

<https://doi.org/10.1038/s44304-025-00098-z>

Causal spatially heterogeneous Bayesian networks with GPs and normalizing flows for seismic multi-hazard estimation

Xuechun Li¹, Shan Gao¹, Runyu Gao¹ & Susu Xu^{1,2}✉

Post-earthquake hazard and impact estimation are critical for effective disaster response, yet current approaches face significant limitations. Traditional models employ fixed parameters regardless of geographical context, misrepresenting how seismic effects vary across diverse landscapes, while remote sensing technologies struggle to distinguish between co-located hazards. We address these challenges with a spatially-aware causal Bayesian network that decouples co-located hazards by modeling their causal relationships with location-specific parameters. Our framework integrates sensing observations, latent variables, and spatial heterogeneity through a novel combination of Gaussian Processes with normalizing flows, enabling us to capture how same earthquake produces different effects across varied geological and topographical features. Evaluations across three earthquakes demonstrate Spatial-VCBN achieves Area Under the Curve (AUC) improvements of up to 35.2% over existing methods. These results highlight the critical importance of modeling spatial heterogeneity in causal mechanisms for accurate disaster assessment, with direct implications for improving emergency response resource allocation.

Earthquakes cause harm not only through direct ground shaking but also by triggering secondary ground failures such as landslides and liquefaction. These combined effects lead to devastating consequences, including structural damage and human casualties. A striking illustration is the 2021 Haiti earthquake, which initiated over 7000 landslides covering more than 80 square kilometers. This catastrophic event resulted in damage or destruction to over 130,000 buildings, claimed 2248 lives, and left more than 12,200 people injured¹. Rapidly identifying where and how severely ground failures and structural damage have occurred following an earthquake is essential for effective victim rescue operations within the crucial “Golden 72 Hour” window, and plays a vital role in developing effective post-disaster recovery plans^{2,3}.

Over the years, researchers have developed various approaches for estimating the location and intensity of earthquake-induced ground failures and building damage. Traditional approaches to seismic hazard assessment fall into two main categories: physical and statistical models^{4–7}. Physical models apply fundamental geotechnical principles, such as the Newmark displacement method for landslides or liquefaction potential indices. While foundational, these approaches require detailed geotechnical data often unavailable during rapid response and frequently oversimplify complex physical dynamics. Statistical models estimate hazards using geospatial susceptibility indicators paired with ground motion data, calibrated on

historical events. However, current approaches employ fixed parameters regardless of location. This one-size-fits-all approach misrepresents how seismic effects vary across different geological compositions, topographical features, and infrastructure, leading to significant estimation inaccuracies.

Remote sensing technologies, particularly Interferometric Synthetic Aperture Radar (InSAR), have revolutionized rapid post-earthquake assessment capabilities. InSAR works by comparing phase differences between radar signals captured at different times, enabling detection of surface deformation with high-resolution precision⁸. Damage Proxy Maps (DPMs) generated by the NASA’s Advanced Rapid Imaging and Analysis (AIRA) team, which is a state-of-the-art InSAR product, identify ground changes by analyzing radar measurement disparities before and after seismic events⁹. Although remote sensing technologies facilitate expedited hazard evaluation¹⁰, these approaches typically address individual hazards and face challenges in differentiating between overlapping signals from multiple hazard types while filtering out environmental noise¹¹. This limitation creates a critical information gap for emergency responders who need a comprehensive understanding of multiple concurrent hazards.

Recent advances have introduced causal Bayesian networks to address the limitations of remote sensing technologies by disentangling cascading earthquake impacts^{12–19}. These probabilistic models leverage variational inference to track causal chains from initial earthquake events through

¹Center for Systems Science and Engineering, Department of Civil and Systems Engineering, Johns Hopkins University, Baltimore, MD, USA. ²Data Science and AI Institute, Johns Hopkins University, Baltimore, MD, USA. ✉e-mail: susuxu@jhu.edu

multiple hazards to ultimate building damage. A recent approach²⁰ attempted to account for spatial variations using bilateral filters, which is a technique that creates weighted averages based on proximity between locations^{21,22}. While this improved upon uniform-parameter models, bilateral filtering fundamentally treats spatial dependency as a simple averaging operation without modeling how the underlying causal mechanisms themselves vary across different geological settings^{23,24}. For example, the same level of ground shaking might have a much stronger causal effect on triggering landslides in areas with steep slopes and loose soil compared to flat areas with stable bedrock²⁵. Similarly, building damage resulting from liquefaction will vary dramatically depending on foundation types and subsurface conditions. These location-specific causal relationships cannot be captured by simple spatial averaging²⁶.

To address these limitations, we propose a spatially-aware variational causal Bayesian network (Spatial-VCBN) that represents a fundamental shift in approach: rather than modeling spatial correlation as an afterthought, we directly model the variation in causal mechanisms themselves. The distinction is crucial because while previous approaches might apply spatial smoothing to already-estimated parameters, our framework incorporates spatial heterogeneity directly into the causal structure itself. Such integration captures the reality that the same earthquake generates dramatically different effects depending on local conditions. The physical intuition behind Spatial-VCBN is that causal relationships in earthquake scenarios are fundamentally location-dependent, a principle well-established in geophysics but rarely incorporated into predictive models. While simpler approaches like linear spatial interpolation or kernel smoothing might seem sufficient, the highly non-linear and potentially multi-modal nature of earthquake causal effects demands more sophisticated methods. Our framework consists of three key components: (1) observable variables including geospatial features and DPMs derived from satellite imagery; (2) latent hazard/impact variables representing unobserved landslides, liquefaction, and building damage; and (3) spatially-varying causal coefficients that quantify the strength of causal relationships at each location.

A key innovation in Spatial-VCBN is the modeling of these spatially-varying causal coefficients using a combination of Gaussian Processes (GPs) with normalizing flows. GPs serve as spatial priors that ensure causal relationships vary smoothly across regions based on geophysical similarities²⁷, while normalizing flows provide flexible, non-linear transformations that capture complex^{28,29}, non-Gaussian distributions of causal effects. Conceptually, normalizing flows transform simple probability distributions into more complex ones through a sequence of invertible mappings, allowing us to represent the richly varied ways that earthquake forces translate into surface effects across different terrains. The use of Gaussian Processes is particularly well-suited for this problem because they naturally model spatial correlation while allowing for location-specific variations. Normalizing flows complement this approach by transforming simple distributions into more complex ones through a series of invertible mappings, enabling Spatial-VCBN to capture the highly non-linear and potentially multi-modal nature of causal effects in earthquake scenarios. Our ablation studies reveal that an optimal flow number of $K = 6$ is required to adequately model these complex relationships, confirming that simpler distributional assumptions would fail to capture the physical reality of earthquake impacts. This is crucial because the relationships between geological features, seismic waves, and resulting hazards often follow complex, non-Gaussian patterns that simpler distributional assumptions would fail to capture.

By combining GPs with normalizing flows, our framework can represent complex, non-Gaussian distributions of causal effects that better reflect the physical reality of how earthquake impacts propagate through diverse environments. This probabilistic approach also enables principled quantification of uncertainty in hazard estimates, which is crucial for effective disaster response planning. For inference, we implement a stochastic variational approach with an expectation-maximization algorithm that alternates between updating posteriors of unobserved variables and refining model parameters. To handle large geographical regions efficiently,

we apply a local pruning strategy that exploits the natural sparsity in real-world causal networks, achieving processing speeds of ~ 0.94 s/km² with GPU acceleration, making Spatial-VCBN viable even for large-scale events like the Haiti earthquake (15,970 km²).

Our empirical evaluations demonstrate that this spatially-variant approach significantly improves hazard and impact estimation accuracy, with AUC improvements of up to 35.2% over prior probability baselines and 5.5% over state-of-the-art methods across three earthquake events (Haiti, Puerto Rico, and Turkey-Syria). These substantial improvements are not merely incremental advances but represent a step-change in our ability to rapidly assess complex disaster scenarios with the precision needed for effective emergency response. The model shows remarkable robustness in signal-constrained environments, successfully extracting coherent hazard patterns even from noisy DPM signals, as demonstrated in mountainous and coastal regions. Particularly noteworthy is the strong performance of Spatial-VCBN at low false positive rates, which is crucial for effective resource allocation in disaster response operations where false alarms can waste limited resources. These results highlight the critical importance of modeling spatial heterogeneity in the actual causal mechanisms themselves, rather than merely accounting for spatial proximity, for effective disaster assessment and response planning.

Results

Data description

We evaluate our framework on three earthquake events: the 2020 Puerto Rico earthquake (M6.4), 2021 Haiti earthquake (M7.2), and 2023 Turkey-Syria earthquake sequence (M7.8).

The 2020 Puerto Rico earthquake A magnitude 6.4 earthquake hit the southwest part of Puerto Rico on January 7, 2020. The ARIA team created DPMs using SAR images from the Sentinel-1 satellite to identify potentially damaged areas²⁹. Researchers from the USGS, the University of Puerto Rico Mayagüez, the GEER team, and the StEER team later conducted field reconnaissance to collect ground truth observations^{30–32}. Post-earthquake reports documented that at least 300 landslides were triggered near the epicenter³⁰.

The 2021 Haiti earthquake On August 14, 2021, a magnitude 7.2 earthquake struck the southern peninsula of Haiti. The StEER team and GEER team later collected ground truth inventories for landslides and building damage^{33–35}. According to post-disaster reports, the earthquake resulted in at least 2248 human fatalities, destroyed 53,815 buildings, and damaged 83,770 structures throughout Grand Anse, Nippes and Sud^{36,37}.

The 2023 Turkey-Syria earthquake sequence On February 6, 2023, an earthquake of Mw 7.8 and its aftershocks caused unparalleled destruction in Turkey and Syria. This disaster resulted in more than 55,000 deaths, displaced 3 million people in Turkey and 2.9 million in Syria, and caused damage or destruction to at least 230,000 buildings^{14,17}. The ARIA team generated DPM derived from synthetic aperture radar (SAR) images on Feb. 10, 2023 by the Copernicus Sentinel-1 satellites operated by the European Space Agency³⁸. Ground truth was later collected and reported by the Turkish Ministry of Environment^{14,39}.

Evaluation metrics and benchmarks

We evaluate our framework using multiple complementary metrics to provide a comprehensive assessment of performance. The primary evaluation is based on the receiver operating characteristics (ROC) curve and its area under curve (AUC) metric⁴⁰. ROC curves plot True Positive Rate against False Positive Rate (FPR) across varying decision thresholds, providing a threshold-independent assessment that is particularly valuable in disaster contexts where optimal classification thresholds may vary by event type, geographic region, or hazard category. This approach aligns well with the probabilistic nature of our framework, as both our spatially-aware causal model and the benchmark methods produce confidence scores rather than hard classifications.

To complement the AUC metric, we also employ the F1 score, which is the harmonic mean of precision and recall. The F1 score is calculated as

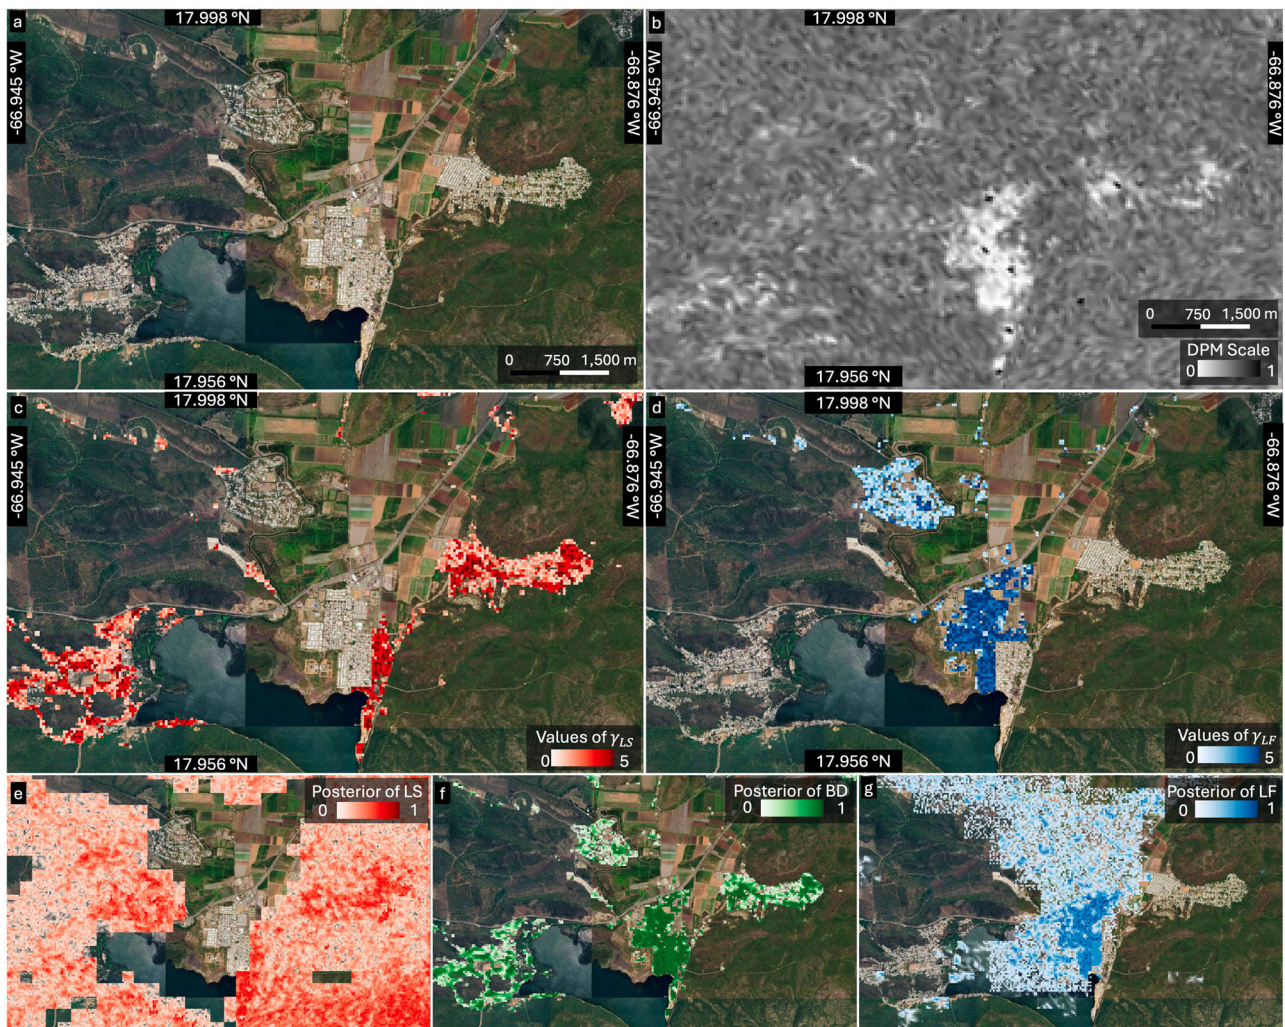


Fig. 1 | Spatial distribution of hazards, causal parameters, and damage following the 2020 Puerto Rico earthquake. **a** Google satellite imagery of the study area with extent of -66.945°W , 17.956°N to -66.876°W , 17.998°N . **b** DPM derived from satellite imagery. **c** Spatial distribution of γ_{LS} values, showing the causal parameter

strength from landslides (LS) to building damage (BD). **d** Spatial distribution of γ_{LF} values, showing the causal parameter strength from liquefaction (LF) to building damage (BD). **e** Posterior probability of landslide occurrence. **f** Posterior probability of building damage occurrence. **g** Posterior probability of liquefaction occurrence.

