Neighbourhood-Scale Typhoon Risk Mapping Using Machine Learning and SAR-Derived Impact Proxies: Cross-Event Evidence from Cebu City, Philippines

Xin Zhou[1], Pavel Hejcik[1], Zhen Jin[1], Wei Chin Teoh[2], Anudari Batsaikhan[1], Shu Tian[3], Junghong Zhou[4]

[1]Institute for Global Environmental Strategies, 2108-11 Kamiyamaguchi, Hayama, Kanagawa, 240-0115 Japan
[2]Regional Centre in Bangkok of the Institute for Global Environmental Strategies, 87 Liberty Square Building, Unit 2103, 21st Floor, Silom Road, Silom, Bangrak, Bangkok 10500, Thailand
[3]Asian Development Bank, ADB Avenue, Mandaluyong City 1550, Metro Manila, Philippines
[4]Singapore Institute of Technology, 1 Punggol Coast Road Singapore 828608, Singapore

zhou@iges.or.jp, hejcik@iges.or.jp, jin@iges.or.jp, teoh@iges.or.jp, batsaikhan@iges.or.jp, stian@adb.org, junhong.zhou@singaporetech.edu.sg

Abstract

Typhoon impacts in the Philippines exhibit strong neighbourhood-scale heterogeneity driven by interactions among hazard intensity, exposure, and vulnerability, yet most operational assessments remain hazard-centric. This study presents a grid-based machine-learning framework for neighbourhood-scale typhoon risk mapping integrating geophysical hazards, exposure and vulnerability indicators, and satellite-derived impact proxies. A Random Forest classification model is implemented at 100 m resolution for Cebu City, using Sentinel-1 Synthetic Aperture Radar (SAR) backscatter change as an all-weather indicator of disturbance. Binary impact labels are generated by applying a disturbance threshold ($\tau = 1.5$) to SAR backscatter change, converting the disturbance signal into a supervised classification target. Because SAR backscatter changes may reflect surface wetness, vegetation disturbance, or flooding effects, the label represents a disturbance proxy rather than a direct measure of structural damage. Predictor variables include ERA5-Land wind and rainfall metrics, population and building exposures, topography, land cover, and proximity to coastal and riverine features. Model performance is evaluated using the Area Under the Receiver Operating Characteristic Curve (AUC) through cross-event generalisation between Typhoon Rai (2021) and Typhoon Phanfone (2019). Results show above-random discrimination (average AUC $\approx$ 0.59) and stronger transferability in built-up areas (AUC $\approx$ 0.69), supporting neighbourhood-scale risk screening under realistic data constraints.

Introduction

Typhoons are the dominant natural hazard affecting the Philippines, generating widespread loss of life, infrastructure damage, and economic disruption. While national agencies provide increasingly accurate forecasts of wind, rainfall, and storm surge, translating these hazards into actionable neighbourhood-scale risk information remains a persistent challenge (Cuadra et al. 2014; Galloway et al. 2025). Local government units (LGU) require spatially explicit guidance on where impacts are likely to occur, not merely where hazards are strongest, particularly under conditions of constrained resources and time pressure.

Importance of Neighbourhood-Scale Typhoon Risk Mapping

Typhoon impacts are highly heterogeneous in space. Even within a single city, adjacent communities can experience very different outcomes because of differences in topography, proximity to coasts and rivers, settlement patterns, land cover, and socio-economic conditions (Cabrera and Lee 2019; Galloway et al. 2025). City-wide indicators can therefore conceal localised hotspots that matter for evacuation planning, resource pre-positioning, and post-event recovery prioritisation.

Yet many operational products remain hazard-centric. In practice, wind, rainfall, flood, and storm surge information are often delivered as separate layers, with limited integration of exposure and vulnerability at neighbourhood scales (UN-GGIM-AP 2017; World Bank 2018). While indispensable for early warning, such products provide limited guidance for rapid decision-making under uncertainty. Neighbourhood-scale risk mapping can address this gap by linking hazard forcing with local exposure and vulnerability in forms directly usable for preparedness and response, particularly in rapidly urbanising typhoon-prone cities where exposure patterns change quickly and institutional capacity is constrained (Galloway et al. 2025; World Bank 2018).

