Empirical Fragility-Based Adjustment of Embodied Carbon Using Machine-Learning-Derived Post-Hurricane Fragility Data

Mana Nemati Aghdam[1], and *Manish* Dixit[1*]

[1]Texas A&M University, College Station, TX 77843, USA

Abstract. Climate mitigation and hazard resilience are often treated as separate goals, even though storm-driven repairs can add substantial embodied greenhouse gas emissions over a building's lifetime. This paper presents a data-driven workflow that links post-disaster field observations (StEER), site-specific wind fields (ARA), and machine-learning fragility models to support hazard-adjusted embodied carbon assessment. A pooled dataset of 979 residential buildings inspected after Hurricanes Laura (2020) and Michael (2018) was harmonized to a common schema and assigned peak 3-second gust wind speed using nearest-neighbor geospatial matching. Gradient boosting classifiers estimate the probability of roof and wall replacement from gust intensity and building attributes. The models achieve AUC = 0.75 (roof) and AUC = 0.92 (wall). Gust speed is the main driver, with roof age and wall cladding group acting as key modifiers. These fragility outputs can be combined with regional hazard frequency and component-specific repair carbon factors to estimate expected repair-related emissions over a service life. The framework treats resilience as an emissions-reduction strategy by translating replacement risk into lifecycle CO_2e consequences.

- **Keywords.** Climate change adaptation, Hazard-adjusted life cycle assessment, Climate resilience, Machine learning fragility model, Embodied greenhouse gas emissions

1 Introduction

Climate policy usually separates two ideas that are linked in real life. One part focuses on lowering carbon through cleaner materials and efficient systems [1]. The other part focuses on protecting buildings from hazards through stronger design and construction [2]. In practice, these two goals cannot be kept apart [2]. When a building is damaged during a storm, the repairs and material replacements release new embodied carbon [3]. If this cycle repeats after every major event, the building carries a growing carbon load over its lifetime, even if it was designed as a "low-carbon" structure at the start.

This gap creates a major problem. Many sustainability standards measure carbon at construction, but they do not account for repeated repair cycles caused by wind, rain, or flood hazards. As a result, buildings that look sustainable on paper can produce high emissions over time simply because they fail more often.

This study presents a method that links observed post-hurricane building performance with lifecycle carbon accounting. Using StEER reconnaissance data and ARA wind fields, we learn component replacement probability and embed it in a hazard-adjusted lifecycle CO_2e model. In this paper, "carbon" refers to cradle-to-gate embodied GHG emissions in CO_2e units.

2 Literature Review

2.1 Sustainability Research: Carbon Without Hazard Effects

Most sustainability studies focus on lowering embodied and operational carbon at the time of construction [4]. LCA tools such as Tally, GaBi, and OneClick calculate material and energy impacts, but they usually assume the building stays intact for its entire service life [5]. They do not include unplanned repairs caused by storms, windborne debris, or flooding [4]. Because of this, a design that appears low-carbon at the start may produce high emissions over decades if it fails often. As a result, embodied carbon is treated as a static quantity, implicitly assuming uninterrupted component performance over the full service life [6].

2.2 Resilience Research: Failure Without Carbon Accounting

Work in structural resilience has developed detailed models such as fragility curves, engineering damage functions, and hazard mapping [7]. Fragility models

* Corresponding author: mana.ne@tamu.edu

express damage likelihood probabilistically as a function of hazard intensity and structural attributes, making them suitable for integration with lifecycle-based metrics [8]. These tools estimate how likely different components are to fail under wind or surge [7]. They are commonly used to predict economic loss, downtime, and safety outcomes [9]. However, they rarely convert failure probability into carbon impact. Most resilience studies stop at describing physical damage and do not link these outcomes to lifecycle emissions [10].

2.3 Gap in the Literature and Need for Integration

A small body of work links resilience and sustainability by adding repair scenarios to LCA or adjusting maintenance cycles under assumed hazards [11]. These efforts show that repeated replacement of roofs, cladding, and other components can create a large "carbon penalty," especially in regions with frequent storms [11].However, many studies rely on hypothetical damage scenarios or deterministic repair schedules. This study fills that gap by learning replacement probability directly from observed post-event building performance, then structuring the output for hazard-adjusted embodied carbon assessment.