$F1 = 2 \times \frac{\text{precision} \times \text{recall}}{\text{precision} + \text{recall}}$, where precision represents the ratio of correctly identified hazards to all predicted hazards, and recall captures the proportion of actual hazards that were correctly identified. Unlike AUC, which evaluates performance across all possible thresholds, F1 score requires selecting a specific classification threshold. For consistent comparison, we use the threshold that maximizes the F1 score for each model and hazard type. This metric is particularly informative for disaster response applications, where balancing false alarms (precision) with missed hazards (recall) is critical for effective resource allocation.

Although our framework is fundamentally an unsupervised model that approximates true posteriors through variational inference, these evaluation metrics enable objective comparison with supervised approaches. Additionally, we utilize the variational bound as an internal metric to optimize model parameters and determine the optimal flow length for the normalizing flow component.

Spatially-varying causal influence mapping between multi-hazards and structural damage

Figure 1 illustrates the complex spatial relationships between different types of hazard, their causal influence, and surface deformation patterns following the 2020 Puerto Rico earthquake with an extent of -66.945°W , 17.956°N to

-66.876°W , 17.998°N . The DPM (Fig. 1b) captures surface deformation and ground changes, serving as an observable proxy that Spatial-VCBN interprets through causal relationships.

The spatial distribution of the causal parameters γ_{LS} and γ_{LF} , which are shown in Fig. 1c, d, reveals how different hazard mechanisms contribute to the damage of buildings throughout the landscape. Areas with elevated γ_{LS} values indicate locations where landslides have a stronger causal influence on building damage, predominantly in the hillier portions of the study area. These high-coefficient regions align well with areas showing a high posterior probability of landslide occurrence (Fig. 1e), suggesting that Spatial-VCBN successfully captures the spatial specificity of landslide impacts.

The effects of liquefaction show distinctly different spatial patterns, with high values of γ_{LF} concentrated in the central lowlands and coastal areas. This clear spatial separation between high γ_{LS} and high γ_{LF} regions validates Spatial-VCBN assumption that landslides and liquefaction generally do not co-occur at the same location due to their different geological requirements. The posterior probability of building damage (Fig. 1f) represents the result of our causal network, showing how the model integrates information from both types of hazard. Importantly, the model does not simply translate DPM signals directly into damage estimates but interprets these signals through the learned causal structure and spatially varying parameters.




Fig. 2 | Spatial distribution of causal parameters and DPM signals in a high-deformation area. **a** Damage Proxy Map (DPM) showing surface deformation with brighter areas indicating higher change detection values in the study area with extent of -66.946°W , 17.958°N to -66.890°W , 18.009°N . **b** Spatial distribution of λ_{LF} ,

quantifying the causal relationship strength from liquefaction to DPM signals. **c** Spatial distribution of λ_{LS} , quantifying the causal relationship strength from landslides to DPM signals. **d** Spatial distribution of λ_{BD} , quantifying the causal relationship strength from building damage to DPM signals.

These results demonstrate the ability of Spatial-VCBN to decouple different causal mechanisms contributing to observed surface changes, providing a more nuanced understanding of disaster impacts than would be possible from DPM interpretation alone. This capability forms the foundation for our subsequent analysis of how these causal parameters influence remote sensing observations.

Building on our understanding of causal relationships between hazards and structural damage, we now examine how these mechanisms present in DPM signals under varying conditions. We analyze both high-fidelity signal regions and challenging signal-constrained environments to demonstrate the robustness of Spatial-VCBN.

Robust spatially-varying causal inference between hazards and remote sensing observations

Figure 2 reveals the complex spatial relationships between different hazard mechanisms and their contributions to the observed DPM signals in an area with significant surface deformation in the study area with an extent of -66.946°W , 17.958°N to -66.890°W , 18.009°N . The DPM presented in Fig. 2a shows a concentrated area with high deformation values. The spatial distribution of the causal parameters provides insight into the primary mechanisms that drive these changes.

In this region, we observe spatial heterogeneity in how different hazards influence the DPM signals. The liquefaction causal parameter λ_{LF} (Fig. 2b) shows stronger values in the lower portion of the deformation area, suggesting that liquefaction processes are a dominant contributor to the DPM signals in that specific zone. This pattern is consistent with the typical geographic distribution of liquefaction in low-lying areas with specific soil conditions. Conversely, the landslide causal parameter λ_{LS} shown in Fig. 2c exhibits higher values in scattered patches, particularly in the right portion of the image where several high-intensity areas appear. This indicates locations where landslide mechanisms have a stronger influence on the observed DPM signals, likely corresponding to areas with steeper slopes or unstable terrain.

Figure 2d displays the building damage parameter λ_{BD} . It shows a distinct spatial pattern with moderate to high values in several discrete clusters. These areas represent locations where building deformation and structural impacts most strongly contribute to the DPM signals, which often correspond to zones with higher density of built structures. This spatial segregation of causal parameters demonstrates the ability of Spatial-VCBN to decouple multiple contributing factors to DPM signals, providing insights into the dominant hazard mechanisms at different locations within the affected area.

This capability to decouple in high-signal regions establishes a baseline for comparison with more challenging signal-constrained environments, which we examine next to demonstrate the robustness of Spatial-VCBN under varying conditions.

While Spatial-VCBN performs well in areas with strong DPM signals, real-world disaster assessment often involves regions with weak, noisy, or

ambiguous observational data. We now demonstrate the ability of Spatial-VCBN to extract meaningful causal parameters even in these challenging signal-constrained environments. Figure 3 highlights the robust capability of Spatial-VCBN to identify meaningful landslide patterns even in challenging conditions where DPM signals are weak, inconsistent, or contaminated by noise. The areas shown exhibit noisy DPM readings, shown in Fig. 3c, g that likely result from environmental factors such as snow cover or the complex reflectance properties of mountainous terrain, which can introduce artifacts in satellite-based change detection. Despite these challenging conditions, our spatially-aware causal Bayesian network successfully extracts coherent landslide susceptibility patterns. The distributions of landslide causal parameter λ_{LS} shown in Fig. 3a, e display distinct spatial organization with higher values concentrated along features that correspond to terrain characteristics associated with landslide risk. This demonstrates that Spatial-VCBN can effectively filter signal from noise by leveraging spatial correlation structures through its Gaussian Process component.

Notably, the uncertainty maps for landslide occurrence (Fig. 3c, h) reveal important patterns that complement our posterior probability estimates. Areas with high landslide probability in Fig. 3b, g generally exhibit low uncertainty (darker regions in the uncertainty maps), indicating high confidence in these estimations. Conversely, transitional zones between high and low probability areas show elevated uncertainty levels. This uncertainty quantification provides critical decision support information, as emergency responders can prioritize areas with both high hazard probability and low uncertainty for immediate intervention, while areas with high uncertainty might warrant additional monitoring or investigation.

The posterior probability maps for landslide occurrence presented in Fig. 3b, f reveal a refined understanding of landslide risk that transcends the limitations of the noisy DPM data. These posterior estimates incorporate both the learned causal strengths and the underlying geophysical context, resulting in coherent spatial patterns that align with landslide-prone landscape features visible in the satellite imagery (Fig. 3d, h). This ability to maintain signal fidelity in challenging conditions is particularly important for comprehensive hazard assessment in mountainous regions, where environmental factors typically confound traditional analysis methods.

Similar resilience is observed in building damage estimation, as shown in Fig. 4. Several factors can introduce noise in building damage estimation from satellite-based DPM, including variable building materials and construction types that respond differently to deformation, vegetation coverage partially obscuring buildings, complex urban geometries creating shadows and radar reflection artifacts, temporal variations in atmospheric conditions affecting satellite measurements, and pre-existing structural modifications unrelated to disaster impacts. Despite these challenges, our spatially-aware causal Bayesian network successfully extracts coherent building damage patterns. The distributions of building damage causal parameter λ_{BD} shown in Fig. 4a, e exhibit clear spatial structure with higher values concentrated along what appear to be developed corridors visible in the satellite imagery. These patterns follow the distribution of built environments rather than




Fig. 3 | Spatial pattern learning in areas with noisy DPM signals following the 2020 Puerto Rico earthquake. This figure demonstrates the ability of Spatial-VCBN to detect landslide patterns even in areas with weak or noisy DPM signals. Top row (a–e) and bottom row (f–j) show two different regions following the 2020 Puerto Rico earthquake: **a, f** Spatial distribution of λ_{LS} , quantifying the causal relationship

strength from landslides to DPM; **b, g** Posterior probability of landslide occurrence; **c, h** Uncertainty of landslide posterior distributions; **d, i** Damage Proxy Map (DPM) showing noisy surface deformation signals (0–1); **e, j** Google satellite imagery showing the corresponding terrain.




Fig. 4 | Building damage estimation in areas with noisy DPM signals following the 2020 Puerto Rico earthquake. This figure demonstrates the ability of Spatial-VCBN to detect building damage patterns despite weak DPM signals. Top row (a–e) and bottom row (f–j) show two different regions. **a, f** Spatial distribution of λ_{BD} ,

quantifying the causal relationship strength from building damage to DPM; **b, g** Posterior probability of building damage occurrence; **c, h** Uncertainty of building damage posterior; **d, i** DPM showing noisy surface deformation signals; **e, j** Google satellite imagery showing developed areas with building structures.

appearing randomly distributed, suggesting that Spatial-VCBN is able to capture true building-related signals despite the noise.

The uncertainty maps for building damage (Fig. 4c, h) show particularly low uncertainty values (typically below 0.2) in both developed and undeveloped areas, with slightly higher uncertainty in transition zones. This spatial consistency in uncertainty estimates further validates our model's robustness. The uncertainty is generally lower for building damage estimation compared to landslide estimation (maximum of 0.2 versus 0.25), suggesting that building damage signatures may be more distinctly identifiable in the DPM signals. This has important implications for disaster response prioritization, as resources can be more confidently allocated to areas with predicted building damage.

The posterior probability maps for building damage presented in Fig. 4b, g show improved spatial patterns that are more structured than what might be inferred from the noisy DPM alone. The larger posterior probabilities appear primarily in areas with visible building clusters in the satellite

imagery (Fig. 4d, h), demonstrating that Spatial-VCBN effectively incorporates contextual information about the built environment.

The robustness of Spatial-VCBN extends to different earthquake events and hazard types, as demonstrated in Fig. 5 for liquefaction estimation in the 2021 Haiti earthquake. Even when DPM signals are compromised by various confounding factors such as water level fluctuations, coastal erosion processes, wave action affecting shoreline appearance, soil moisture variations in near-shore environments, and varying sediment compositions that respond differently to seismic shaking, Spatial-VCBN successfully identifies coherent liquefaction patterns across different coastal areas. The distributions of liquefaction causal parameter λ_{LF} , which is presented in Fig. 5a, e, show distinctive spatial patterns with higher values concentrated in areas near coastlines where geological conditions favor liquefaction. The consistency of these patterns across different coastal regions within the same earthquake event suggests that Spatial-VCBN is able to capture fundamental physical relationships rather than location-specific anomalies.




Fig. 5 | Liquefaction estimation in coastal areas of the 2021 Haiti earthquake. Top row (a–e) and bottom row (f–j) show two different coastal areas. **a, f** Spatial distribution of λ_{LF} , quantifying the causal relationship strength from liquefaction to DPM; **b, g** Posterior probability of liquefaction occurrence; **c, h** Uncertainty of liquefaction posterior; **d, i** DPM showing noisy surface deformation signals; **e, j** Google satellite imagery showing coastal areas susceptible to liquefaction.