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Limitations of Existing Approaches and Research Gaps

Conventional typhoon risk assessments are commonly framed as *Risk = Hazard × Exposure × Vulnerability* and implemented through mechanism-based models that simulate physical processes such as wind fields, rainfall–runoff, and storm surge, coupled with exposure inventories and vulnerability functions (Cabrera and Lee 2019; World Bank 2018). These approaches are essential for long-term planning and engineering design, but they face well-known challenges at neighbourhood scales: high-resolution data are unevenly available, vulnerability functions are difficult to calibrate locally, and computational and expertise requirements limit operational uptake in many local planning contexts.

Machine learning (ML) and Earth observation offer complementary pathways by enabling empirical integration of heterogeneous datasets and capturing non-linear interactions among risk drivers. Tree-based ensembles are widely used in hazard and impact applications because they can accommodate mixed predictors and complex interactions (Kooshki Forooshani et al. 2024). However, several gaps persist in ML-based typhoon risk studies. Many analyses focus on single events or use random spatial splits, which can inflate performance due to spatial autocorrelation and provide limited evidence of cross-event generalisation. Evaluation often relies on threshold-dependent metrics (e.g., accuracy or recall) that can be misleading under severe class imbalance. Moreover, transferability across land-cover regimes is rarely tested, despite distinct impact mechanisms in built-up versus non-built-up environments. Finally, satellite-derived change indicators—especially SAR-based metrics—are not always framed as relative impact proxies, encouraging over-interpretation as direct damage observations.

Positioning and Contributions of This Study

This study addresses these gaps by developing a neighbourhood-scale typhoon risk mapping framework that combines open geospatial datasets, SAR-derived impact proxies, and transparent Random Forest (RF) classification. The objective is not absolute loss estimation, but identification of relative spatial patterns of potential impact to support prioritisation and early decision-making.

Cebu City is selected because it represents a highly exposed coastal urban environment where population and infrastructure are concentrated in low-lying coastal and riverine areas. Catastrophe-risk assessments estimate average annual losses of about PHP 6.4 billion, with a 40% probability of losses exceeding PHP 67.7 billion within the next 50 years, highlighting substantial asset exposure to typhoon hazards (World Bank 2018). Cebu also typifies the data and capacity constraints faced by many LGUs, making it a relevant testbed for operationally feasible approaches.

The study contributes four elements. First, it develops a 100 m grid-based RF classifier integrating hazard, exposure, and vulnerability indicators. Second, it introduces an all-weather Sentinel-1 SAR disturbance proxy with an empirically justified threshold to generate consistent binary impact labels across events. Third, it evaluates cross-event generalisation between Typhoon Rai (2021) and Typhoon Phanfone (2019), reducing the risk of inflated performance from spatial autocorrelation. Fourth, it examines land-cover-stratified transferability, demonstrating stronger stability in built-up environments than in non-built areas. By relying on open data and transparent modelling, the framework provides a replicable approach for neighbourhood-scale typhoon risk screening in data-constrained urban contexts.

Data and Methodology

Study Area, Events, and Spatial Framework

This study focuses on Cebu City, a densely populated coastal city in the Central Visayas region of the Philippines. Cebu City is exposed to multiple typhoon-related hazards, including strong winds, intense rainfall, coastal flooding, and riverine flooding, while exhibiting strong spatial heterogeneity in exposure and vulnerability due to rapid urbanisation, complex topography, and extensive coastal development (World Bank 2018). These characteristics make the city a suitable testbed for neighbourhood-scale typhoon risk mapping relevant to local government decision-making.

Two recent typhoon events were analysed: Typhoon Phanfone (Ursula, December 2019) and Typhoon Rai (Odette, December 2021). Both events caused substantial impacts in Cebu City and are well covered by Sentinel-1 SAR observations and reanalysis-based hazard datasets. Rai represents a more destructive event dominated by extreme winds and widespread infrastructure damage (AHA Centre 2021), whereas Phanfone provides a contrasting case with comparatively lower intensity but notable rainfall- and flooding-related impacts. Rather than detailing these differences here, they are explicitly examined through cross-event transfer experiments in Section 3.

To enable spatial integration of heterogeneous datasets, the study area was discretised into a regular grid. Two candidate grid resolutions—30 m × 30 m and 100 m × 100 m—were tested. The 30 m grid matches the native resolution of key exposure and vulnerability datasets (e.g. Copernicus DEM at 30 m, ESA WorldCover at 10 m, building footprints at metre scale). However, core hazard variables derived from ERA5-Land have a much coarser native resolution (~9–11 km). When mapped to a very fine grid, these hazard variables show limited spatial variability between neighbouring cells, increasing noise and class imbalance in the classification task. Comparative testing showed that the 100

m grid yields more stable model behaviour, clearer spatial risk patterns, and improved discrimination while retaining sufficient spatial detail for neighbourhood-scale applications. The 100 m grid was therefore adopted for final modelling.