3 Methodology

This methodology is organized into three stages: (1) construction of a pooled empirical dataset and geospatial hazard assignment, (2) supervised learning of component-level failure probabilities, and (3) calculation of expected lifecycle GHG emissions (CO_2e) conditioned on hazard frequency.

3.1 Empirical Dataset and Geospatial Hazard Assignment

The empirical basis of the analysis is a pooled dataset of residential buildings inspected after Hurricane Laura (2020) and Hurricane Michael (2018) through the Structural Extreme Events Reconnaissance (StEER) network [12,13]. Records from both campaigns were harmonized to a common schema. After excluding entries without assigned gust speeds or construction year, the final modeling dataset contains 979 buildings.

Component-level damage targets were defined from StEER damage flags. For roofs, either a generic "DS_roof_any" flag or a logical maximum of "roof cover replacement" and "roof structure replacement" was used to create a binary target for "any roof replacement." For walls, a similar maximum across cladding, substrate, and structural replacement flags was constructed. This yields 320 roof failures and 60 wall failures in the pooled sample. A foundation damage indicator was computed using scour depth and eroded footprint, but only two buildings satisfy this criterion; therefore, foundation performance is not modeled statistically and is used only to interpret the role of foundation type as a predictor in roof and wall models.

To move beyond county-level wind categories, each building was assigned site-specific wind intensities from the Applied Research Associates (ARA) wind fields for Hurricanes Laura and Michael by nearest-neighbor matching (k = 1). Matching distances reflect wind-field grid resolution differences (mean: 0.70 km for Laura; 19.95 km for Michael).

For each building, the spatial join assigns a peak 3-second gust speed V_{gust} (Vg_mph_assigned) and a 1-minute sustained speed V_{sus} (Vsus_mph_assigned). In the predictive models, the primary hazard predictor is the V_{gust}, while the sustained speed is retained in the raw dataset and used for optional derived indices and quality checks.

Structural and architectural attributes are taken from StEER forms. To reduce sparsity while preserving meaningful performance classes, detailed material labels are grouped into broader categories. Roof cover types are mapped to Shingle, Metal, Tile, Built-up, or Other; wall cladding to Vinyl, Brick/Masonry, Fiber Cement, Wood, or Other; and foundations to Slab, Pier/Pile, Crawlspace, or Other. Additional numeric predictors include first-floor elevation, overhang length, number of stories, and a wind damage rating score, which summarizes observed wind-related damage at the site.

The final predictor set used in the machine learning models therefore consists of:

- **Numerical:** assigned peak gust wind speed, roof age, first–floor elevation, roof overhang length, number of stories, and the observed wind damage rating when available.
- **Categorical:** roof shape, roof structural system, grouped roof material type, grouped wall cladding type, and foundation type.

All categorical variables are subsequently transformed using one-hot encoding.

3.2 Machine Learning Pipeline

Component-level failure probabilities were estimated using a Gradient Boosting Classifier (GBM) implemented in scikit-learn. GBM constructs an additive ensemble of decision trees, where each tree is trained to reduce the residual error of the current ensemble. This architecture is well-suited to tabular engineering data, as it can capture non-linear interactions between wind intensity, geometry, and material properties without assuming a parametric functional form.

A common preprocessing pipeline was applied before fitting the GBM models for roof and wall targets. Numerical predictors were imputed using the median of each feature to handle missing entries, which are typical in post-disaster reconnaissance datasets. Categorical attributes were imputed using the most frequent category and then converted into a one-hot representation, yielding a fully numeric design matrix compatible with gradient boosting.

The GBM models were trained separately for the Roof and Wall targets using the combined numeric and one-hot encoded categorical features. For each component, the fitted classifier outputs a predicted probability of failure

$Pf(X)$ for each building, conditioned on its hazard intensity and structural attributes. For each component (roof and wall), the dataset was split into training (75%) and testing (25%) sets using a stratified split, and their discrimination ability was assessed using receiver operating characteristic (ROC) curves and the area under the curve (AUC). The ROC curves for the roof and wall models are reported in Figure 1, showing the trade-off between true positive and false positive rates.