Table 1 | AUC values comparison across earthquake events

Model	Landslides			Liquefaction			Building Damage		
	HT	PR	TK	HT	PR	TK	HT	PR	TK
Spatial-VCBN	0.9703 ± 0.0124	0.9527 ± 0.0097	NGA	NGA	0.9489 ± 0.0108	NGA	0.9774 ± 0.0085	0.9532 ± 0.0091	0.9674 ± 0.0076
Prior Model ^{41,42,59}	0.9228	0.9158	NGA	NGA	0.8662	NGA	0.7755	0.7050	0.9051
VBCI ¹²	0.9413	0.9233	NGA	NGA	0.9097	NGA	0.9327	0.9145	0.9361
Bilateral filter ²⁰	0.9508	0.9413	NGA	NGA	0.9201	NGA	0.9406	0.9319	0.9486
ANN ⁴³	0.8428	0.8058	NGA	NGA	0.7897	NGA	–	–	–
GBM ⁴³	0.8819	0.8501	NGA	NGA	0.7454	NGA	–	–	–
Ensemble ⁴⁴	–	–	–	–	–	–	0.8619	0.8703	0.8555

Bold values represent the best performance for each column. Values for Spatial-VCBN are shown as mean ± standard deviation.

LS landslide, LF liquefaction, BD building damage, HT Haiti, PR Puerto Rico, TK Turkey, NGA no ground truth available.

The uncertainty maps for liquefaction (Fig. 5c, h) exhibit the lowest overall uncertainty values among the three hazard types (maximum of 0.15), particularly near shorelines where liquefaction is most likely to occur. This lower uncertainty for liquefaction estimation likely reflects the stronger geophysical constraints on where liquefaction can occur, primarily in saturated, unconsolidated sediments near water bodies. The clear spatial correlation between low uncertainty and coastal proximity offers valuable information for risk-aware decision making, as emergency responders can confidently prioritize resources for coastal communities where liquefaction risk is both high and certain.

The posterior probability maps for liquefaction, shown in Fig. 5b, f, exhibit refined spatial patterns that align with coastal geomorphology visible in the satellite imagery (Fig. 5d, h). Higher probabilities appear in low-lying coastal areas with likely unconsolidated sediments that are typically more susceptible to liquefaction during seismic events. This correspondence between model estimations and physical geography further validates the ability of Spatial-VCBN to capture meaningful hazard patterns despite noisy observations.

Our comprehensive uncertainty quantification extends beyond spatial visualizations to the performance metrics themselves, as shown in Table 1. By incorporating standard deviation values alongside mean AUC metrics, we provide a robust assessment of model reliability. Spatial-VCBN consistently achieves the highest AUC values with relatively small standard deviations (between ±0.0076 and ±0.0124), demonstrating both superior performance and reliability. The standard deviations tend to be smaller for building damage estimation compared to landslide estimation, which aligns with our spatial uncertainty observations. This performance consistency

across different hazard types and earthquake events, coupled with explicit uncertainty quantification, enhances the practical utility of our framework for real-world disaster response applications where decision-makers must consider both hazard likelihood and estimation confidence.

These qualitative assessments across different hazard types, signal conditions, and earthquake events demonstrate the robust nature of our spatially-aware causal framework. In the following section, we provide quantitative evaluation metrics that further validate these observations.

Figure 5 demonstrates the effectiveness of Spatial-VCBN in detecting liquefaction patterns across different coastal regions affected by the 2021 Haiti earthquake, even when DPM signals are compromised by various confounding factors such as water level fluctuations, coastal erosion processes, wave action affecting shoreline appearance, soil moisture variations in near-shore environments, and varying sediment compositions that respond differently to seismic shaking. Despite these challenges, our spatially-aware causal Bayesian network successfully identifies coherent liquefaction patterns across different coastal areas. The distributions of liquefaction causal parameter λ_{LF} , which is presented in Fig. 5a, e, show distinctive spatial patterns with higher values concentrated in areas near coastlines where geological conditions favor liquefaction. The consistency of these patterns across different coastal regions within the same earthquake event suggests that Spatial-VCBN is able to capture fundamental physical relationships rather than location-specific anomalies. The posterior probability maps for liquefaction, shown in Fig. 5b, f, exhibit refined spatial patterns that align with coastal geomorphology visible in the satellite imagery (Fig. 5d, h). Higher probabilities appear in low-lying coastal areas with likely unconsolidated sediments that are typically more susceptible to

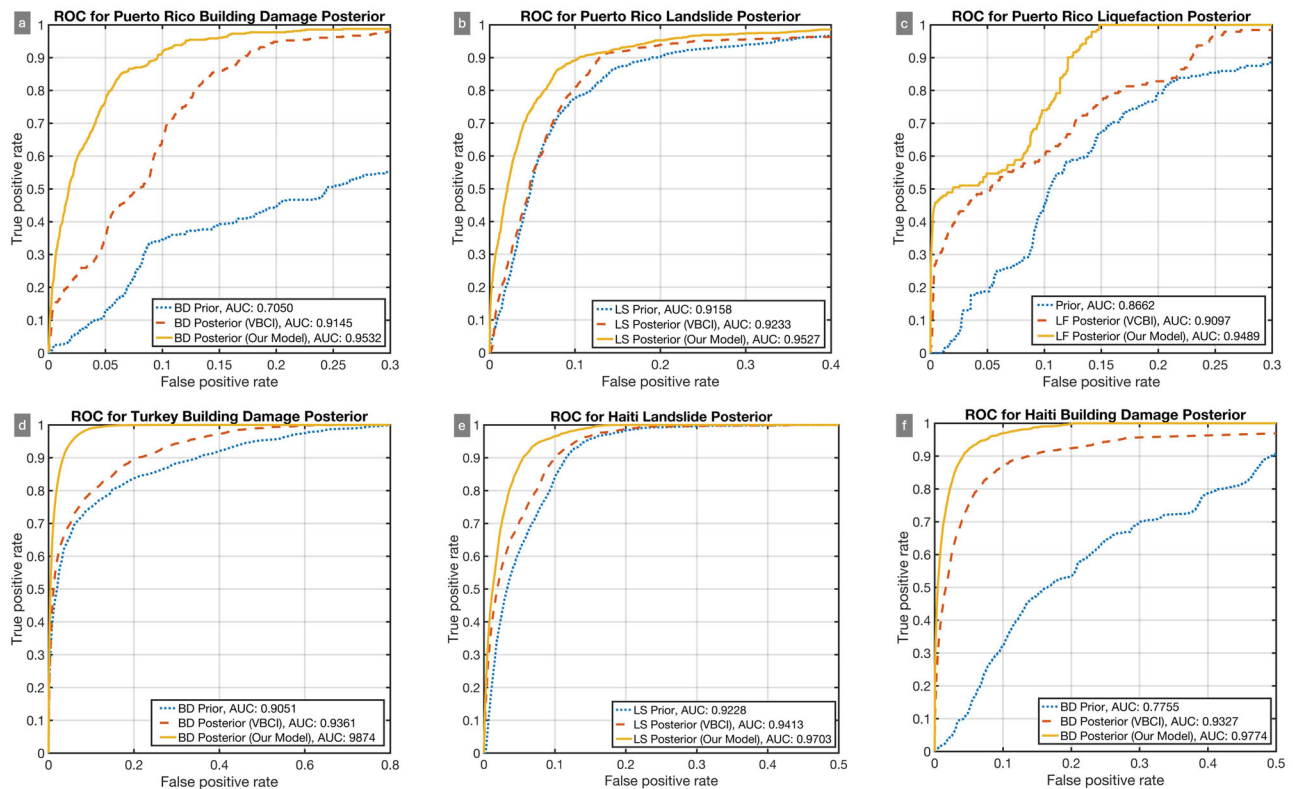


Fig. 6 | ROC (Receiver Operating Characteristic) curves and AUC (Area Under the Curve) values comparing the performance of Spatial-VCBN against baselines for different hazard types across three recent earthquakes. a–c Building damage, Landslide, and Liquefaction ROC curves of the 2020 Puerto Rico earthquake.

d Building damage for 2023 Turkey-Syria earthquake, **e, f** Building damage and Landslide ROC curves of the 2021 Haiti earthquake. Each plot compares three approaches: prior probability^{41,42} (blue dotted line), posterior from VBCI baseline¹² (red dashed line), and posterior from Spatial-VCBN (solid yellow line).

liquefaction during seismic events. This correspondence between model estimations and physical geography further validates the ability of Spatial-VCBN to capture meaningful hazard patterns despite noisy observations. The ability of Spatial-VCBN to produce consistent, physically plausible liquefaction assessments in challenging coastal environments, where traditional DPM analysis would be heavily compromised by noise, demonstrates its value for comprehensive multi-hazard assessment in diverse settings.

Cross-event validation and baseline comparison

Our evaluation framework includes a comprehensive array of comparative methods across multiple hazard types. For our landslide (LS) and liquefaction (LF) estimation performance, we evaluate against several established approaches: the ground failure models developed by the USGS^{41,42}, the VBCI methodology proposed by¹², as well as both the Artificial Neural Network (ANN) and Gradient Boosting Machine (GBM) techniques outlined by⁴³. In assessing building damage (BD) estimation capabilities, we contrast Spatial-VCBN with conventional building fragility curves, the VBCI framework¹², ensemble methodology⁴⁴, and the bilateral filtering approach recently introduced by²⁰. Figure 6 and Tables 1 and 2 demonstrate the superior performance of our spatially-aware causal Bayesian network across multiple hazard types and earthquake events. The consistently higher metric values achieved by Spatial-VCBN validate the effectiveness of Spatial-VCBN in accurately identifying disaster impacts.

For the 2020 Puerto Rico earthquake, Spatial-VCBN achieves remarkable improvements in building damage (Fig. 6a), with an AUC of 0.9532 compared to 0.9145 for the VBCI and only 0.7050 for the prior. This represents a substantial 35.2% improvement over the prior probability baseline. As shown in Table 1, Spatial-VCBN consistently outperforms six competitive baselines, including both traditional statistical approaches and deep learning methods. The F1 score results shown in Table 2 further

validate these findings, with Spatial-VCBN achieving 0.9167 for building damage in Puerto Rico, compared to the next best performer (Bilateral filter²⁰) at 0.8945. This performance gain is particularly noteworthy for building damage, where distinguishing between different damage mechanisms can be challenging. For landslides (Fig. 6b), Spatial-VCBN also outperforms competitors with an AUC of 0.9527, though the margin is smaller as all methods perform relatively well on this hazard type. Looking at Table 1, we observe that the closest competitor for landslide estimation is the Bilateral filter approach (0.9413), which still falls short of the performance of our method by 1.2%. The consistency between AUC and F1 metrics (0.9173 for Spatial-VCBN versus 0.9031 for Bilateral filter) indicates that our performance improvements are robust across different evaluation criteria. Liquefaction estimation (Fig. 6c) shows intermediate improvements, with Spatial-VCBN achieving an AUC of 0.9489, representing a 9.5% improvement over the prior (0.8662). Table 2 shows a similar pattern for F1 scores, with Spatial-VCBN achieving 0.9107 compared to 0.8815 for Bilateral filter and 0.8714 for VBCI, highlighting the balanced performance of Spatial-VCBN in both estimation accuracy and false alarm reduction.

The cross-event evaluation further substantiates the robustness of Spatial-VCBN. For the 2023 Turkey-Syria earthquake (Fig. 6d), Spatial-VCBN achieves an impressive AUC of 0.9874 for building damage, significantly outperforming the prior (0.9051) with improvement of 9.1%. Tables 1 and 2 reveal that this superiority is maintained across metrics, with our F1 score of 0.9412 substantially outperforming all alternatives including the Bilateral filter (0.9089). Notably, Spatial-VCBN demonstrates particularly strong improvements over learning-based methods such as the Ensemble approach, which achieves only F1 score of 0.8213 for the Turkey-Syria earthquake, highlighting the advantage of our causal framework over pure statistical learning approaches. Similar improvements are observed for the 2021 Haiti earthquake, with Spatial-VCBN achieving AUCs of 0.9703 for landslide (Fig. 6e) and 0.9774 for building damage (Fig. 6f), representing

Table 2 | F1 score comparison across earthquake events

Model	Landslides			Liquefaction			Building Damage		
	HT	PR	TK	HT	PR	TK	HT	PR	TK
Spatial-VCBN	0.9297	0.9173	NGA	NGA	0.9107	NGA	0.9325	0.9167	0.9412
Prior Model	0.8845	0.8797	NGA	NGA	0.8301	NGA	0.7423	0.6756	0.8678
VBCI	0.9015	0.8864	NGA	NGA	0.8714	NGA	0.8947	0.8778	0.8973
Bilateral filter	0.9102	0.9031	NGA	NGA	0.8815	NGA	0.9023	0.8945	0.9089
ANN	0.8102	0.7732	NGA	NGA	0.7582	NGA	–	–	–
GBM	0.8465	0.8160	NGA	NGA	0.7149	NGA	–	–	–
Ensemble	–	–	–	–	–	–	0.8274	0.8356	0.8213

Bold values represent the best performance for each column.

LS landslide, LF liquefaction, BD building damage, HT Haiti, PR Puerto Rico, TK Turkey, NGA no ground truth available.

improvements of 26.0% over the prior (0.7755) for building damage, and 5.15% over the prior (0.9228) for landslide. Table 1 shows that even compared to the stronger Bilateral filter approach, Spatial-VCBN maintains advantages of 2.0% and 3.9% for landslide and building damage estimation, respectively.

A comprehensive examination of Tables 1 and 2 reveals several additional insights not immediately apparent from the ROC curves alone. First, the performance gap between Spatial-VCBN and alternatives is consistently larger for building damage than for landslide estimation across all earthquake events, suggesting that our spatial-causal modeling approach particularly excels at capturing the complex mechanisms underlying building vulnerability. Second, machine learning methods like ANN and GBM show substantially worse performance than causal modeling approaches across all metrics, with AUC differences of up to 0.1475 (comparing Spatial-VCBN to ANN for landslide estimation in Puerto Rico) and F1 score differences of up to 0.2041 (comparing Spatial-VCBN to GBM for liquefaction estimation in Puerto Rico). This highlights the limitations of purely data-driven approaches when physical causal mechanisms are not explicitly modeled.