All datasets were projected to UTM Zone 51N (EPSG:32651). Raster and vector inputs were aggregated to grid cells using zonal statistics and spatial joins, and only grid cells intersecting the Cebu City administrative boundary were retained.

Feature Construction

Predictor variables were constructed following the standard hazard–exposure–vulnerability (HEV) framework. Hazard variables describe typhoon forcing and include wind speed and rainfall metrics derived from ERA5-Land reanalysis, calculated over fixed five-day event windows encompassing landfall and immediate post-landfall conditions. Although ERA5-Land has a coarse native resolution (~9 km), these variables are treated as smoothly varying covariates appropriate for statistical downscaling when combined with high-resolution exposure and vulnerability information.

Exposure variables represent the presence and concentration of people and built assets, including population counts and building footprint metrics. Vulnerability variables capture site-specific susceptibility factors, including elevation, slope, distance to coastlines and rivers, and dominant land-cover type. Hazard variables are event-specific, whereas exposure and vulnerability variables are treated as temporally static, enabling both post-event assessment and forward-looking risk screening. Although this simplification does not capture rapid urbanisation, the two analysed typhoon events occurred within a short time window (2019–2021), during which large-scale exposure patterns in Cebu remained broadly stable. Future work could incorporate time-varying exposure datasets.

Table 1 summarises all feature variables, data sources, native resolutions, and conceptual roles in the model. To avoid redundancy, individual predictors are not discussed further in the text.

Potential multicollinearity among related predictors was examined using pairwise correlations. Terrain variables derived from the Copernicus DEM show a moderate correlation between elevation and slope ($r = 0.450$), indicating that they capture complementary terrain characteristics (absolute elevation versus local steepness). Exposure indicators exhibit similarly moderate relationships, including population versus building area ($r = 0.475$), population versus building density ($r = 0.381$), and building area versus building density ($r = 0.482$). All correlations are well below commonly used thresholds for problematic multicollinearity ($|r| \geq 0.7$–0.8). Given the moderate correlations and the robustness of RF models to correlated predictors, all variables were retained to represent distinct aspects of exposure and vulnerability.

Feature group	Representative variables	Data source	Native resolution	Role
Exposure	Population, building density, building area	WorldPop; Microsoft Buildings	100 m; ~1–3 m	Exposure
Topography	Elevation, slope	Copernicus DEM	30 m	Vulnerability
Proximity	Distance to coast, distance to rivers	Derived vectors	Vector	Vulnerability
Land cover	Dominant land-cover class	ESA WorldCover	10 m	Vulnerability
Rainfall hazard	Event totals, extremes, antecedent rainfall	ERA5-Land	~9 km	Hazard
Wind hazard	Mean, extremes, duration metrics	ERA5-Land	~9 km	Hazard
Antecedent condition	Soil moisture (0–7 cm)	ERA5-Land	~9 km	Context

Table 1. Feature Variables, Data Sources, and Native Resolutions.

SAR-Based Impact Proxy and Label Definition

Because spatially exhaustive ground-based damage observations were unavailable, a Sentinel-1 SAR–based impact proxy was used to characterise relative surface disturbance associated with typhoon impacts. SAR backscatter is sensitive to surface roughness, structural geometry, and moisture conditions, making it suitable for detecting changes caused by wind damage, debris, inundation, and vegetation disturbance under all-weather conditions.

For each event, pre- and post-typhoon VV and VH backscatter differences (ΔVV, ΔVH) were computed and normalised using robust z-scores based on the median and interquartile range (IQR) across all grid cells:

$$Z_{p,i} = \frac{\Delta p_i - \text{median}(\Delta p)}{\text{IQR}(\Delta p)}, p \in \{VV, VH\}$$

A composite disturbance magnitude, $L2_{rz}$, was then calculated as the Euclidean norm of the normalised VV and VH changes:

$$L2_{rz,i} = \sqrt{Z_{VV,i}^2 + Z_{VH,i}^2}$$

Robust normalisation and the $L2$ norm place VV and VH on a comparable, scale-independent basis, suppress outliers, and integrate complementary polarisation responses into a single indicator of overall disturbance magnitude.