Model interpretability is needed for engineering use. We quantify driver importance with aggregated feature importance and validate directionality with SHAP summary patterns. This highlights how wind intensity, age, material group, and other attributes influence replacement probability.

Partial dependence plots (PDPs) were generated to examine how failure probability changes with key predictors while averaging over others. This checks for physically consistent trends.

The empirical damage indicators in the StEER dataset are binary observations describing whether replacement occurred in a specific event. While these observations capture past outcomes, they do not directly provide a probabilistic measure suitable for estimating future replacement frequency. The machine learning models generalize these binary observations into a conditional probability of replacement, $Pf(X)$, which can be evaluated for arbitrary combinations of hazard intensity and building attributes and subsequently embedded in lifecycle emission calculations.

3.3 Hazard-Adjusted Lifecycle Greenhouse Gas (CO_2e) Model

The third stage embeds the learned fragility into a probabilistic lifecycle greenhouse gas model. All emissions are expressed as carbon dioxide equivalents (CO_2e); for brevity, we refer to them as "carbon" throughout the paper. The goal is to estimate how much embodied GHG emissions are expected to be released over a building's service life due to hazard-driven repairs, in addition to the initial construction.

Let GWP_{initial} denote the embodied GHG emissions associated with initial construction. For a given wind hazard intensity characterized by a mean recurrence interval MRI (years), the annual occurrence rate is:

$$\lambda = 1/MRI \quad (1)$$

For a building with service life T, and for each component i, the machine-learning model provides a conditional failure probability $P_f(X_i)$. The embodied emissions associated with repairing or replacing that component are denoted $GWP_{\text{repair},i}$ (kg CO_2e).

The expected lifecycle emissions are expressed as:

$$E[CO_2e] = GWP_{initial} + \sum_i (\lambda \cdot T \cdot P_f(X_i) \cdot GWP_{repair\ i}) \quad (2)$$

which generalizes conventional deterministic LCA by explicitly incorporating the stochastic nature of failure. This structure allows the method to be applied under any hazard environment and for any selection of materials once $GWP_{\text{repair},i}$ values become available. Because the present study focuses on the methodological development, no numerical LCA results are reported here.

4 Results

4.1 Model Performance and Predictive Capability

The predictive capability of the component-level fragility models was evaluated using ROC curves and AUC. Figure 1 presents ROC curves for roof and wall replacement prediction.

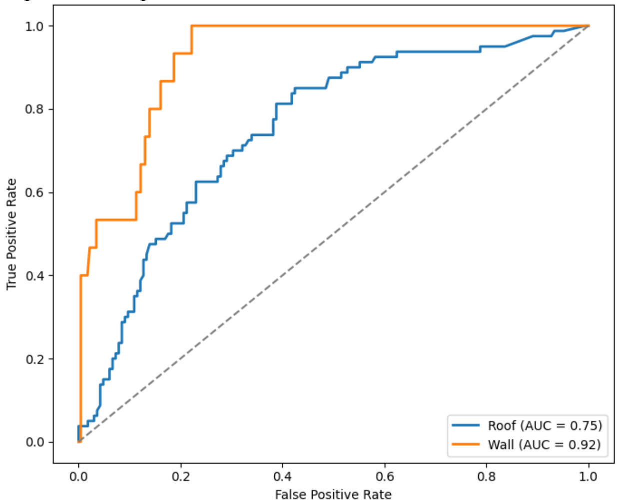

Fig. 1. ROC curves for component damage models.

The roof model achieves an AUC of approximately 0.75, and the wall model achieves an AUC of approximately 0.92. These results indicate that the models capture meaningful relationships between hazard intensity, exposure, and envelope system characteristics.

These results indicate that the gradient boosting framework effectively captures the dominant relationships between wind hazard intensity and building attributes governing envelope performance. While the high AUC values suggest robust predictive capability, the models are not intended for point prediction at the individual-building scale. Rather, they are used to estimate probabilistic failure tendencies across the building stock, which form the basis for the hazard-adjusted lifecycle greenhouse gas assessment.

4.2 Dominant Drivers of Roof and Wall Failure

To understand how structural and hazard-related variables govern component vulnerability, feature importance metrics were extracted from the trained gradient boosting models. Importance values were computed using the mean decrease in impurity and subsequently aggregated to their parent physical attributes to avoid bias from one-hot encoding. Figure 2 summarizes the normalized importance of the dominant predictors for roof and wall failure.