These results highlight three key strengths of Spatial-VCBN. First, Spatial-VCBN demonstrates consistent performance advantages across different hazard types, with particularly notable improvements for building damage estimation where distinguishing between different damage mechanisms is especially challenging. Second, Spatial-VCBN shows strong transferability across different earthquake events in diverse geographical settings, maintaining its performance edge in the Puerto Rico, Haiti, and Turkey-Syria earthquakes despite their varying geological contexts and built environment characteristics. Third, Spatial-VCBN achieves robust performance improvements particularly at low FPRs, which is crucial for practical deployment in disaster response scenarios where false alarms can waste limited resources and undermine confidence in automated hazard assessments. The consistent outperformance across multiple events, hazard types, and evaluation metrics demonstrates that modeling spatially-varying causal relationships through combined Gaussian Processes and normalizing flows provides a powerful framework for disaster impact assessment that generalizes well across different scenarios. This performance advantage is from the ability of Spatial-VCBN to capture both the spatial correlation structure inherent in disaster impacts and the complex, non-Gaussian distributions of causal effects that characterize real-world hazard-damage relationships. By learning location-specific causal parameters rather than assuming uniform relationships across the study area, Spatial-VCBN can adapt to the unique geological, structural, and environmental factors that influence how different hazards manifest in different locations.

Ablation study and model robustness

Computational efficiency analysis. Table 3 provides insights into the computational requirements of our framework across three earthquake events with varying affected areas. We implement a pruning strategy that maintains efficiency by focusing computational resources on active nodes

Table 3 | This table shows the time cost of running our framework using the same batch size in three earthquake events using real-world data

Method	Haiti EQ.	Puerto Rico EQ.	Turkey-Syria EQ.
Map size	15,970 km ²	1305 km ²	4676 km ²
Time(s)	15,029	1198	4543

while eliminating inactive ones from the processing pipeline. For the Haiti earthquake with the largest affected area (15,970 km²), Spatial-VCBN required 15,029 s (~4.2 h), while the smaller Puerto Rico earthquake area (1305 km²) was processed in just 1198 s (~20 min). This near-linear scaling (with a processing rate of ~0.94 s/km²) indicates that our framework maintains computational efficiency as the geographical scope increases.

The consistency in processing time per unit area across different earthquake events suggests that the computational demands of Spatial-VCBN are primarily determined by the spatial extent rather than being affected by the complexity or specific characteristics of different regions. This predictable scaling behavior is particularly valuable for emergency response scenarios where estimation of required computational resources is crucial for timely deployment. Our implementation leverages GPU acceleration, equipped with an NVIDIA Tesla T4 GPU (15 GB VRAM) and 51 GB of RAM, enabling rapid parallel processing of spatial data. The results indicate that even for large-scale events like the Turkey-Syria earthquake (4676 km²), analysis can be completed within reasonable timeframes (~1.3 h), making our approach practical for operational use in disaster response.

Hyperparameter sensitivity. Table 4 presents a systematic evaluation of the sensitivity of Spatial-VCBN to the flow number K in the normalizing flow component, which directly affects the expressiveness of the spatially-varying causal parameter distributions. The results reveal several important patterns across different earthquake events and hazard types. For all hazard types and earthquake events, performance improves substantially as K increases from 2 to 6, with the significant improvements observed in the early stages. For example, in the Haiti earthquake building damage estimation, AUC increases from 0.8915 with $K = 2$ –0.9774 with $K = 6$, an improvement of 9.6%. This pattern indicates that the spatial distribution of causal parameters in disaster scenarios exhibits complexity that cannot be adequately captured by simpler flow architectures.

Notably, we observe that $K = 6$ represents an optimal balance point across all hazards and events, achieving the highest AUC values for landslide estimation in Haiti (0.9703), building damage in Haiti (0.9774), landslide in Puerto Rico (0.9527), liquefaction in Puerto Rico (0.9489), building damage in Puerto Rico (0.9532), and building damage in Turkey-Syria (0.9874). Beyond $K = 6$, we observe a slight performance degradation, with AUC

Table 4 | Ablation study on flow number K in normalizing flows

K	Haiti EQ.			Puerto Rico EQ.			Turkey-Syria EQ.		
	AUC_{LS}	AUC_{LF}	AUC_{BD}	AUC_{LS}	AUC_{LF}	AUC_{BD}	AUC_{LS}	AUC_{LF}	AUC_{BD}
$K = 2$	0.9239	NGA	0.8915	0.9013	0.8896	0.8714	NGA	NGA	0.9089
$K = 3$	0.9425	NGA	0.9178	0.9187	0.9057	0.8982	NGA	NGA	0.9302
$K = 4$	0.9568	NGA	0.9452	0.9311	0.9257	0.9214	NGA	NGA	0.9587
$K = 5$	0.9638	NGA	0.9641	0.9478	0.9375	0.9422	NGA	NGA	0.9755
$K = 6$	0.9703	NGA	0.9774	0.9527	0.9489	0.9532	NGA	NGA	0.9874
$K = 7$	0.9701	NGA	0.9772	0.9525	0.9486	0.9530	NGA	NGA	0.9871
$K = 9$	0.9695	NGA	0.9767	0.9522	0.9481	0.9527	NGA	NGA	0.9869

NGA means no ground truth available. Bold values represent the best performance.

Table 5 | Comparison of normalizing flow architectures

Flow Architecture	Expressiveness	Memory	AUC (Haiti LS)	AUC (PR BD)	Running Time
Planar Flows (Current)	Moderate	1.0×	0.9703	0.9532	1.0×
RealNVP	High	2.8×	0.9712	0.9547	1.8×
MAF	Very High	4.7×	0.9719	0.9556	3.2×
Radial Flows	Moderate	1.2×	0.9681	0.9511	1.1×
IAF	High	3.2×	0.9714	0.9550	2.5×

Memory and Running Time values are reported relative to Planar Flows (1×). All architectures were evaluated using the same dataset and hardware configuration. AUC values represent area under the ROC curve for landslide detection in Haiti (LS) and building damage detection in Puerto Rico (BD).

decreasing marginally for $K = 7$ and $K = 9$. This pattern suggests a form of overfitting when the flow becomes too expressive relative to the available constraints.

The impact of K varies across different hazard types, with building damage showing the highest sensitivity to flow complexity. For the Puerto Rico earthquake, increasing K from 2 to 6 improves building damage AUC by 0.0818, compared to improvements of 0.0514 for landslide and 0.0593 for liquefaction. This difference likely reflects the more complex causal mechanisms involved in building damage, which depends on both the primary hazards and structural characteristics. Cross-event comparison reveals consistent patterns in hyperparameter sensitivity, with all three earthquake events showing similar optimal flow complexity despite their geographical and geological differences. This consistency suggests that Spatial-VCBN architecture captures fundamental aspects of the spatially-varying causal relationships in disaster scenarios rather than simply fitting to dataset-specific patterns.

Normalizing flow architecture evaluation. To justify our choice of planar flows in the Spatial-VCBN framework, we conducted a comparative analysis of several normalizing flow architectures, evaluating their performance, computational efficiency, and memory requirements. Table 5 presents the results of this comparison across five different flow architectures: Planar Flows (our implemented approach), RealNVP, Masked Autoregressive Flows (MAF), Radial Flows, and Inverse Autoregressive Flows (IAF).

While more complex architectures such as MAF demonstrate slightly higher performance with AUC improvements of 0.0016 (0.17%) for Haiti landslide detection and 0.0024 (0.25%) for Puerto Rico building damage detection, these marginal gains come at substantial computational costs. MAF requires 4.7× the memory of planar flows and 3.2× longer running time. Similarly, RealNVP and IAF show modest performance improvements (less than 0.2% AUC increase) while requiring significantly more computational resources. For our spatially-aware causal modeling

framework, which often needs to process large geographical regions, computational efficiency is a critical consideration. The landslide detection area for the Haiti earthquake alone covers ~15,970 km², making efficient processing essential for practical deployment in time-sensitive disaster response scenarios. Planar flows provide an optimal balance between model expressiveness and computational efficiency, enabling our framework to handle large-scale disaster events while maintaining high accuracy ($AUC \geq 0.95$ across all test cases). Radial flows, despite their comparable computational efficiency to planar flows, demonstrated slightly lower performance in our experiments. This suggests that the specific transformation properties of planar flow, which are well-suited for capturing the spatial causal relationships in earthquake-induced hazard scenarios. Based on this analysis, we selected planar flows for our final implementation, prioritizing practical application while maintaining performance for multi-hazard estimation.

Discussion

The development of a spatially-aware causal Bayesian network with normalizing flows represents a significant advancement in multi-hazard disaster impact assessment. Our results demonstrate that this approach not only improves estimation accuracy across multiple hazards and impacts induced by different earthquake events, with AUC improvements of up to 35.2% over prior probability baselines and 5.5% over state-of-the-art VBCI methods. Additionally, it provides interpretable insights into the complex causal mechanisms underlying disaster impacts. In this discussion, we contextualize our findings within the broader disaster science literature and examine their implications for both theoretical understanding and practical applications.

Traditional approaches to hazard assessment often rely on either purely data-driven methods that lack causal interpretability or physical models that struggle to incorporate the complex spatial heterogeneity of real-world disaster contexts. Our framework bridges this gap by explicitly modeling spatially-varying causal relationships through a combination of Gaussian Processes and normalizing flows. This innovation allows us to capture both the spatial correlation structure inherent in disaster impacts and the non-Gaussian, potentially multimodal distributions of causal effects that characterize real-world hazard-damage relationships.

The spatial patterns observed in our causal parameters (γ_{LS} , γ_{LF} , λ_{LS} , λ_{LF} , and λ_{BD}) highlight the fundamental importance of accounting for spatial heterogeneity in disaster modeling. For instance, in the Puerto Rico earthquake (Fig. 1), the clear differentiation between landslide-dominated hillsides and liquefaction-dominated lowlands demonstrates how terrain characteristics fundamentally alter causal mechanisms. This spatial variability, along with the distinct patterns of building damage influence, suggests that assuming spatial homogeneity in causal relationships, as many existing models do, may lead to substantial inaccuracies in hazard assessment. The optimal flow number of $K = 6$ identified in our ablation study (Table 4) suggests that real-world causal relationships in disaster contexts exhibit complexity that cannot be adequately captured by simpler distributional

assumptions, underscoring the importance of our methodological innovation.

Remote sensing plays a crucial role in disaster response, but the interpretation of complex signals like DPM remains challenging. The ability of Spatial-VCBN to decouple multiple contributing factors to DPM signals addresses a significant gap in existing approaches. As demonstrated in Fig. 2, our framework can identify distinct causal pathways from different hazard types to observed DPM signals within the same geographic region, providing critical insights that would be lost in approaches that treat DPM as a uniform indicator of damage. The results demonstrating robust causal parameter extraction even in signal-constrained environments (Figs. 3 and 4) highlight the potential of causal modeling to enhance the utility of remote sensing data in challenging conditions. For example, in mountainous regions where snow cover and complex topography introduce noise into DPM measurements, Spatial-VCBN successfully extracts coherent landslide susceptibility patterns that align with terrain characteristics visible in satellite imagery. The consistent performance across different hazard types and earthquake events suggests that our approach captures fundamental physical relationships rather than only fitting to dataset-specific patterns. This is particularly notable given the diverse geological and built environment contexts represented in our study areas (Puerto Rico, Haiti, and Turkey-Syria), as evidenced by the comparable AUC values achieved across these events (Table 1). The transferability of Spatial-VCBN suggests that the spatially-varying causal relationships we identify reflect genuine physical processes that generalize across different disaster scenarios.

The practical implications of our work extend beyond theoretical advancements to offer tangible benefits for disaster risk reduction and emergency response. Particularly noteworthy is the strong performance of Spatial-VCBN at low FPR, as shown in the ROC curves (Fig. 6). In practical disaster response, false alarms can waste limited resources and undermine stakeholder confidence in automated assessment systems. The ability of our approach to maintain high detection rates while minimizing false positives addresses this crucial operational concern. The reasonable computational requirements demonstrated in our efficiency analysis (Table 3), with processing times of ~ 0.94 s/km² with GPU acceleration, suggest that our approach is viable for operational deployment, even for large-scale events. The near-linear scaling with geographical area provides predictable resource requirements for emergency management agencies planning post-disaster assessments. For example, our results indicate that even a large-scale event like the Haiti earthquake (15,970 km²) can be analyzed in ~ 4.2 h using modest computational resources. Beyond immediate response applications, our spatially-explicit causal framework offers valuable insights for long-term disaster risk reduction planning. The identification of areas where specific hazard mechanisms dominate could inform targeted mitigation strategies. For example, authorities might prioritize slope stabilization in regions with high γ_{LS} values or focus on liquefaction-resistant foundation designs in areas with elevated γ_{LF} parameters.

While we focused on seismic events in this study, the fundamental design of Spatial-VCBN makes it readily adaptable to non-seismic hazards and multi-hazard scenarios. The causal Bayesian network architecture is inherently flexible and can be reconfigured to model alternative causal chains, such as those involved in rainfall-induced landslides, flooding, or hurricane impacts. For instance, in modeling rainfall-triggered landslides, precipitation intensity could replace seismic shaking as the primary triggering node, while retaining similar geophysical feature nodes. The key requirement for applying our framework to other hazard types is the availability of (1) relevant remotely sensed coherence change products similar to DPMs, (2) geospatial features that influence the hazard mechanisms, and (3) a conceptual understanding of the causal relationships between triggers, environmental conditions, and resulting impacts. Our spatially-varying approach using Gaussian Processes with normalizing flows would be particularly valuable for modeling non-seismic hazards like hurricanes or floods, where impacts often show strong spatial heterogeneity based on topography, infrastructure quality, and socioeconomic factors. Extending Spatial-VCBN to these scenarios would primarily require

adapting the causal graph structure and incorporating hazard-specific geospatial features, while the core inference methodology would remain applicable.