RF Classification Model

A RF classifier was used to model the relationship between hazard, exposure, and vulnerability features and the binary SAR-derived impact label. Classification was preferred over regression due to noise and uncertainty in SAR backscatter change, severe class imbalance at fine spatial scales, and the operational objective of identifying relative high-risk locations rather than predicting exact damage magnitudes. RF models are used for their robustness to noise and mixed predictors, although label uncertainty remains an important limitation.

To enable classification-based modelling, the continuous SAR-derived impact proxy $L2_{rz}$ was converted into a binary label by applying a disturbance threshold τ. Grid cells with $L2_{rz} > \tau$were labelled as potentially impacted, while remaining cells were labelled as non-impacted.

A sensitivity analysis across candidate thresholds ($\tau = 0.5, 1, 1.5, 2, 2.5$) was conducted to examine class prevalence and cross-event consistency. A value of $\tau = 1.5$ was selected as it yields comparable proportions of impacted grid cells across events while avoiding extreme class imbalance. The use of a fixed disturbance threshold (τ = 1.5) introduces potential label uncertainty, particularly for marginal cases near the threshold, which may introduce label noise and affect classification performance.

For each grid cell, the RF model outputs a risk probability, defined as the proportion of trees voting for the impacted class. The Random Forest classifier was implemented in Python (version 3.10.19) using scikit-learn (version 1.7.2) with n_estimators=900, max_depth=None, min_samples_split=10, min_samples_leaf=3, max_features="sqrt", bootstrap=True, class_weight="balanced_subsample", oob_score=True, random_state=42, and n_jobs=-1; all other parameters followed library defaults. Binary labels were defined using the SAR disturbance threshold \tau=1.5. The classification threshold was selected from out-of-bag training probabilities by maximizing F1, and risk classes were assigned using training-set probability percentiles (top 1%, 1–5%, and 5–20%).

Experimental Design and Model Evaluation

Experimental Design

To assess cross-event transferability, models were trained and evaluated on independent typhoon events rather than random spatial splits, avoiding inflated performance due to spatial autocorrelation. Three experimental configurations were implemented across land-surface conditions with differing damage mechanisms and SAR responses: (i) all land types (Exp1), (ii) built-up areas (Exp2), and (iii) non-built land (Exp3). For each configuration, models were trained on Typhoon Rai (2021) and tested on Typhoon Phanfone (2019), and vice versa, ensuring symmetric bidirectional evaluation (Table 2).

Experiment	Land surface coverage	Training event	Testing event
Exp 1a	All land types	Rai	Phanfone
Exp 1b	All land types	Phanfone	Rai
Exp 2a	Built-up land	Rai	Phanfone
Exp 2b	Built-up land	Phanfone	Rai
Exp 3a	Non-built land	Rai	Phanfone
Exp 3b	Non-built land	Phanfone	Rai

Table 2. Experimental Design for Cross-Event Transfer.

This reciprocal design ensures that performance reflects generalisation across events rather than event-specific artefacts. Comparing results in both directions allows assessment of the robustness and symmetry of learned relationships under differing hazard conditions.

Evaluation Metrics

Model performance was evaluated using complementary, threshold-independent and diagnostic metrics suited to extreme class imbalance.

The Area Under the Receiver Operating Characteristic Curve (AUC) is used as the primary performance indicator. AUC measures the probability that a randomly selected impacted grid cell is assigned a higher risk probability than a non-impacted cell and is insensitive to the prevalence of positive samples. This makes it appropriate for cross-event comparison and risk ranking.

Recall (true positive rate) is reported to aid physical interpretation, as it reflects the model's ability to identify grid cells exhibiting strong SAR-based disturbance. Recall is discussed in conjunction with false negatives to highlight trade-offs between missed impacts and conservative predictions.

Confusion matrices are used as a diagnostic tool to examine the balance between true positives, false positives, true negatives, and false negatives, supporting interpretation of error types under different experimental settings.

Feature Importance Analysis

To improve interpretability, feature importance was assessed using impurity-based (Gini) importance which quantifies each feature's contribution to reducing classification uncertainty across the RF ensemble. This metric provides a

global overview of influential predictors but may favour variables with higher variance or more split opportunities.

Risk Probability Mapping and Classification

For each grid cell, the trained RF classifier outputs a risk probability:

$$\text{risk\_prob}_i = P(y = 1 \mid X_i),$$

representing the likelihood that SAR-based disturbance exceeds the selected threshold τ, given local hazard intensity and exposure–vulnerability conditions.

