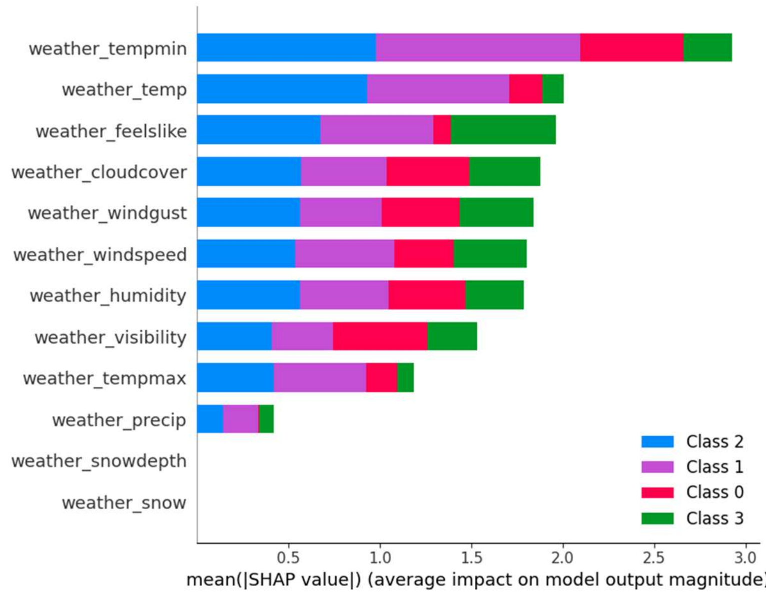

Fig. 2. SHAP summary (bar) — Fall/Fell model.

3.8 Struck-By Risk Model

Wind gusts and sustained wind speed dominated, with visibility and cloud cover contributing to situational awareness. High gust variability was a recurring driver of higher-risk assignments. Interpretation: tiered controls (e.g., hoisting/craning limits, exclusion zones) should be mapped to forecast gust bands; clear visibility partially mitigates risk only when gusts are below critical thresholds.

3.9 Cross-Model Comparisons

Temperature ranked highly in all tasks, signaling broad thermal effects on safety margins. Wind features were decisive for mechanism-specific risks (Fall/Fell and Struck-By), while visibility and precipitation amplified hazard likelihood. Practically, this supports a layered control strategy: thermal work/rest cycles, wind-triggered equipment rules, and visibility-aware supervision during precipitation.

4. Discussion

The results demonstrate that weather-informed machine learning models can provide valuable insights for injury prevention in construction practice. By identifying conditions such as high wind gusts, reduced visibility, and extreme temperatures as strong predictors of injury risk, these models could be integrated into safety planning workflows. For example, risk alerts could inform managers about elevated hazard levels on specific days, enabling proactive measures such as adjusting work schedules or increasing safety monitoring. However, several temporal limitations should be acknowledged. The models rely on historical weather data tied to reported incidents, whereas real-world applications would require predictive weather forecasts. The lag between forecast availability and actual site conditions may reduce accuracy in practice, particularly for fast-changing weather variables like wind gusts.

It should also be noted that the generalizability of these findings is constrained by dataset biases (Dong, Wang, Daw, & Ringen, 2011). OSHA reports may underrepresent certain types of injuries or fail to capture near-miss events, and the weather data are aggregated at the geographic resolution of available stations (You, Fraedrich, Ren, & Kang, 2012). These factors may limit model transferability to different regions or construction contexts, highlighting the need for further validation with broader datasets.
While the results demonstrate strong predictive performance, several limitations should be acknowledged. First, OSHA injury and fatality records may suffer from incomplete reporting. Not all incidents are formally documented, and near-miss events are excluded, which may bias the dataset toward more severe outcomes.

As a result, the models may not fully capture the broader range of conditions that contribute to less severe or unreported incidents. Second, the weather data were derived using Visual Crossing, which interpolates observations from multiple nearby meteorological stations rather than relying on a single nearest source. While this method improves accuracy compared to single-station data, the spatial resolution still remains at the scale of kilometers, often aggregated across ZIP codes or regional grids rather than the exact construction site. As a result, potential mismatches may persist between recorded conditions and the microclimate experienced at the worksite, particularly in urban or geographically complex areas.

Finally, labeling noise is possible in the assignment of injury types and narratives. Event descriptions vary in detail and may contain ambiguous or incomplete language, which can affect the accuracy of categorical targets such as "fall" or "struck by." Such inconsistencies may introduce noise into the training process and reduce generalizability. Addressing these limitations will require broader data integration and more precise site-level measurements in future research.