The spatially-varying causal parameters identified by our framework could serve as valuable calibration or validation data for detailed physical simulations, potentially improving their spatial accuracy. Conversely, insights from physical models could inform prior distributions in our Bayesian framework, creating a virtuous cycle of model improvement. This integration of data-driven causal inference with physical understanding represents a promising direction for advancing multi-hazard assessment methodologies. The ability of Spatial-VCBN to identify meaningful patterns even in areas with weak or noisy DPM signals suggests that it could extend the utility of physics-based models to regions where observational data is limited or compromised by environmental factors. This is particularly relevant for global-scale hazard assessment, where data quality varies substantially across different regions.

While our approach complements rather than replaces physics-based hazard models, the identified spatially-varying causal parameters could serve as valuable calibration data for physical simulations. Conversely, insights from physical models could inform prior distributions in our Bayesian framework. Despite promising results, limitations remain. Future work should validate our approach across a broader range of disaster types and incorporate temporal dynamics of hazard evolution.

In conclusion, our spatially-aware causal Bayesian network advances disaster impact assessment through its ability to capture spatially heterogeneous causal relationships, maintain robust performance under varying signal conditions, and provide interpretable insights into complex disaster mechanisms. The framework addresses critical limitations in existing approaches that either lack causal interpretability^{43,44} or struggle with spatial heterogeneity^{3,12}. Our integration of Gaussian Processes with normalizing flows builds upon advances in spatially-varying coefficient models⁴⁵ while extending their applicability to non-Gaussian posterior distributions. The performance improvements demonstrated across multiple hazard types and earthquake events, particularly the substantial gains in building damage assessment, validate the practical utility of our approach for disaster response operations.

Methods

Prior geospatial models

Earthquake shaking intensity serves as the fundamental catalyst for structural damage and geohazards like landslides and liquefaction. We incorporate this critical factor through USGS ShakeMap data, which provides comprehensive shaking metrics. The composite ShakeMap effectively captures maximum ground motion intensities, enabling more accurate assessments of structural and geologic impacts, which is an invaluable resource for post-earthquake forensic analysis.

Beyond basic shaking information, USGS expanded their rapid post-earthquake information products in 2018 to include ground-failure estimations⁴⁶. These specialized models predict the likelihood and distribution of earthquake-triggered landslides and liquefaction³⁰. The modeling framework integrates multiple environmental factors including terrain slope, geological susceptibility, and soil characteristics to generate probability maps^{41,42}. An illustrative example appears in Fig. 7, showing the USGS ground failure projections following the 2023 Turkey-Syria earthquake. Within our Bayesian network architecture (Fig. 9), these established USGS ground failure products are incorporated as prior probability distributions for landslide and liquefaction occurrence, represented by nodes α_{LS} and α_{LF} respectively.

InSAR data and damage proxy maps

For our analysis, we utilized DPMs, which are specialized remote sensing products derived from InSAR data. DPMs identify ground surface changes by analyzing coherence differences between radar signals captured before and after seismic events⁹. These maps were generated by NASA's ARIA team following three major earthquakes: the 2020 Puerto Rico earthquake

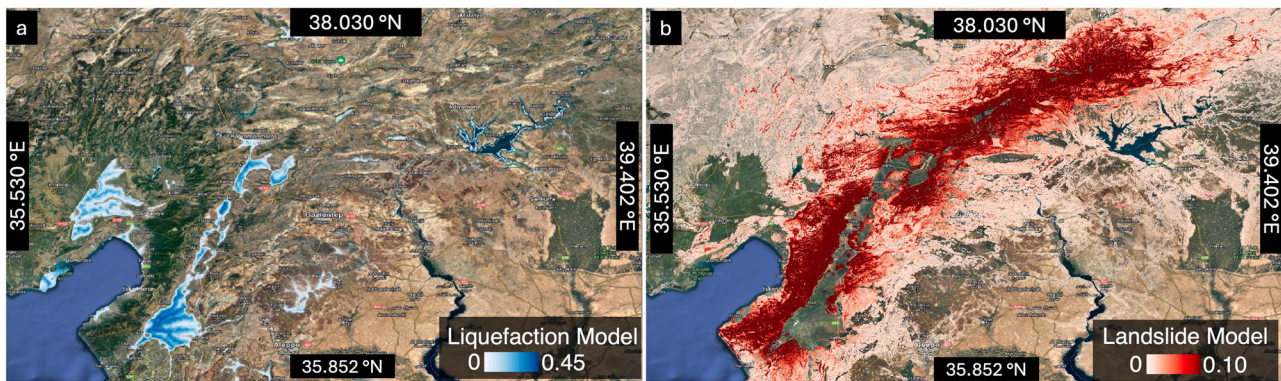


Fig. 7 | Ground Failure Probability Maps Generated by the USGS. After the 2023 Turkey-Syria earthquake sequence, the USGS produced example ground failure models for liquefaction (shown in panel b) and landslide (shown in panel a)⁶⁰. The probability of ground failure models is what the legend colors represent.




Fig. 8 | Damage proxy maps (DPM) generated by the ARIA team following major earthquakes. a 2021 Haiti earthquake (M7.2) with red boxes highlighting areas of significant surface deformation; b 2020 Puerto Rico earthquake (M6.4) with red boxes indicating concentrated damage zones; c 2023 Turkey-Syria earthquake

sequence (M7.8). Brighter areas in all maps indicate greater surface deformation detected by InSAR. The geographical extents of these DPMs define our study regions.

(M6.4), the 2021 Haiti earthquake (M7.2), and the 2023 Turkey-Syria earthquake sequence (M7.8).

This InSAR product works by measuring changes in radar coherence, which represents the consistency of radar reflections from Earth's surface between two time periods. Earthquake-induced changes to the surface disrupt this coherence, allowing detection of surface deformation. The resulting DPMs, shown in Fig. 8, represent these coherence changes where brighter pixels indicate greater surface disruption potentially caused by landslides, liquefaction, or building damage. While DPMs provide valuable high-resolution (30-m pixel) observations, they present interpretation challenges as they cannot inherently distinguish between different hazard types and contain environmental noise. In this work, the Spatial-VCBN framework addresses these limitations.

Our study regions include (a) from -74.771°W to -72.581°W , from 18.118°N to 18.719°N for the 2021 Haiti earthquake (Fig. 8a); (b) from -67.022°W to -66.570°W , from 17.908°N to 18.153°N for the 2020 Puerto Rico earthquake (Fig. 8b); (c) from 36.681°E to 37.467°E , from 37.075°N to 37.676°N for the 2023 Turkey-Syria earthquake sequence (Fig. 8a).

Geophysical features

We incorporate seven geophysical features in Spatial-VCBN: (1) *Vs30*, (2) *Slope from DEM*, (3) *Land Cover*, (4) *DEM (Elevation)*, (5) *CTI (Compound Topographic Index)*, and (6) *Water Body Distance*, and (7) *Lithology*. Below, we provide a brief rationale for including each feature.

- **Vs30 (Shear-Wave Velocity over 30 m):** Vs30 is widely recognized as a key parameter for characterizing local site conditions and capturing

near-surface amplification effects in seismic hazard assessments. Higher Vs30 values generally correspond to stiffer soils or rocks, reducing the likelihood of significant ground amplification^{47,48}.

- **Slope from DEM:** Slope, derived from a Digital Elevation Model (DEM), is crucial in identifying areas susceptible to landslides and other mass-movement hazards triggered by seismic activity. Regions with steep slopes tend to experience higher landslide hazard potential^{49,50}.
- **Land Cover:** Land cover information helps discern the presence of vegetation, urban infrastructure, or water bodies. Different land cover classes (e.g., dense vegetation vs. built-up areas) can significantly influence ground response and post-earthquake soil stability⁵¹.
- **DEM (Elevation):** The base elevation data from DEM are essential for capturing broader topographic variations. Elevation can affect not only slope angle but also local climate, drainage patterns, and material properties, all of which have implications for earthquake-induced geophysical hazards⁵².
- **CTI (Compound Topographic Index):** CTI, sometimes referred to as the Topographic Wetness Index, quantifies the potential for water accumulation in the terrain. High CTI values often indicate persistently wetter soils, influencing both liquefaction susceptibility and potential post-seismic runoff or slope failure⁵³.
- **Water Body Distance:** Proximity to the Water Body plays a significant role in liquefaction evaluations because saturated coastal regions and near-shore sediments are more prone to ground failure under strong seismic shaking^{41,54}.
- **Lithology:** Lithology plays a crucial role in assessing earthquake-related hazards as it directly influences seismic wave propagation, amplification, and ground shaking intensity. Different rock types exhibit varying mechanical properties such as stiffness, strength, and porosity, which affect the degree of energy dissipation and resonance during an earthquake. Additionally, lithological features control the susceptibility of slopes to landslides and liquefaction in earthquake-prone areas. Understanding lithology is therefore essential for accurate hazard zoning and designing resilient infrastructure in seismically active regions^{42,55}.

These geophysical features collectively represent critical factors governing seismic hazard distributions, spanning soil stiffness, topography, hydrology, and land use. By explicitly modeling their spatial variations and interactions, we aim to more accurately capture the localized nature of seismic impacts across heterogeneous terrains.

Causal modeling for disaster impacts

Disasters propagate through complex chains of causation, where initial events trigger cascading hazards that ultimately result in observable impacts. To accurately model these processes, we develop a spatially-aware causal Bayesian network that captures both the structural relationships between variables and the spatial heterogeneity of these relationships. Figure 9 presents our causal Bayesian network architecture, which consists of three key component types: (1) observable variables (represented as orange rectangles) including geospatial features (GF) and DPMs; (2) latent hazard/impact variables (shown as blue circles) representing unobserved landslides (LS), liquefaction (LF), and building damage (BD); and (3) spatially-varying causal coefficients that quantify the strength of causal relationships at each location.

For each location l in our study area, we define the leaf node y^l as the damage proxy map observation. Its relationship with parent nodes follows a log-linear model:

$$\log y^l | \mathcal{P}(y^l), \epsilon_y^l, \lambda_{\mathcal{P}(y^l)}^l \sim N \left(w_{0y} + \sum_{k \in \mathcal{P}(y^l)} \lambda_k^l x_k^l, w_{\epsilon_y}^2 \right) \quad (1)$$

where x_k^l represents the hidden hazard/damage variables, $\lambda_{\mathcal{P}(y^l)}^l$ denotes the spatially-varying causal coefficients that quantify the causal relationship

from latent variables $x_k \in \{LS, LF, BD\}$ to y , and ϵ_y accounts for environmental noise. The hidden nodes $x_i^l \in \{0, 1\}$ for $i \in \{LS, LF, BD\}$ are binary variables indicating hazard occurrence, with activation probabilities governed by their respective parent nodes and spatially-varying coefficients γ_i^l that quantify the causal relationship from the parent nodes $\mathcal{P}(x)$ to x :

$$\begin{aligned} \log \frac{p(x_i^l = 1 | x_{\mathcal{P}(x_i^l)}^l, \epsilon_i^l, \gamma_{\mathcal{P}(x_i^l)}^l)}{1 - p(x_i^l = 1 | x_{\mathcal{P}(x_i^l)}^l, \epsilon_i^l, \gamma_{\mathcal{P}(x_i^l)}^l)} &= \sum_{k \in \mathcal{P}(x_i^l)} \gamma_k^l x_k^l + w_{\epsilon_i} \epsilon_i + w_{0i} \\ p(x_i^l | x_{\mathcal{P}(x_i^l)}^l, \epsilon_i^l, \gamma_{\mathcal{P}(x_i^l)}^l) &= [p(x_i^l = 1 | x_{\mathcal{P}(x_i^l)}^l, \epsilon_i^l, \gamma_{\mathcal{P}(x_i^l)}^l)]^{x_i^l} \\ &\times [1 - p(x_i^l = 1 | x_{\mathcal{P}(x_i^l)}^l, \epsilon_i^l, \gamma_{\mathcal{P}(x_i^l)}^l)]^{1-x_i^l} \end{aligned} \quad (2)$$

Modeling spatially heterogeneous causal effects using normalizing flows with Gaussian process

A key innovation in Spatial-VCBN is the modeling of spatially-varying causal coefficients λ^l and γ^l using Normalizing Flows with Gaussian Process. This combination allows us to capture both the spatial correlation structure inherent in these coefficients while accommodating potentially complex, non-Gaussian posterior distributions. The following sections detail our methodology for posterior inference in this model and demonstrate its effectiveness for multi-hazard impact estimation.

To capture the complex spatial variation in causal coefficients $v^l \in \{\lambda_a^l, \gamma_b^l\}$ where $a \in \{LS, LF, BD\}$ and $b \in \{\alpha_{LS}, \alpha_{LF}, LS, LF\}$, we propose a novel approach that combines Gaussian Processes (GPs) with normalizing flows. This approach effectively captures both the spatial correlation structure and the complex, potentially non-Gaussian distributions of causal effects.

We introduce latent spatial variables z_v^l which serve as the building blocks for our spatially varying causal coefficients. These latent variables represent underlying spatial patterns that, after transformation, will yield the causal parameters used in Spatial-VCBN. By working with these latent variables rather than directly modeling the coefficients, we gain mathematical tractability while maintaining expressiveness.