To enhance robustness and comparability across events, probabilities were converted into categorical risk classes using percentile-based thresholds derived from the training probability distribution: *Very High* (top 1%), *High* (1–5%), *Medium* (5–20%), or *Low* (remaining cells) risk. This approach defines risk classes by relative ranking, ensuring consistent proportions of high-risk cells across events. It captures the upper tail of predicted risk and supports prioritisation of high-risk locations rather than deterministic prediction of damage magnitudes.

Results

Cross-Event Transferability Across Land-Cover Regimes

Cross-event experiments reveal systematic differences in model transferability across spatial domains (Table 3). The positive class rate ranges from approximately 7% to 16% across experiments, indicating that SAR-detected disturbance is spatially concentrated and affects a limited proportion of grid cells.

Experiment	Domain	N	AUC	Recall	AP
Exp 1a	All grids	29540	0.57	0.56	0.18
Exp 1b	All grids	29540	0.57	0.98	0.19
Exp 2a	Built-up	3604	0.69	0.65	0.26
Exp 2b	Built-up	3604	0.62	0.18	0.21
Exp 3a	Non-built-up	25936	0.53	0.47	0.18
Exp 3b	Non-built-up	25936	0.56	1.00	0.19

Table 3. Cross-Event Performance of the RF Classifier. Note: N is the number of grid cells. AUC is the Area Under the Receiver Operating Characteristic, AP denotes average precision, and Recall denotes the true positive rate.

When trained on Typhoon Rai and tested on Typhoon Phanfone using all grid cells (Exp 1a), the RF classifier achieves moderate discrimination performance, with ROC–AUC values around 0.57 and AP around 0.18. This indicates that learned relationships between hazard, exposure, and vulnerability generalise beyond a single event, but with limited strength. Reciprocal transfer (Exp 1b) yields similar AUC but substantially higher recall, likely reflecting differences in event intensity, where the stronger Typhoon Rai produces more pronounced and widespread disturbance signals that are more easily detected by a model trained on the weaker Phanfone event.

Performance improves markedly when analysis is restricted to built-up areas (Exp 2a). In this regime, cross-event AUC increases to 0.69, with the highest AP (≈0.26), demonstrating stronger and more stable transferability within dense urban environments. Built-up areas exhibit more consistent exposure structures and relatively uniform SAR scattering behaviour, enabling the model to learn transferable risk patterns across events. However, reciprocal transfer (Exp 2b) shows substantially lower recall, suggesting that models trained on weaker events may fail to capture more severe structural damage patterns in stronger events. This contrast indicates that, in built-up areas, transferability depends not only on spatial consistency but also on the range of damage intensity represented in the training data.

In contrast, transferability deteriorates in non-built-up areas (Exp 3a), where AUC values fall to 0.53 and AP remains low. These areas are characterised by heterogeneous surface processes, including vegetation flattening, soil moisture variability, and flood-induced backscatter changes, which reduce the consistency of SAR-based impact signatures across events. Reciprocal experiments (Exp 3b) again show very high recall but similarly low AP, reflecting increased false positives under class imbalance and weaker separability of impact signals in these environments.

Confusion matrix diagnostics (Figure 1, illustrated for Exp 2a) indicate that the classifier tends to produce conservative predictions, balancing relatively high recall with low precision. This behaviour is consistent across experiments and reflects the study's objective of probabilistic risk ranking rather than deterministic damage detection.

Taken together, these results highlight two key factors governing transferability: (i) land-cover regime, with built-up areas exhibiting stronger and more stable performance, and (ii) event intensity, where models trained on weaker events may not capture the full range of impact signals present in stronger events. Neighbourhood-scale risk mapping is therefore most reliable in built-up urban contexts when training data adequately represent the spectrum of impact severity.

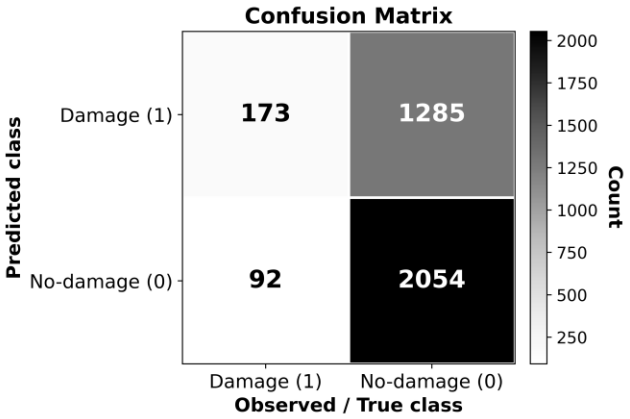

Figure 1: Confusion Matrix for Built-up Land Cross-Event Transfer (Exp 2a); Rows = Predicted, Columns = Observed, $\tau = 1.5$.