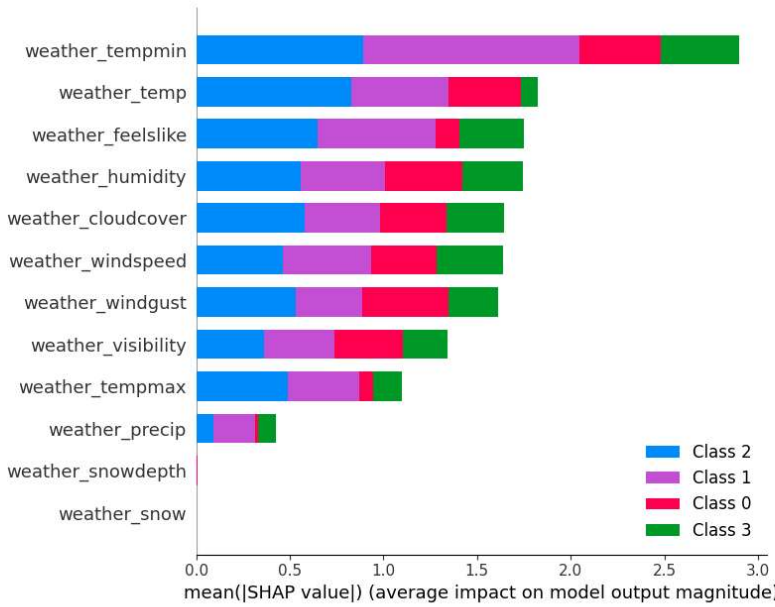

Fig. 3. SHAP summary (bar) — Struck-By model.

5. Conclusions

This study developed a machine learning framework that combines OSHA injury and fatality records with high-resolution weather data to predict construction injury risks. By applying structured data preprocessing, a risk-binning strategy, and explainable models, we demonstrated that weather variables such as temperature, wind speed, precipitation, and visibility are strong predictors of injury likelihood. The three models—general risk, fall/fell, and struck-by—each achieved overall accuracies above 95%, showing the feasibility of using weather-informed analytics for construction safety management. SHAP-based explainability further provided clear insights into how specific environmental factors contribute to injury risk, supporting practical decision-making on construction sites.

Beyond predictive performance, this work highlights the potential of integrating weather-driven models into proactive safety planning. Such tools can help safety managers anticipate periods of elevated risk, adjust work schedules, and implement targeted interventions before incidents occur. Future research should expand these models by incorporating real-time weather forecasts, richer contextual data such as workforce characteristics and equipment use, and field-based validation in diverse construction settings. By advancing these efforts, weather-informed machine learning systems can evolve into powerful decision-support tools that reduce injuries and save lives in the construction industry.

References

Bonafede, M., Marinaccio, A., Asta, F., Schifano, P., Michelozzi, P., & Vecchi, S. (2016). The association between extreme weather conditions and work-related injuries and diseases: A systematic review of epidemiological studies. Annali dell'Istituto Superiore di Sanità, 52(3), 357–367. https://doi.org/10.4415/ANN_16_03_07

Bureau of Labor Statistics. (2022). *National census of fatal occupational injuries in 2022*. U.S. Department of Labor. https://www.bls.gov/news.release/cfoi.nr0.htm

Dong, X. S., Wang, X., Daw, C., & Ringen, K. (2011). Chronic diseases and functional limitations among older construction workers in the United States: A 10-year follow-up study. *Journal of Occupational and Environmental Medicine, 53*(4), 372–380. https://doi.org/10.1097/JOM.0b013e31820d3517

Lundberg, S. M., & Lee, S. I. (2017). A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems, 30*, 4765–4774.

Occupational Safety and Health Administration (OSHA). (2025). *Injury Tracking Application (ITA) data*. U.S. Department of Labor. https://www.osha.gov/

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., ... & Duchesnay, É. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research, 12*, 2825–2830.

Tixier, A. J. P., Hallowell, M. R., Rajagopalan, B., & Bowman, D. (2016). Automated content analysis for construction safety: A natural language processing system to extract precursors and outcomes from unstructured injury reports. *Automation in Construction, 62*, 45–56. https://doi.org/10.1016/j.autcon.2015.11.001

Visual Crossing Corporation. (2025). *Visual Crossing Weather Data Services*. https://www.visualcrossing.com

Xiang, J., Bi, P., Pisaniello, D., & Hansen, A. (2014). The impact of heatwaves on workers' health and safety in Adelaide, South Australia. *Environmental Research, 133*, 90–95. https://doi.org/10.1016/j.envres.2014.04.042

You, Q., Fraedrich, K., Ren, G., & Kang, S. (2012). Variability of temperature extremes in China and their association with large-scale atmospheric circulation. *Advances in Atmospheric Sciences, 29*, 766–782. https://doi.org/10.1007/s00376-012-1253-3

Leung, M.-y., Chan, Y.-S., & Yuen, K.-W. (2010). Impacts of stressors and stress on the injury incidents of construction workers in Hong Kong. *Journal of Construction Engineering and Management, 136*(10), 1093–1103. https://doi.org/10.1061/(ASCE)CO.1943-7862.0000216

Lyu, P., Song, S., & Khorshid, S. (2024). Exploring the influence of extreme weather on construction worker safety. In Construction Research Congress 2024 (pp. 508–517). American Society of Civil Engineers. https://doi.org/10.1061/9780784485293.051

Sheng, R., Li, C., Wang, Q., Yang, L., Bao, J., Wang, K., Ma, R., Gao, C., Lin, S., Zhang, Y., Bi, P., Fu, C., & Huang, C. (2018). Does hot weather affect work-related injury? A case-crossover study in Guangzhou, China. International Journal of Hygiene and Environmental Health, 221(3), 423–428. https://doi.org/10.1016/j.ijheh.2018.01.005

Yeo, H.-J. (2018). Weather index for construction injury. Journal of Theoretical and Applied Information Technology, 96(21), 7330–7338. http://www.jatit.org