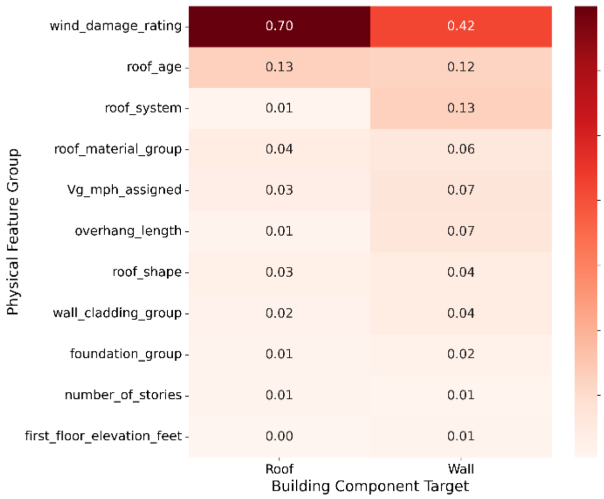

Fig. 2. Normalized feature importance heatmap (aggregated by physical attribute).

Wind hazard intensity, represented by the assigned peak 3-second gust speed, emerges as the most influential predictor for both roof and wall damage. This result is consistent with wind engineering theory, confirming that increasing aerodynamic demand is the primary driver of envelope failure across construction types. The dominance of gust speed across both targets supports the physical validity of the learned fragility relationships.

For roof performance, structural aging is identified as a secondary but significant contributor. Roof age exhibits a strong influence on failure probability, reflecting degradation mechanisms such as material embrittlement, fastener fatigue, and cumulative weathering that reduce resistance over time. In contrast, roof geometric attributes and material groupings show moderate influence, indicating that while material choice matters, its effect is conditional on both hazard intensity and age.

Wall performance is governed by a different hierarchy of predictors. After wind intensity, wall cladding type is the most influential variable. Lightweight cladding systems, particularly vinyl-based assemblies, contribute disproportionately to wall failure probability compared to heavier masonry or fiber cement systems. This finding highlights the sensitivity of wall systems to pressure-driven detachment and debris impact mechanisms under extreme winds.

Elevation-related variables, including first-floor elevation and foundation type, exhibit comparatively lower direct influence on roof and wall failure probabilities. Their contribution is indirect, primarily modifying exposure rather than governing component resistance. The consistently low importance of foundation-related variables further supports the decision to exclude foundation damage as a statistically modeled target due to its rarity in the dataset.

Figure 3 presents a grouped comparison of the relative importance of major physical drivers across roof and wall models. While wind intensity dominates both components, the divergence in secondary drivers underscores the need for component-specific resilience strategies. Roof vulnerability is strongly tied to aging and maintenance cycles, whereas wall vulnerability is more sensitive to material selection and system type.

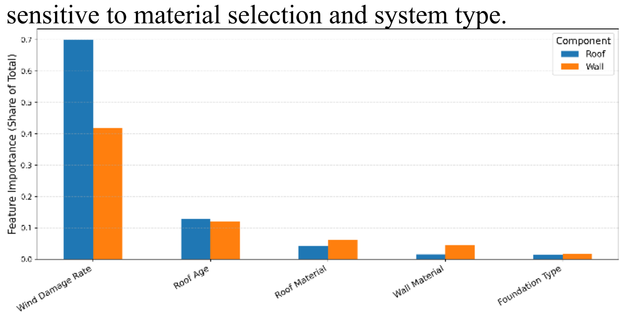

Fig. 3. SHAP summary plots for roof and wall damage models

4.3 Feature Attribution and Directionality (SHAP Analysis)

To complement aggregated feature importance, SHapley Additive exPlanations (SHAP) were used to examine the direction and variability of individual predictor contributions to failure probability. Figure 4 presents SHAP summary plots for the roof and wall damage models, where each point represents a building observation and its contribution to the predicted probability of failure.