Spatial Patterns of Predicted Risk

Figure 2 presents the grid-level categorical risk map for Typhoon Phanfone generated using the model trained on Typhoon Rai (Exp 1a). Predicted higher-risk grids are concentrated along coastal zones and dense urban corridors, particularly in low-lying areas characterised by high building density and population exposure. Inland areas with lower exposure generally exhibit lower predicted risk, even where hazard forcing is non-negligible.

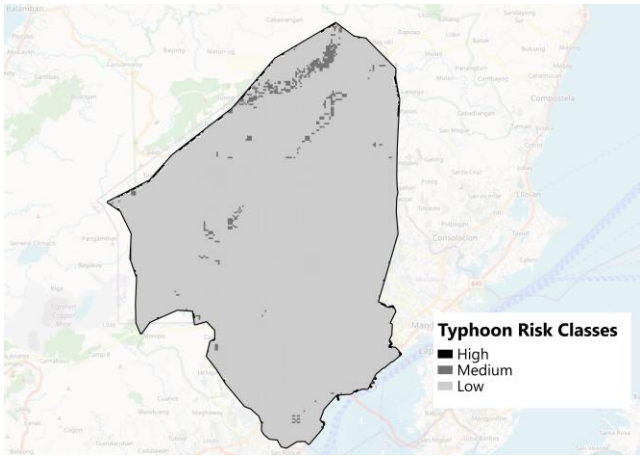

Figure 2: Predicted Risk Classes for Exp 1a (Rai → Phanfone) Across Cebu City, Displayed Over a Geographic Basemap with City Boundary for Spatial Reference.

These spatial patterns differ from raw hazard fields such as wind speed or rainfall intensity, which typically exhibit smoother gradients. In contrast, the predicted risk surface reflects the combined influence of hazard, exposure, and vulnerability, demonstrating that hazard intensity alone is insufficient to explain spatial variations in potential impact. This highlights the importance of integrated modelling for neighbourhood-scale risk prioritisation.

Figure 3 shows the corresponding risk map restricted to built-up areas (Exp 2a). Compared with the city-wide result, the built-up-only map exhibits clearer spatial coherence and more stable clustering of higher-risk grids along coastal and urbanised zones. This pattern is consistent with the stronger cross-event transferability observed in built-up areas and reflects more stable exposure conditions and SAR scattering behaviour.

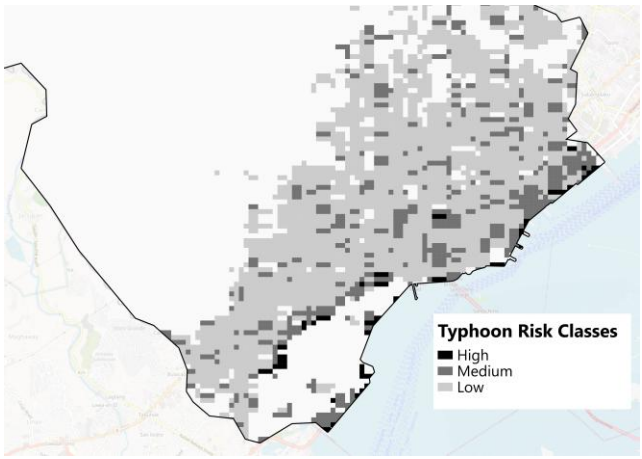

Figure 3: Predicted Risk Classes for Exp 2a (Rai → Phanfone) in Built-up Areas of Cebu City, Displayed Over a Geographic Basemap with City Boundary.

Together, these maps illustrate how the proposed framework supports neighbourhood-scale risk prioritisation by identifying spatially concentrated hotspots of potential impact.

Feature Importance and Interpretation

Feature importance analysis for Exp 1a (Figure 4) indicates that exposure and vulnerability variables dominate model decisions. Distance to the coastline, population density, distance to rivers, terrain slope, and elevation consistently rank among the most influential predictors, capturing the spatial concentration of people and assets and their susceptibility to typhoon impacts.

Hazard-related variables, including wind and rainfall metrics, contribute to model performance but are generally of secondary importance. This does not imply that hazards are unimportant; rather, it reflects their role as triggers whose impacts are strongly mediated by exposure and vulnerability. Areas experiencing similar hazard intensity can exhibit markedly different risk levels depending on settlement patterns and terrain characteristics.