We model these latent variables z_v^l across different locations using a Gaussian Process prior, which naturally captures correlations between locations with similar geophysical characteristics:

$$\mathbf{z}_v \sim \mathcal{GP}(m_v(\mathbf{GF}), k_v(\mathbf{GF}, \mathbf{GF}')) \quad (3)$$

where $\mathbf{z}_v = [z_v^1, z_v^2, \dots, z_v^L]$ is the vector of latent variables across all locations, $m_v(\mathbf{GF})$ is the mean function that depends on the geophysical features \mathbf{GF} at all locations, and $k_v(\mathbf{GF}, \mathbf{GF}')$ is the kernel function that defines the covariance structure between different geophysical features. We use a Matern kernel, which is a widely used covariance function in spatial statistics whose smoothness parameter ν directly controls the mean-square differentiability of the process⁵⁶⁻⁵⁸, with $\nu = 3/2$:

$$k_v(\mathbf{GF}_i, \mathbf{GF}_j) = \sigma_v^2 \frac{2^{1-\nu}}{\Gamma(\nu)} \left(\sqrt{2\nu} \frac{\|\mathbf{GF}_i - \mathbf{GF}_j\|}{\ell_v} \right)^\nu K_\nu \left(\sqrt{2\nu} \frac{\|\mathbf{GF}_i - \mathbf{GF}_j\|}{\ell_v} \right) \quad (4)$$

where σ_v^2 controls the variance, ℓ_v is the length scale parameter, K_ν is the modified Bessel function, and $\mathbf{GF}_i, \mathbf{GF}_j$ represent the geophysical feature vectors at locations i and j . This formulation ensures that locations with similar geophysical features will have similar coefficients, which better captures the underlying physical relationships in Spatial-VCBN.

To allow for complex non-Gaussian distributions of causal coefficients, we transform the GP latent variables through a series of invertible normalizing flow transformations:

$$v^l = f_{K_v} \circ \dots \circ f_2 \circ f_1(z_v^l) \quad (5)$$

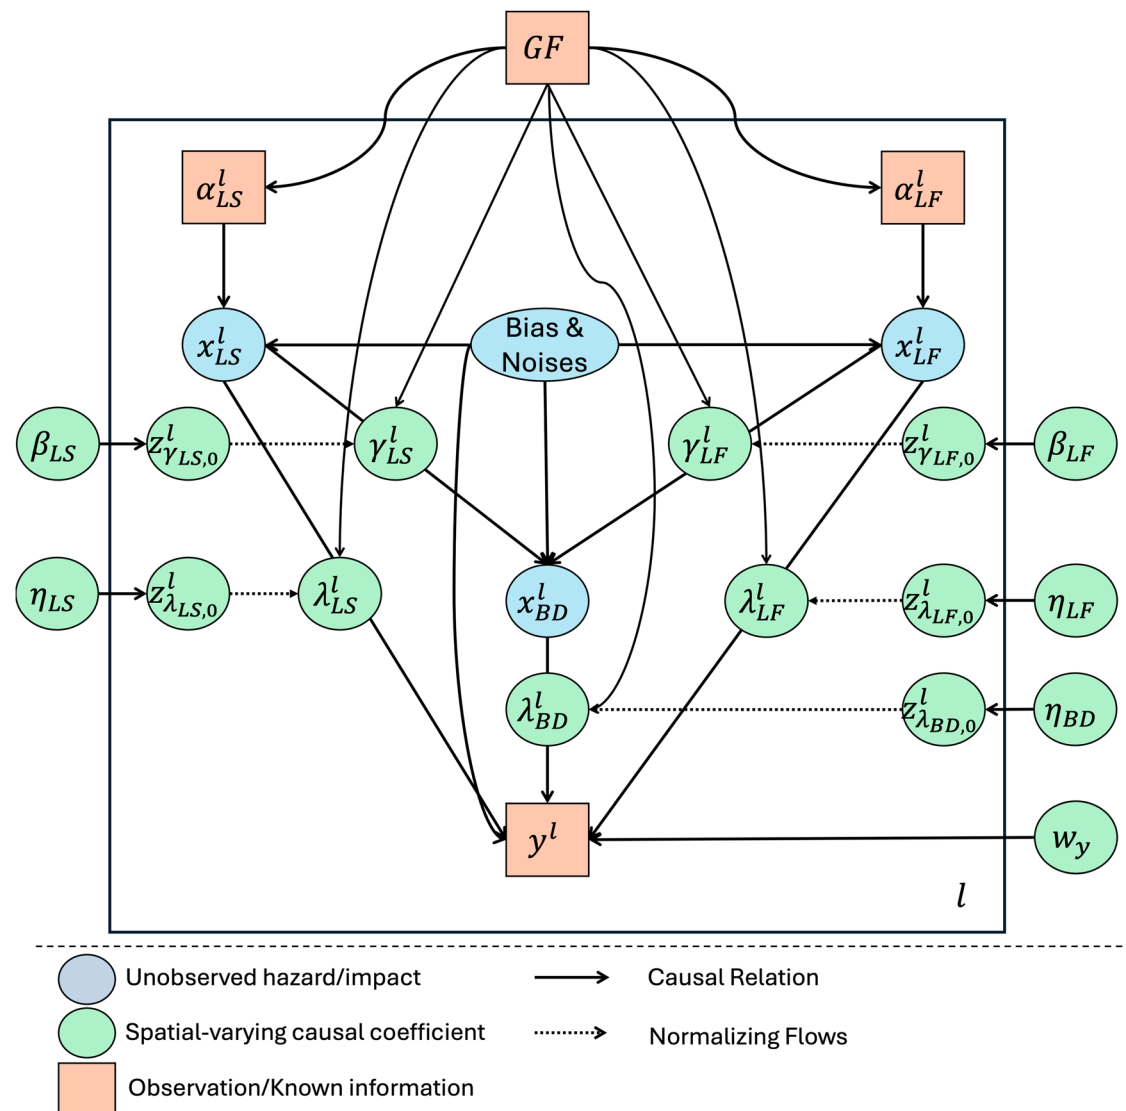


Fig. 9 | Overview of our causal Bayesian inference framework for seismic multi-hazard and impacts estimation. l in the figure refers to the l th location in a target area. Blue circles refer to latent hazard/impact variables. Green circles refer to the spatial-varying causal coefficients. Orange rectangles refer to the observations or known information.

where f_1, \dots, f_{K_v} are invertible transformations and K_v is the number of flow layers. This combination allows us to model both spatial correlation (through the GP) and distributional complexity (through the normalizing flow).

For inference, we approximate the true posterior of the GP latent variables with a multivariate Gaussian variational distribution:

$$q(\mathbf{z}_v) = \mathcal{N}(\boldsymbol{\mu}_v, \boldsymbol{\Sigma}_v) \quad (6)$$

where $\boldsymbol{\mu}_v$ represents the posterior mean vector and $\boldsymbol{\Sigma}_v$ the posterior covariance matrix across all locations. For computational efficiency, we parameterize $\boldsymbol{\Sigma}_v$ as a low-rank plus diagonal structure:

$$\boldsymbol{\Sigma}_v = \mathbf{L}_v(\mathbf{L}_v)^T + \text{diag}(\boldsymbol{\delta}_v) \quad (7)$$

where \mathbf{L}_v is a lower-triangular matrix of rank $r \ll n$ (with n being the number of spatial locations), and $\boldsymbol{\delta}_v$ is a vector of positive diagonal elements.

To enable efficient gradient computation during optimization, we employ the reparameterization trick:

$$\mathbf{z}_v = \boldsymbol{\mu}_v + \mathbf{L}_v \boldsymbol{\epsilon}_1 + \text{diag}(\sqrt{\boldsymbol{\delta}_v}) \boldsymbol{\epsilon}_2, \quad \boldsymbol{\epsilon}_1 \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_r), \boldsymbol{\epsilon}_2 \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_n) \quad (8)$$

This allows us to sample from $q(\mathbf{z}_v)$ while maintaining differentiability with respect to the variational parameters $\boldsymbol{\mu}_v$, \mathbf{L}_v and $\boldsymbol{\delta}_v$.

When transforming the GP latent variables through the normalizing flow, we apply the change of variables formula to compute the density of the transformed variables:

$$\log q(v^l | GF) = \log q(z_v^l) - \sum_{k=1}^{K_v} \log \left| \det \left(\frac{\partial f_k}{\partial z_{v,k-1}^l} \right) \right| \quad (9)$$

where $z_{v,k-1}^l$ represents the output of the $(k-1)$ -th transformation, with $z_{v,0}^l = z_v^l$.

Computing expectations under the transformed distribution is done via Monte Carlo sampling:

$$\begin{aligned} \mathbb{E}_{q(v^l)}[h(v^l)] &= \mathbb{E}_{q(z_v^l)}[h(f_{K_v} \circ \dots \circ f_2 \circ f_1(z_v^l))] \\ &\approx \frac{1}{M} \sum_{m=1}^M h(f_{K_v} \circ \dots \circ f_2 \circ f_1(z_v^{l(m)})) \end{aligned} \quad (10)$$

where $z_v^{l(m)}$ are samples drawn from $q(z_v^l)$ using the reparameterization in Equation (8), and M is the number of Monte Carlo samples.

In our implementation, we utilize planar flows of the form:

$$f_k(z^l; u_k, c_k, b_k) = z^l + u_k h(c_k^T z^l + b_k) \quad (11)$$

where h is a tanh activation function with derivative $h'(\cdot)$, and u_k, c_k, b_k are learnable parameters of the k -th flow. The log-determinant of the Jacobian for this transformation is:

$$\log |\det \left(\frac{df_k}{dz_{k-1}^l} \right)| = \log |1 + u_k^T \psi(z_{k-1}^l)| \quad (12)$$

where $\psi(z) = c_k \cdot h'(c_k^T z + b_k)$.

We parameterize the mean function $m_v(\mathbf{GF})$ as a function of the geophysical features:

$$m_v(\mathbf{GF}_l) = \text{NN}_v(\mathbf{GF}_l) \quad (13)$$

where NN_v is a neural network that maps geophysical features at location i to the mean of the corresponding GP latent variable z_v^l .

This GP-Normalizing Flow approach provides several advantages for modeling spatially heterogeneous causal effects: (1) it naturally captures spatial correlation through the GP prior, (2) it allows for complex, non-Gaussian distributions through the normalizing flow transformation, and (3) it provides a principled way to incorporate uncertainty in the causal parameter estimates.

Stochastic variational inference with spatial-variant causal parameters

The ultimate goal is to jointly infer the true posteriors of multiple unobserved target variables $x_i \in X$, which represents the target hazards and impacts with unknown parameters of causal dependencies, and the unknown spatial-varying causal parameters λ_a^l , where $a \in \{LS, LF, BD\}$, and γ_b^l , where $b \in \{\alpha_{LS}, \alpha_{LF}, LS, LF\}$, across different locations that quantify the causal dependencies among parent nodes and child nodes. Therefore, we use variational inference to approximate the true posteriors of X^l using $q(X^l)$ by optimizing the variational lower bound. Since the variable X^l is binary, we can further factorize $q(X^l)$ over hidden (unobserved) nodes by:

$$q(X^l) = \prod_{i \in \{LS, LF, BD\}} q(x_i^l) = \prod_{i \in \{LS, LF, BD\}} (q_i^l)^{x_i^l} (1 - q_i^l)^{1-x_i^l} \quad (14)$$

where q_i^l is defined to approximate the posterior probability that node i is active in location l .

We derive the lower bound of the marginal log-likelihood of the observed Y by Jensen's inequality as follows:

$$\log P(Y, GF) = \underbrace{\mathbb{E}_{q(X^l, \epsilon^l, \mathbf{z}^l)} [\log p(y^l, GF, X^l, \mathbf{z}^l)]}_{[1]} - \underbrace{\mathbb{E}_{q(X^l, \epsilon^l, \mathbf{z}^l)} [\log q(X^l, \epsilon^l, \mathbf{z}^l)]}_{[2]} \quad (15)$$

Here \mathbf{z}^l represents the latent Gaussian Process variables, which are transformed through a normalizing flow to produce the spatially-varying parameters λ^l and γ^l shown in Equation (5).

The item [1] is further expanded as:

$$\begin{aligned} \mathbb{E}_{q(X^l, \epsilon^l, \mathbf{z}^l)} [\log p(y^l, GF, X^l, \mathbf{z}^l)] &= \underbrace{\mathbb{E} [\log p(y^l | x_{P(y)}^l, f_{K_1} \circ \dots \circ f_2 \circ f_1(\mathbf{z}_\lambda^l, \epsilon_\gamma^l, GF))]}_{[3]} \\ &+ \underbrace{\sum_{i \in \{LS, LF, BD\}} \mathbb{E} [\log p(x_i^l | f_{K_i} \circ \dots \circ f_2 \circ f_1(\mathbf{z}_\gamma^l, \mathcal{P}(x_i^l), \epsilon_i^l, GF))]}_{[4]} \\ &+ \underbrace{\mathbb{E} [\log p(\mathbf{z}_\lambda^l | GF)]}_{[5]} + \underbrace{\mathbb{E} [\log p(\mathbf{z}_\gamma^l | GF)]}_{[6]} \\ &+ \underbrace{\sum_{i \in \{LS, LF, BD\}} \mathbb{E} [\log p(\epsilon_i^l)] + \mathbb{E} [\log p(\epsilon_\gamma^l)]}_{C_1} \end{aligned} \quad (16)$$

Where $p(\mathbf{z}_\lambda^l | GF)$ and $p(\mathbf{z}_\gamma^l | GF)$ are Gaussian Process priors shown in Equation (3). As for item [2], we have:

$$\begin{aligned} \mathbb{E}_{q(X^l, \epsilon^l, \mathbf{z}^l)} [\log q(X^l, \epsilon^l, \mathbf{z}^l)] &= \underbrace{\sum_i \mathbb{E}(x_i^l) \log q_i^l + \sum_i \mathbb{E}(1 - x_i^l) \log (1 - q_i^l)}_{[7]} \\ &+ \underbrace{\mathbb{E} [\log q(\mathbf{z}_\lambda^l)]}_{[8]} + \underbrace{\mathbb{E} [\log q(\mathbf{z}_\gamma^l)]}_{[9]} \\ &+ \underbrace{\sum_i \mathbb{E} [\log q(\epsilon_i^l)] + \mathbb{E} [\log q(\epsilon_\gamma^l)]}_{C_2} \end{aligned} \quad (17)$$

where C_1 and C_2 cancel out. $q(\mathbf{z}_\lambda^l)$ and $q(\mathbf{z}_\gamma^l)$ are Gaussian Process posteriors presented in Equation (6).