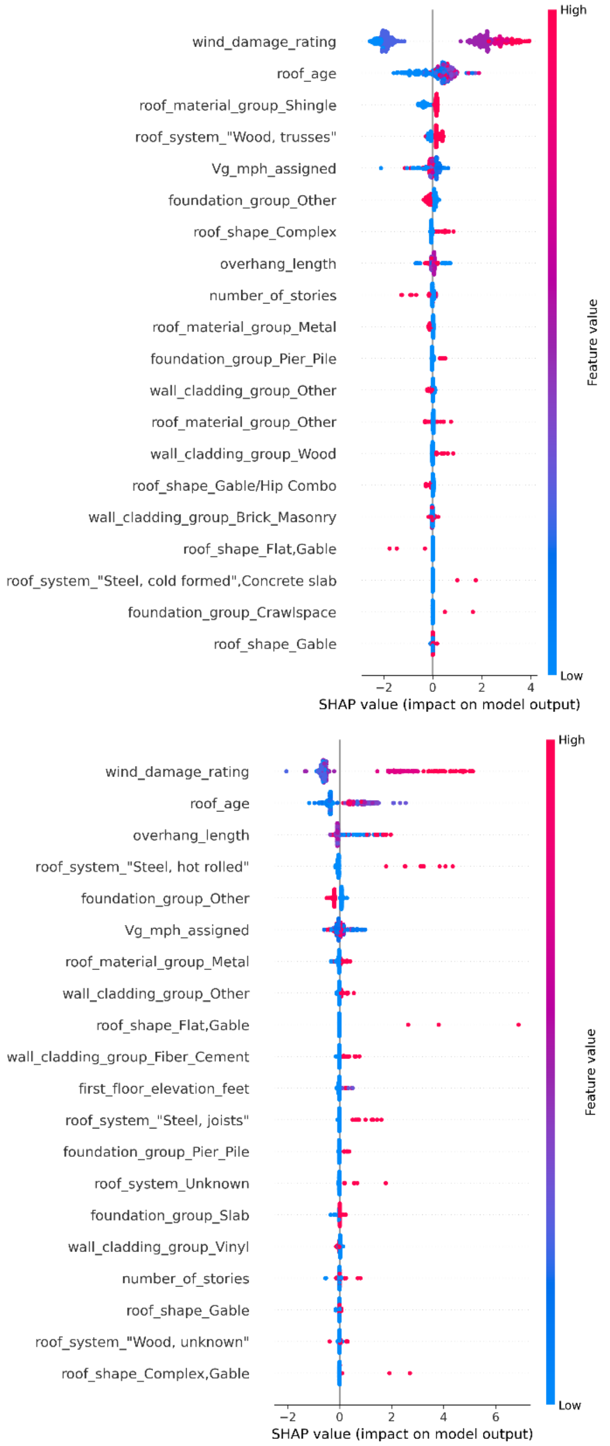

Fig. 4. Relative importance of major physical drivers across components.

Across both components, wind-related variables exhibit the largest SHAP magnitudes and consistent positive contributions at higher values, confirming their role as dominant global drivers of failure probability. Roof age shows a positive but more dispersed influence for roof damage, indicating that aging effects are conditional and interact with other design attributes rather than acting uniformly across the dataset.

Material-related predictors, including grouped roof material and wall cladding type, display narrower SHAP distributions centered closer to zero. This suggests that material choice influences failure probability primarily under specific hazard and configuration conditions rather than exerting a uniform effect across the building stock. Overall, the SHAP analysis confirms that hazard intensity governs baseline failure risk, while construction attributes modulate that risk locally.

4.4 Shape of Hazard–Failure Relationships (PDPs)

Beyond identifying dominant predictors, it is essential to examine how failure probability varies with changes in key hazard and structural variables. To this end, partial dependence plots (PDPs) were generated for selected continuous predictors, allowing the marginal effect of each variable on the predicted probability of failure to be evaluated while averaging over all other inputs.