Importance values are interpreted qualitatively rather than causally. Although multicollinearity among key predictors was found to be limited, impurity-based (Gini) importance may still be influenced by feature correlation. Nevertheless, the observed ranking is consistent with the spatial risk patterns and supports the conceptual framing of typhoon risk as the interaction of hazard, exposure, and vulnerability.

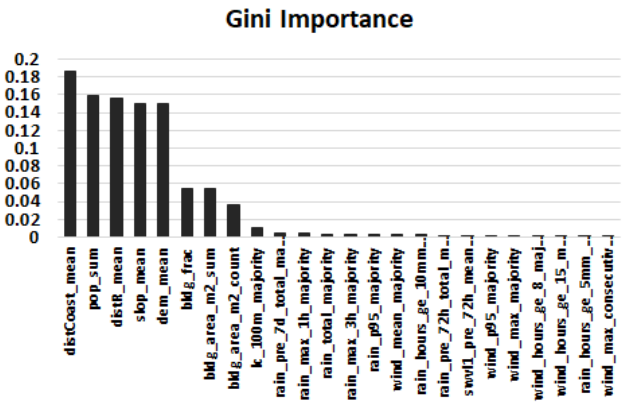

Figure 4: Gini Feature Importance for Exp 1a Showing Relative Contribution of Input Variables to the RF Model

Discussion and Future Research Agenda

Risk vs Hazard

The divergence between hazard fields and predicted risk surfaces is not a modelling artefact but an intrinsic property of integrated risk assessment. High wind speeds or rainfall intensities do not necessarily translate into high risk unless they coincide with concentrated exposure and vulnerability. By jointly learning hazard, exposure, and vulnerability indicators, the RF classifier captures these interactions and produces spatial risk patterns that differ fundamentally from hazard maps. This distinction is critical for preparedness and prioritisation, as it shifts attention from where hazards are strongest to where impacts are most likely to occur.

The moderate AUC values (~0.57) reflect the difficulty of predicting heterogeneous proxy-based impacts and are consistent with the objective of relative risk ranking.

Value and Limitations of SAR-Based Impact Proxies

SAR-derived backscatter change provides a scalable, all-weather proxy for screening potential impacts, particularly in contexts where spatially exhaustive damage data are unavailable. The cross-event experiments indicate that SAR-based disturbance metrics encode transferable information on relative impact likelihood, supporting consistent risk ranking across events when combined with exposure and vulnerability indicators.

At the same time, SAR signals are inherently ambiguous with respect to specific damage mechanisms. Backscatter changes may reflect a mixture of structural damage, vegetation disturbance, debris accumulation, surface wetness, or flooding, and should therefore not be interpreted as direct observations of physical damage. This limitation underscores the importance of treating SAR-derived metrics as relative impact proxies rather than absolute damage measures.

The absence of spatially explicit ground-truth damage data remains a key limitation. While SAR provides a scalable proxy, incorporating post-disaster surveys, high-resolution optical imagery, or official damage assessments would be essential to improve label reliability and model validation. Validation through complementary data sources remains essential. On-site damage assessments, pre- and post-event optical imagery, and post-disaster reports provide critical reference information for calibrating thresholds, interpreting SAR signals, and strengthening confidence in model outputs.

In addition, both SAR-derived disturbance metrics and downscaled hazard variables introduce sources of uncertainty. SAR backscatter changes may be affected by speckle noise, viewing geometry, and surface moisture conditions unrelated to damage, while ERA5-Land variables, due to their coarse native resolution (~9 km), may not fully capture fine-scale variability in wind and rainfall when downscaled to neighbourhood grids. These factors may introduce spatial smoothing or local misrepresentation of hazard intensity, particularly in complex urban and coastal environments.

Implications for Local Governments

The proposed framework has direct relevance for LGUs. It supports pre-event risk screening by linking forecast hazards with existing exposure and vulnerability patterns, and post-event prioritisation by identifying neighbourhoods likely to require rapid assessment or intervention. The use of percentile-based risk classes avoids reliance on unstable absolute probability thresholds and facilitates transparent, relative prioritisation under uncertainty. Importantly, the cross-event transferability demonstrated in this study suggests that the framework can be applied beyond a single typhoon event, enabling LGUs to reuse trained models and risk representations to support preparedness and response planning across future events, particularly in typhoon-prone, rapidly urbanising cities in developing-country contexts.