For terms [5], [6], [8], and [9] involving the Gaussian Process, we apply the KL divergence formula for multivariate Gaussians:

$$\begin{aligned} [5] - [8] &= -KL[q(\mathbf{z}_\lambda^l) || p(\mathbf{z}_\lambda^l | GF)] \\ &= -\frac{1}{2} [\text{tr}(\mathbf{K}_\lambda^{-1} \Sigma_\lambda) + (\boldsymbol{\mu}_\lambda - \mathbf{m}_\lambda)^T \mathbf{K}_\lambda^{-1} (\boldsymbol{\mu}_\lambda - \mathbf{m}_\lambda) + \log |\mathbf{K}_\lambda| - \log |\Sigma_\lambda| - n_\lambda] \end{aligned} \quad (18)$$

$$\begin{aligned} [6] - [9] &= -KL[q(\mathbf{z}_\gamma^l) || p(\mathbf{z}_\gamma^l | GF)] \\ &= -\frac{1}{2} [\text{tr}(\mathbf{K}_\gamma^{-1} \Sigma_\gamma) + (\boldsymbol{\mu}_\gamma - \mathbf{m}_\gamma)^T \mathbf{K}_\gamma^{-1} (\boldsymbol{\mu}_\gamma - \mathbf{m}_\gamma) + \log |\mathbf{K}_\gamma| - \log |\Sigma_\gamma| - n_\gamma] \end{aligned} \quad (19)$$

where \mathbf{K} represents the prior covariance matrices derived from the Matern kernel over geophysical features, Σ denotes the variational posterior covariance matrices, $\boldsymbol{\mu}$ and \mathbf{m} are the posterior and prior mean vectors respectively, and n is the dimensionality of the corresponding latent variable vectors. These equations quantify the divergence between our variational approximation and the true GP prior distributions for both λ and γ coefficients.

Based on our conditional distribution assumptions for $y^l | \mathcal{P}(y^l), \epsilon_\gamma^l, \lambda_{\mathcal{P}(y^l)}^l$ as defined in Equation (1), we can calculate item [3] in Equation (16) as:

$$\begin{aligned} \mathbb{E} [\log p(y^l | \mathcal{P}(y^l), \epsilon_\gamma^l, \lambda_{\mathcal{P}(y^l)}^l)] &= -\log y^l - \log |w_{\epsilon_\gamma}| - \frac{(\log y^l) + w_{0y}^2 + \sum_{k \in \mathcal{P}(y^l)} \mathbb{E}[(\lambda_k^l)^2] q_k^l - 2 \log y^l \cdot w_{0y}}{2w_{\epsilon_\gamma}^2} \\ &- \frac{\sum_{j \neq k} \mathbb{E}(\lambda_j^l) \cdot \mathbb{E}(\lambda_k^l) \cdot q_j^l \cdot q_k^l + (w_{0y} - \log y^l) (\sum_{k \in \mathcal{P}(y^l)} \mathbb{E}(\lambda_k^l) \cdot q_k^l)}{w_{\epsilon_\gamma}^2} \end{aligned} \quad (20)$$

As for item [4] in Equation (16), we consider two scenarios—when x_i is a leaf node and a non-leaf node. First we can describe the conditional distribution of LS, LF, and BD as follows:

$$\begin{aligned} p(x_i^l | x_{\mathcal{P}(x_i)}^l, \epsilon_i^l, \gamma_{\mathcal{P}(x_i)}^l) &= \\ \left[\frac{1}{1 + \exp \left(-\sum_{k \in \mathcal{P}(x_i)} \gamma_k^l x_k^l - w_{\epsilon_i} \epsilon_i^l - w_{0i} \right)} \right]^{x_i^l} \cdot \left[\frac{1}{1 + \exp \left(\sum_{k \in \mathcal{P}(x_i)} \gamma_k^l x_k^l + w_{\epsilon_i} \epsilon_i^l + w_{0i} \right)} \right]^{1-x_i^l} \end{aligned} \quad (21)$$

The expectation of Equation (21) can be formulated as:

$$\begin{aligned} \mathbb{E} [\log p(x_i^l | x_{\mathcal{P}(x_i)}^l, \epsilon_i^l, \gamma_{\mathcal{P}(x_i)}^l)] &= q_i^l \mathbb{E} \left[-\log \left(1 + \exp \left(-\sum_{k \in \mathcal{P}(x_i)} \gamma_k^l x_k^l - w_{\epsilon_i} \epsilon_i^l - w_{0i} \right) \right) \right] \\ &+ (1 - q_i^l) \mathbb{E} \left[-\log \left(1 + \exp \left(\sum_{k \in \mathcal{P}(x_i)} \gamma_k^l x_k^l + w_{\epsilon_i} \epsilon_i^l + w_{0i} \right) \right) \right] \end{aligned} \quad (22)$$

However, the distribution of $-\log [1 + \exp (\sum_{k \in \mathcal{P}(x_i)} \gamma_k^l x_k^l + w_{\epsilon_i} \epsilon_i^l + w_{0i})]$ is intractable, as it involves a log-sum-exp function that mixes both discrete and continuous variables. Consequently, we need to derive a tight

lower bound for its expectation. Without loss of generality, we begin with the case where node i has a single active parent. Given that the function $-\log x$ is convex, Jensen's inequality, combined with Taylor's theorem, allows us to establish the following relationship:

$$\mathbb{E}[-\log(1 + \exp(x))] \geq -\log(1 + \mathbb{E}[\exp(x)]) \quad (23)$$

Therefore, we obtain the lower bound of Equation (22) as:

$$\begin{aligned} \mathbb{E}[\log p(x_i^l | x_{\mathcal{P}(x_i)}^l, \epsilon_i^l, \gamma_{\mathcal{P}(x_i)}^l)] &\geq q_i^l \left[-\log \left(1 + \mathbb{E} \left[\exp \left(- \sum_{k \in \mathcal{P}(x_i)} \gamma_k^l x_k^l - w_{\epsilon_i} \epsilon_i^l - w_{0i} \right) \right] \right) \right] \\ &\quad + (1 - q_i^l) \left[-\log \left(1 + \mathbb{E} \left[\exp \left(\sum_{k \in \mathcal{P}(x_i)} \gamma_k^l x_k^l + w_{\epsilon_i} \epsilon_i^l + w_{0i} \right) \right] \right) \right] \\ &= -q_i^l \left\{ \log \left[1 + \left(\prod_{k \in \mathcal{P}(x_i)} [(1 - q_k^l) + q_k^l \mathbb{E}[\exp(-\gamma_k^l)]] \right) \cdot \exp \left(\frac{w_{\epsilon_i}^2}{2} - w_{0i} \right) \right] \right\} \\ &\quad - (1 - q_i^l) \left\{ \log \left[1 + \left(\prod_{k \in \mathcal{P}(x_i)} [(1 - q_k^l) + q_k^l \mathbb{E}[\exp(\gamma_k^l)]] \right) \cdot \exp \left(\frac{w_{\epsilon_i}^2}{2} + w_{0i} \right) \right] \right\} \end{aligned} \quad (24)$$

As for item [7] in Equation (17), we have:

$$\begin{aligned} \sum_i \mathbb{E}(x_i^l) \log q_i^l + \sum_i \mathbb{E}(1 - x_i^l) \log(1 - q_i^l) \\ = \sum_i q_i^l \log q_i^l + \sum_i (1 - q_i^l) \log(1 - q_i^l) \end{aligned} \quad (25)$$

Given a map containing a set of locations, $l \in L$, we further derive a tight lower bound for the log-likelihood as follows:

$$\begin{aligned} \mathcal{L} &= \log P(Y, GF) \\ &= -\log y^l - \log |w_{\epsilon_y}| - \frac{(\log y^l + w_{0y}^2 + \sum_{k \in \mathcal{P}(y^l)} \mathbb{E}[(\lambda_k^l)^2] q_k^l - 2 \log y^l \cdot w_{0y})}{2w_{\epsilon_y}^2} \\ &\quad - \frac{\sum_{j \neq k} \mathbb{E}(\lambda_k^l) \cdot \mathbb{E}(\lambda_j^l) \cdot q_k^l \cdot q_j^l + (w_{0y} - \log y^l) \left(\sum_{k \in \mathcal{P}(y^l)} \mathbb{E}(\lambda_k^l) \cdot q_k^l \right)}{w_{\epsilon_y}^2} \\ &\quad - \sum_{i \in \{LS, LF, BD\}} \left\{ q_i^l \left[\log \left[1 + \left(\prod_{k \in \mathcal{P}(x_i)} [(1 - q_k^l) + q_k^l \mathbb{E}[\exp(-\gamma_k^l)]] \right) \cdot \exp \left(\frac{w_{\epsilon_i}^2}{2} - w_{0i} \right) \right] \right] \right\} \\ &\quad - (1 - q_i^l) \left\{ \log \left[1 + \left(\prod_{k \in \mathcal{P}(x_i)} [(1 - q_k^l) + q_k^l \mathbb{E}[\exp(\gamma_k^l)]] \right) \cdot \exp \left(\frac{w_{\epsilon_i}^2}{2} + w_{0i} \right) \right] \right\} \\ &\quad - \sum_{a \in \{LS, LF, BD\}} \frac{1}{2} \left[\text{tr}(\mathbf{K}_{\lambda_a}^{-1} \Sigma_{\lambda_a}) + (\mu_{\lambda_a} - \mathbf{m}_{\lambda_a})^T \mathbf{K}_{\lambda_a}^{-1} (\mu_{\lambda_a} - \mathbf{m}_{\lambda_a}) + \log |\mathbf{K}_{\lambda_a}| - \log |\Sigma_{\lambda_a}| - n_{\lambda_a} \right] \\ &\quad - \sum_{b \in \{\alpha_{LS}, \alpha_{LF}, LS, LF\}} \frac{1}{2} \left[\text{tr}(\mathbf{K}_{\gamma_b}^{-1} \Sigma_{\gamma_b}) + (\mu_{\gamma_b} - \mathbf{m}_{\gamma_b})^T \mathbf{K}_{\gamma_b}^{-1} (\mu_{\gamma_b} - \mathbf{m}_{\gamma_b}) + \log |\mathbf{K}_{\gamma_b}| - \log |\Sigma_{\gamma_b}| - n_{\gamma_b} \right] \\ &\quad - \sum_i q_i^l \log q_i^l - \sum_i (1 - q_i^l) \log(1 - q_i^l) \end{aligned} \quad (26)$$

Training and stochastic optimization with sparse Gaussian processes

We aim to minimize our loss function in order to find optimal combinations of posteriors and parameters of causal dependencies (including the weights of parent nodes, and parameters of the normalizing flows). To achieve this, we develop an expectation-maximization (E-M) algorithm that alternates

between updating the posteriors of unobserved variables (e.g., LS, LF, BD) with causal effects, and flow parameters.

To manage computational complexity when working with large geographical regions, we employ a sparse Gaussian Process formulation using inducing points. Rather than modeling the GP latent field directly at every location, we introduce a set of M inducing points \mathbf{u}_v at strategic locations in feature space. This approach reduces the computational complexity from $O(N^3)$ to $O(NM^2)$, where N is the number of locations and $M \ll N$ is the number of inducing points.

We model spatial correlation in the causal parameter field using a Gaussian Process with a Matern kernel operating in feature space rather than geographical space. This approach allows the model to capture correlations between locations with similar geophysical characteristics (elevation, slope, lithology, etc.) even when they are not physically adjacent. This is motivated by the observation that similar feature combinations tend to exhibit similar causal relationships in earthquake-triggered hazard scenarios regardless of their geographic proximity.

Within each iteration, we sample a mini-batch of locations and perform the following two steps:

- **Expectation step:** Update the posterior probability estimates of the unobserved latent variables (LS, LF, BD) at each location by conditioning on the current GP parameters, neural network weights, and the most recent samples of causal coefficients derived from the GP

latent variables through normalizing flows. The posterior probabilities are updated according to the causal dependencies encoded in the Bayesian network structure.

- **Maximization step:** Update the parameters of the sparse Gaussian Process, specifically the variational parameters for inducing points \mathbf{u}_u , along with the neural network parameters that map geophysical features to the mean function of the GP. Specifically, for iteration $t + 1$, these parameters are updated as

$$\mu_u^{(t+1)} = \mu_u^{(t)} + \rho \mathbf{A} \nabla \mathcal{L}(\mu_u), \quad \Sigma_u^{(t+1)} = \Sigma_u^{(t)} + \rho \mathbf{A} \nabla \mathcal{L}(\Sigma_u), \quad (27)$$

$$\theta_{NN}^{(t+1)} = \theta_{NN}^{(t)} + \rho \mathbf{A} \nabla \mathcal{L}(\theta_{NN}), \quad (28)$$

where μ_u and Σ_u are the parameters of the variational distributions over inducing points, θ_{NN} represents the neural network parameters, \mathbf{A} is a positive definite preconditioner, ρ is the learning rate, and $\nabla \mathcal{L}(\cdot)$ denotes gradients of the loss function with respect to the corresponding parameters. The causal coefficients at observed locations are then derived through conditioning on the inducing points and applying normalizing flow transformations to the resulting GP latent variables. This gradient update scheme is guaranteed to converge to a local maximum of \mathcal{L} if ρ satisfies appropriate decay conditions.