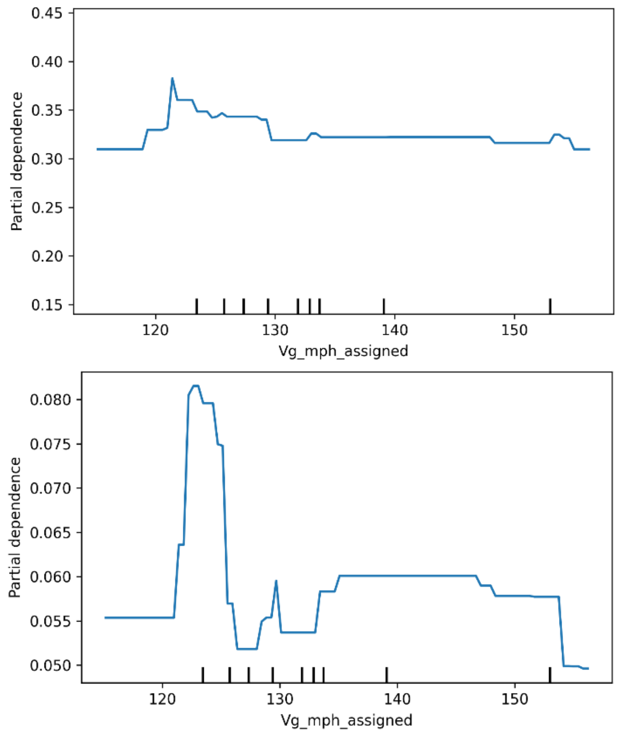

Fig. 5. Partial dependence plots for roof and wall failure probability as a function of peak gust speed.

Figure 5 shows the partial dependence of roof and wall failure probability on peak gust speed. For both components, failure probability increases with wind intensity. Roofs show a gradual rise followed by a sharper increase at higher gusts, while walls show a steeper transition over a narrower range. Additional PDPs for roof age show higher roof failure probability for older roofs, consistent with aging and code-vintage effects.

4.1 Implications for Hazard-Adjusted Carbon Assessment

The preceding results establish that component failure probability is primarily governed by hazard intensity, with secondary modulation by condition and system attributes such as roof age and envelope material class. Within the hazard-adjusted lifecycle formulation (Eq. 2), these model outputs act as the probabilistic link between wind exposure and repair-related greenhouse gas emissions. Specifically, for a given service life T and hazard occurrence rate λ, the expected repair contribution for component i scales with the product $\lambda T P_f(X_i)$, implying that differences in predicted fragility translate directly into differences in expected repair frequency.

The dominance of gust wind speed across both roof and wall models indicates that hazard environment largely sets the baseline emissions risk from repairs. In high-wind regions, even modest reductions in P_f can meaningfully reduce expected lifecycle emissions because the repair term is accumulated over repeated hazard exposure. Conversely, in low-hazard regions the same design change yields smaller lifecycle benefits, even if it reduces failure probability at the event scale. This highlights that the carbon implications of resilience interventions are inherently site-dependent through λ, while the fragility models quantify the design-dependent portion through $P_f(X_i)$.

Material and system attributes influence the lifecycle term in two coupled ways: by altering the probability of failure and by determining the emissions intensity of repair actions through $GWP_{\text{repair},i}$. The feature attribution results indicate that these effects are conditional rather than uniform, suggesting that low-emission materials are not automatically low-emission over the lifecycle if they are associated with higher repair likelihood under strong winds. Accordingly, the proposed framework provides a consistent mechanism to evaluate design alternatives by jointly considering (i) hazard-driven replacement likelihood from the fragility models and (ii) component-specific repair emissions from LCA inventories. Because this study emphasizes methodological development, numerical lifecycle emissions are not reported; however, the results demonstrate that the model structure is sufficient to quantify hazard-adjusted embodied emissions once repair inventories are specified.

5 Discussion

5.1 What the learned fragility structure implies about envelope vulnerability

Across both components, the models consistently attribute failure risk primarily to wind hazard intensity, with gust speed emerging as the dominant driver and showing a monotonic increase in predicted failure probability. This is an important check on physical credibility: the response shape aligns with wind loading fundamentals, and the SHAP attribution does not indicate nonphysical sign reversals at high hazard levels. In practical terms, this means that the hazard environment sets a baseline risk that cannot be “designed away” without materially changing the component’s resistance.