Regime Dependence and Transferability

The land-cover-stratified experiments highlight that model transferability is regime dependent. Built-up areas exhibit

more stable exposure patterns and SAR scattering behaviour, leading to substantially higher cross-event performance than non-built-up areas. This finding reinforces the importance of explicitly accounting for land-cover context when developing and applying neighbourhood-scale risk models, and cautions against uniform interpretation of risk estimates across heterogeneous surface regimes.

Future Research Agenda

Future research should focus on: (i) closer co-design with local governments to develop operational dashboards linking risk maps with early warning systems; (ii) extension to additional typhoon events to strengthen assessment of model generalisability and uncertainty; (iii) systematic validation using ground surveys, post-disaster reports, and optical imagery; and (iv) integration of richer and time-varying socio-economic and structural indicators (e.g. building height, construction materials, and dynamic population or infrastructure data). Performance uncertainty (e.g., AUC confidence intervals) was not explicitly quantified and should be addressed in future work.

Conclusions

This study demonstrates that neighbourhood-scale typhoon risk mapping can be implemented using open-access geospatial data, SAR-based impact proxies, and interpretable machine-learning classification. Cross-event and land-cover-stratified experiments show that RF models generalise beyond single events and capture physically meaningful risk patterns, particularly in urban environments. Performance is moderate overall (AUC ≈ 0.57), improving to ≈ 0.69 in built-up areas and declining to ≈ 0.53 in non-built environments. By focusing on relative risk ranking rather than absolute damage estimation, the proposed framework complements hazard forecasts and provides actionable information to support preparedness, response prioritisation, and urban resilience planning. This study contributes to ongoing funded efforts to advance scalable, data-driven disaster risk assessment and climate resilience.

Acknowledgments

This work is funded by the Asian Development Bank under the technical assistance project “TA-10344 REG: Role of Artificial Intelligence in Asia's Growth and Development (Typhoon Risk Mapping in the Phils)”, and the Institute for Global Environmental Strategies (IGES) under FY2025 Strategic Research Fund project entitled “Machine-Learning-Based Community-Level Typhoon Risk Mapping: A Proof-of-Concept for Cebu City”.

References

ASEAN Coordinating Centre for Humanitarian Assistance on disaster management (AHA Centre). 2021. Situation Update No. 2– Typhoon Rai in the Philippines. https://ahacentre.org/wp-content/uploads/2021/12/AHA-Situation_Update-no2-TC-Rai-Philippines_23Dec2021.pdf. Accessed: 2026-01-15.

Cabrera, J. S.; and Lee, H. S. 2019. Flood-Prone Area Assessment Using GIS-Based Multi-Criteria Analysis: A Case Study in Davao Oriental, Philippines. *Water* 11(11): 2203. doi.org/10.3390/w11112203.

Cuadra, C.; Biton, N. I.; Cabacaba, K. M.; Santiago, J.; Suarez, J. K.; Lapidez, J. P.; Lagmay, A. M. F.; and Malano, V. 2014. Development of Inundation Map for Bantayan Island, Cebu Using Delft3D-Flow Storm Surge Simulations of Typhoon Haiyan. Project NOAH Open-File Reports 3: 37–44.

Galloway, E. G.; Catto, J. L.; Luo, C.; and Siegert, S. 2025. Tropical Cyclone Impact Data in the Philippines: Implications for Disaster Risk Research. *Natural Hazards* 121(13): 15275–96. doi.org/10.1007/s11069-025-07394-x.

Forooshani, M. K.; van den Homberg, M.; Kalimeri, K.; Kaltenbrunner, A.; Mejova, Y.; Milano, L.; Ndirangu, P.; Paolotti, D.; Teklesadik, A.; and Turner, M. L. 2024. Towards a Global Impact-Based Forecasting Model for Tropical Cyclones. *Natural Hazards and Earth System Sciences* 24(1): 309–29. doi.org/10.5194/nhess-24-309-2024.

Regional Committee of United Nations Global Geospatial Information Management for Asia and the Pacific (UN-GGIM-AP). 2017. Collection of Best Practices on the Use of Geospatial Information for Disaster Risk Reduction. https://un-ggim-ap.org/sites/default/files/media/docs/Knowledge%20Base/Disaster%20Risk%20Reduction/Best%20Practices_Geospatial%20Information%20for%20Disaster%20Risk%20Reduction.pdf. Accessed: 2025-12-10.

World Bank. 2018. Province Risk Profile: Cebu Philippines Catastrophe Risk Assessment and Modeling. https://www.financialprotectionforum.org/third-party/microsite_3/provinces/cebu.html. Accessed: 2026-01-20.