The secondary drivers differ by component, which supports the need for component-specific resilience strategies rather than a single generic “resilience score.” For roofs, roof age is a strong contributor, indicating that deterioration and code vintage effects meaningfully shift fragility, even under similar wind intensity. For walls, the prominence of cladding group indicates that system choice and attachment behavior are central to failure risk, consistent with pressure-driven detachment and debris interaction mechanisms. Importantly, construction attributes act conditionally rather than universally: SHAP patterns show that many system and material variables cluster near zero for most observations and become influential under higher hazard exposure or in specific configurations. This conditional structure matters for decision support because it implies that the benefit of an intervention is not uniform across a building stock.

5.2 Implications for hazard-adjusted embodied emissions

Even without reporting numerical lifecycle emissions, the results support a clear methodological implication for hazard-adjusted embodied GHG assessment. In Eq. (2), the expected repair-related emissions scale with $\lambda T P_f(X_i) GWP_{\text{repair},i}$. Because the models show that P_f increases strongly with gust intensity, the same material or retrofit choice can produce very different lifecycle outcomes depending on the local hazard rate λ. This directly challenges a common implicit assumption in embodied carbon discussions, namely that lower initial emissions necessarily correspond to lower lifecycle emissions in hazard-prone regions. Under repeated exposure, a modest change in failure probability can accumulate into a large expected repair term over the service life, particularly for frequently replaced components such as roof coverings.

The results also clarify why “material-only” rankings are insufficient for climate adaptation decisions framed as sustainability actions. Material and system selection affects lifecycle emissions through two coupled channels: (i) changing failure probability via resistance and deterioration behavior and (ii) changing repair emissions intensity through $GWP_{\text{repair},i}$. The conditional nature of material contributions observed in the attribution results suggests that lifecycle comparisons should be formulated as scenario-dependent, not global. In other words, a material with a low embodied carbon label is not automatically a low-lifecycle-emissions solution if its fragility under the local hazard environment is high. Conversely, a higher-initial-emissions intervention may be justified if it produces a meaningful reduction in expected replacement frequency. This is the central rationale for treating resilience as a sustainability strategy in a quantitative way.

5.3 Robustness, limitations, and next research steps

Several limitations affect how the present results should be interpreted and how the framework can be extended. Although a train–test split was used to evaluate model performance, further assessment using cross-validation or event-based holdout testing would strengthen confidence in model generalization across different hurricane contexts. Future journal extensions can address this by explicitly separating training and testing by event to evaluate transferability.

Generalizability beyond these two hurricanes comes from the fact that the classifier is driven primarily by gust intensity and building attributes, which are event-agnostic inputs. In principle, the same workflow can be applied to any wind event if comparable post-event observations and a site-specific wind field are available to assign V_{gust}. However, because the model was trained on only two events, external validity should be confirmed through event-based holdout testing and, when possible, re-training or calibration with additional hurricanes covering other regions and construction vintages.

Uncertainty in hazard assignment also varies between events. As shown by the nearest-neighbor matching distances for the ARA wind fields, Hurricane Michael exhibits substantially coarser spatial resolution than Hurricane Laura. This difference introduces exposure misclassification that may attenuate true hazard–damage relationships in pooled models. Future work should propagate this uncertainty into failure probability estimates through distance-weighted interpolation, multi-node assignment, or explicit uncertainty bounds on assigned gust speeds.

indicators of "any replacement" and do not distinguish damage severity or repair quantity. Foundation damage is too rare in the pooled dataset to support stable supervised modeling, and the analysis focuses exclusively on wind-driven envelope performance rather than multi-hazard interactions. Most importantly for lifecycle carbon assessment, component-level repair emissions $GWP_{\text{repair},i}$ are not specified in this study. As a result, the framework currently demonstrates the probabilistic linkage between hazard exposure and replacement likelihood rather than producing comparative lifecycle emission totals.

Despite these limitations, the proposed workflow establishes a clear and extensible structure for hazard-adjusted embodied greenhouse gas assessment. Future research will pair the learned fragility functions with component-level repair inventories from EPD-based or LCA databases and apply regional hazard rates to compute expected lifecycle CO_2e. This progression will enable quantitative comparison of design and retrofit strategies based on both resilience performance and long-term environmental impact.

6 Conclusion

Finally, the paper defined a hazard-adjusted lifecycle CO_2e approach where expected repair-related emissions build up over time based on three things: (1) how often damaging hazard events occur, (2) the probability that each component will need replacement under those events (estimated from the ML models), and (3) the CO_2e tied to repairing or replacing that component. Component repair CO_2e values were outside the scope of this study, so we did not report hazard-adjusted embodied emissions. Instead, this paper delivers the data linkage and failure-probability models needed to compute those emissions in a follow-up step once repair CO_2e factors are defined. Future work will (i) populate repair CO_2e using EPDs or LCA databases, (ii) add out-of-sample validation (cross-validation and event holdout testing), and (iii) test sensitivity to post-event predictors such as wind damage rating and to uncertainty in wind field assignment, particularly for Hurricane Michael.

Acknowledgment

This work was supported by the New Frontiers in Research Fund - International (NFRF-International) through the project "Participatory collaborative research to enhance climate change adaptation and mitigation in underserved communities in Asia and North- and South America" (Application ID: NFRFI-2023-00561), administered by the Tri-Agency Institutional Programs Secretariat (SSHRC, NSERC, and CIHR).

References

1. T. Lützkendorf and M. Balouktsi, Embodied carbon emissions in buildings: explanations, interpretations, recommendations. Build. Cities **3**, (2022). https://doi.org/10.5334/bc.257
2. X. Tanguay and B. Amor, Assessing the sustainability of a resilient built environment: research challenges and opportunities. J. Clean. Prod. **458**, 142437 (2024). https://doi.org/10.1016/j.jclepro.2024.142437
3. E. Sesana, C. Bertolin, A. Gagnon and J. Hughes, Mitigating climate change in the cultural built heritage sector. Climate **7**, 90 (2019). https://doi.org/10.3390/cli7070090
4. C. De Wolf, F. Pomponi and A. Moncaster, Measuring embodied carbon dioxide equivalent of buildings: a review and critique of current industry practice. Energy Build. **140**, 68–80 (2017). https://doi.org/10.1016/j.enbuild.2017.01.075
5. F. Greer, P. Raftery, G. Brager and A. Horvath, A perspective on tools for assessing the building sector's greenhouse gas emissions and beyond. Environ. Res. Infrastruct. Sustain. **3**, 043001 (2023). https://doi.org/10.1088/2634-4505/ad064d

6. G. Jayathilaka, N. Thurairajah and A. Rathnasinghe, Digital data management practices for effective embodied carbon estimation: a systematic evaluation of barriers for adoption in the building sector. Sustainability **16**, 236 (2023). https://doi.org/10.3390/su16010236
7. Serrano-Fontova et al., A comprehensive review and comparison of the fragility curves used for resilience assessments in power systems. IEEE Access **11**, 108050–108067 (2023). https://doi.org/10.1109/ACCESS.2023.3320579
8. J.P. Camayang, O. Dela Cruz and R. Grutas, Integrating building- and site-specific and generic fragility curves into seismic risk assessment: a PRISMA-based analysis of methodologies and applications. CivilEng **5**, 1011–1041 (2024). https://doi.org/10.3390/civileng5040050
9. E. Fereshtehnejad and A. Shafieezadeh, Multiple hazard incidents lifecycle cost assessment of structural systems considering state-dependent repair times and fragility curves. Earthq. Eng. Struct. Dyn. **45**, 2327–2347 (2016). https://doi.org/10.1002/eqe.2764
10. S.A. Argyroudis, S.A. Mitoulis, M.G. Winter and A.M. Kaynia, Fragility of transport assets exposed to multiple hazards: state-of-the-art review toward infrastructural resilience. Reliab. Eng. Syst. Saf. **191**, 106567 (2019). https://doi.org/10.1016/j.ress.2019.106567
11. E.C. Matthews, C.J. Friedland and F. Orooji, Integrated environmental sustainability and resilience assessment model for coastal flood hazards. J. Build. Eng. **8**, 141–151 (2016). https://doi.org/10.1016/j.jobe.2016.08.002
12. D. Roueche et al., StEER Field Assessment Structural Team (FAST): in StEER - Hurricane Michael. DesignSafe-CI (2020). https://doi.org/10.17603/DS2-5AEJ-E227
13. D. Roueche et al., Hybrid preliminary virtual reconnaissance report, early access reconnaissance report (PRJ-2888 | StEER - Hurricane Laura). DesignSafe-CI (2021). https://doi.org/10.17603/ds2-dha4-g845