Our variational evidence lower bound (ELBO) with inducing points can be formulated as:

$$\mathcal{L} = \mathbb{E}_q[\log p(\mathbf{y}|\mathbf{X}, \mathbf{z}_v)] + \mathbb{E}_q[\log p(\mathbf{X}|\mathbf{z}_v)] - KL[q(\mathbf{u}_v)||p(\mathbf{u}_v)] - \mathbb{E}_q[\log q(\mathbf{X})] \quad (29)$$

where $q(\mathbf{X}, \mathbf{z}_v, \mathbf{u}_v, \epsilon) = q(\mathbf{X}) \times p(\mathbf{z}_v|\mathbf{u}_v) \times q(\mathbf{u}_v) \times p(\epsilon)$ is our variational distribution. By using the exact conditional $p(\mathbf{z}_v|\mathbf{u}_v)$ in our variational approximation, we ensure that the model maintains valid probabilistic semantics over the entire spatial field while only needing to represent distributions at the inducing points explicitly.

Once the model converges, we obtain the final posterior estimates of LS, LF, and BD for each location, and also the causal effects between hidden hazards and observations, reflecting the spatially heterogeneous causal relationships inferred from the observed data.

We also apply a local pruning algorithm to accelerate the computation over a large region. This strategy is motivated by the observation that real-world causal graphs are typically sparse: only a small subset of nodes stay active. For example, locations without building footprints will not have damaged buildings, i.e., building damage nodes are inactive. Therefore, we can prune these inactive nodes while keeping the active ones crucial for parameter updates^{3,12}.

Data availability

Data used in this study were collected from several publicly accessible sources. The primary observational data consists of Damage Proxy Maps (DPMs) generated by NASA's Advanced Rapid Imaging and Analysis (ARIA) team using InSAR imagery from Sentinel-1 satellites, available at <https://aria-share.jpl.nasa.gov/>. These DPMs capture correlation changes between pre- and post-event images, providing valuable information for rapid hazard and impact estimation. For ground truth validation, we collected data from multiple sources across the three earthquake events studied. For the 2021 Haiti earthquake (M7.2), building damage and landslide inventories were provided by StEER (available at: <https://www.steer.network/haiti-response>) and GEER teams. Field reconnaissance data for the 2020 Puerto Rico earthquake (M6.4) was collected by USGS, University of Puerto Rico Mayagüez, GEER, and StEER teams, available at <https://www.sciencebase.gov/catalog/item/5eb5b9dc82ce25b5135ae83a>. For the 2023 Turkey-Syria earthquake sequence (M7.8), we utilized building damage inventory data from the Turkish Ministry of Environment. Additional data sources used in

this study include USGS ShakeMap and ground failure models (<https://earthquake.usgs.gov/>) as prior models, along with building footprints from OpenStreetMap (<https://www.openstreetmap.org/>). Any data not available through these public repositories may be obtained from the corresponding author upon reasonable request.

Code availability

The underlying code and training/validation datasets for this study are available in the repository and can be accessed via <https://github.com/PaperSubmissionFinal/SpatialBN>.

Received: 31 March 2025; Accepted: 13 May 2025;

Published online: 23 July 2025

References

1. Havenith, H. -B. et al. First analysis of landslides triggered by the August 14, 2021, Nippes (Haiti) earthquake, compared with the 2010 event. In *EGU General Assembly Conference Abstracts* (EGU, 2022).
2. Jang, H. -C., Lien, Y. N. & Tsai, T. C. Rescue information system for earthquake disasters based on MANET emergency communication platform. In *Proc. 2009 International Conference on Wireless Communications and Mobile Computing: Connecting the World Wirelessly*, 623–627 (ACM, 2009).
3. Li, X. et al. Disasternet: causal Bayesian networks with normalizing flows for cascading hazards estimation from satellite imagery. In *Proc. 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining* (ACM) 4391–4403 (2023).
4. Toprak, S. & Holzer, T. L. Liquefaction potential index: field assessment. *J. Geotech. Geoenviron. Eng.* **129**, 315–322 (2003).
5. Marc, O., Hovius, N., Meunier, P., Gorum, T. & Uchida, T. A seismologically consistent expression for the total area and volume of earthquake-triggered landsliding. *J. Geophys. Res. Earth Surf.* **121**, 640–663 (2016).
6. Newmark, N. M. Effects of earthquakes on dams and embankments. *Geotechnique* **15**, 139–160 (1965).
7. Ghofrani, H. & Atkinson, G. M. Site condition evaluation using horizontal-to-vertical response spectral ratios of earthquakes in the Nga-West 2 and Japanese databases. *Soil Dyn. Earthq. Eng.* **67**, 30–43 (2014).
8. Yun, S.-H. et al. Rapid damage mapping for the 2015 M w 7.8 Gorkha earthquake using synthetic aperture radar data from COSMO-SkyMed and ALOS-2 satellites. *Seismol. Res. Lett.* **86**, 1549–1556 (2015).
9. Fielding, E. J. & Jung, J. Damage proxy mapping with SAR interferometric coherence change. *Procedia Comput. Sci.* **239**, 2322–2328 (2024).
10. Yu, X. et al. Intelligent assessment of building damage of 2023 Turkey-Syria earthquake by multiple remote sensing approaches. *npj Nat. Hazards* **1**, 3 (2024).
11. Kongar, I., Rossetto, T. & Giovinazzi, S. Evaluating simplified methods for liquefaction assessment for loss estimation. *Nat. Hazards Earth Syst. Sci.* **17**, 781–800 (2017).
12. Xu, S., Dimasaka, J., Wald, D. J. & Noh, H. Y. Seismic multi-hazard and impact estimation via causal inference from satellite imagery. *Nat. Commun.* **13**, 7793 (2022).
13. Wald, D. J., Xu, S., Dimasaka, J. & Noh, H. Causality-informed Bayesian inference for rapid seismic ground failure and building damage estimation. In *Proc. 12th National Conference on Earthquake Engineering* (EERI, 2023).
14. Li, X. et al. Rapid building damage estimates from the M7.8 Turkey earthquake sequence via causality-informed Bayesian inference from satellite imagery. *Earthq. Spectra* **41**, 5–33 (2025).
15. Wang, C. et al. Scalable and rapid building damage detection after Hurricane Ian using causal Bayesian networks and InSAR imagery. *Int. J. Disaster Risk Reduct.* **104**, 104371 (2024).

16. Wang, C. et al. Causality-informed rapid post-hurricane building damage detection in large scale from InSAR imagery. In *Proc. 8th ACM SIGSPATIAL International Workshop on Security Response Using GIS 7–12* (2023).
17. Li, X. et al. M7.8 Turkey-Syria Earthquake Impact Estimates from Near-real-time Crowdsourced and Remote Sensing Data. DesignSafe-CI. <https://doi.org/10.17603/ds2-vnsc-y870 v2> (2023).
18. Li, X. et al. Normalizing flow-based deep variational Bayesian network for seismic multi-hazards and impacts estimation from InSAR imagery. *arXiv preprint arXiv:2310.13805* (2023).
19. Li, X., Gao, R., Burgi, P. M., Wald, D. J. & Xu, S. Optimizing rapid seismic building damage assessment: Integrating enhanced radar change detection maps with variational Bayesian networks. *IGARSS 2024-2024 IEEE International Geoscience and Remote Sensing Symposium* 3791–3795 (IEEE, 2024).
20. Li, X. & Xu, S. Spatial-variant causal Bayesian inference for rapid seismic ground failures and impacts estimation. *arXiv preprint arXiv:2412.00026* (2024).
21. Tomasi, C. & Manduchi, R. Bilateral filtering for gray and color images. In *Sixth International Conference on Computer Vision* (IEEE Cat. No. 98CH36271) 839–846 <https://api.semanticscholar.org/CorpusID:14308539> (1998).
22. Paris, S. et al. Bilateral filtering: theory and applications: Series: foundations and trends in computer graphics and vision <https://api.semanticscholar.org/CorpusID:67417024> (2009).
23. Gelfand, A. E. & Schliep, E. M. Spatial statistics and Gaussian processes: a beautiful marriage. *Spat. Stat.* **18**, 86–104 (2016).
24. Diggle, P. J. & Ribeiro, P. J. Model-based geostatistic <https://api.semanticscholar.org/CorpusID:134130701> (2013).
25. Xu, C., Xu, X., Dai, F. & Saraf, A. K. Comparison of different models for susceptibility mapping of earthquake triggered landslides related with the 2008 Wenchuan earthquake in China. *Comput. Geosci.* **46**, 317–329 (2012).
26. Banerjee, S., Carlin, B. P. & Gelfand, A. E. Hierarchical modeling and analysis for spatial data <https://api.semanticscholar.org/CorpusID:62708858> (2003).
27. Ray, A. & Myer, D. Bayesian geophysical inversion with trans-dimensional Gaussian process machine learning. *Geophys. J. Int.* **217**, 1706–1726 (2019).
28. Rezende, D. & Mohamed, S. Variational inference with normalizing flows. *International Conference on Machine Learning* 1530–1538 (PMLR, 2015).
29. Advanced Rapid Imaging and Analysis (ARIA)—Center for Natural Hazards. ARIA Data Share https://aria-share.jpl.nasa.gov/20200106-Puerto_Rico_EQ/DPM/ (2020).
30. Allstadt, K. et al. Ground failure triggered by the 7 January 2020 Mw 6.4 Puerto Rico earthquake. *Seismol. Res. Lett.* **93**, 594–608 (2022).
31. Miranda, E., Acosta, A., Aponte-Bermudez, L. D. & Roueche, D. StEER-07 Jan. 2020 Puerto Rico mw6. 4 earthquake: preliminary virtual reconnaissance report (PVRr). DesignSafe-CI <https://doi.org/10.17603/ds2-xfhz-fz88> (2020).
32. Miranda, E. et al. StEER—Puerto Rico Earthquake Sequence December 2019 to January 2020: Field Assessment Structural Team (FAST) Early Access Reconnaissance Report (EARR). DesignSafe-CI <https://doi.org/10.17603/ds2-h0kd-5677> (2020).
33. Zhao, B., Wang, Y., Li, W., Lu, H. & Li, Z. Evaluation of factors controlling the spatial and size distributions of landslides, 2021 nippees earthquake, haiti. *Geomorphology* **415**, 108419 (2022).
34. Network, S. StEER's preliminary virtual reconnaissance report for m7.2 Nippes, Haiti earthquake released. <https://www.steer.network/haiti-response>.
35. Reconnaissance, G. E. E. 2021 Haiti earthquake https://geerassociation.org/component/geer_reports/?view=geerreports&id=100&layout=build (2021).
36. Haiti, O. Haiti: Earthquake situation report no. 4 (7 september 2021). <https://digitalcommons.fiu.edu/cgi/viewcontent.cgi?article=1854&context=srhreports> (2021).
37. U. G. Survey. United States Geological Survey. m 7.2—Nippes, Haiti. <https://earthquake.usgs.gov/earthquakes/eventpage/us6000f65h/executive> (2021).
38. ARIA Data Share. Advanced Rapid Imaging and Analysis (ARIA)—Center for Natural Hazards. Online <https://aria-share.jpl.nasa.gov/> (2020).
39. Dilsiz, A. et al. StEER: 2023 mw7.8 Kahramanmaraş, Türkiye earthquake sequence preliminary virtual reconnaissance report (PVRr) <https://www.designsafe-ci.org/data/browser/public/designsafe.storage.published/PRJ-3824v2/#details-942732811040452115-242ac11b-0001-012> (2023).
40. Fawcett, T. An introduction to ROC analysis. *Pattern Recognit. Lett.* **27**, 861–874 (2006).
41. Zhu, J., Baise, L. & Thompson, E. Updated geospatial liquefaction model for global use. *Seismol. Res. Lett.* **107** (2016).
42. Nowicki Jessee, M. et al. A global empirical model for near-real-time assessment of seismically induced landslides. *J. Geophys. Res.: Earth Surf.* **123**, 1835–1859 (2018).
43. Novellino, A. et al. Slow-moving landslide risk assessment combining machine learning and InSAR techniques. *Catena* **203**, 105317 (2021).
44. Rao, A., Jung, J., Silva, V., Molinario, G. & Yun, S.-H. Earthquake building damage detection based on synthetic-aperture-radar imagery and machine learning. *Nat. Hazards Earth Syst. Sci.* **23**, 789–807 (2023).
45. Banerjee, S., Carlin, B. P. & Gelfand, A. E. *Hierarchical Modeling And Analysis For Spatial Data* (Chapman and Hall/CRC, 2003).
46. Wald, D., Wald, L., Worden, B. & Goltz, J. Shakemap, a tool for earthquake response. *Fact Sheet* <https://doi.org/10.3133/fs08703> (2003).
47. Wald, D. J., Quitoriano, V., Heaton, T. H. & Kanamori, H. Relationships between peak ground acceleration, peak ground velocity, and modified Mercalli intensity in California. *Earthq. Spectra* **15**, 557–564 (1999).
48. Seyhan, E. & Stewart, J. P. Semi-empirical nonlinear site amplification from nga-west2 data and simulations. *Earthq. Spectra* **30**, 1241–1256 (2014).
49. Kamp, U., Growley, B. J., Khattak, G. A. & Owen, L. A. GIS-based landslide susceptibility mapping for the 2005 Kashmir earthquake region. *Geomorphology* **101**, 631–642 (2008).
50. Danielson, J. J. & Gesch, D. B. *Global Multi-resolution Terrain Elevation Data 2010* (gmted2010). Technical Reports (US Geological Survey, 2011).
51. Arino, O. et al. Global Land CoverMap for 2009 (Glob Cover 2009) [dataset]. © European Space Agency (ESA) & Université catholique de Louvain (UCL), PANGAEA. <https://doi.org/10.1594/PANGAEA.787668> (2012).
52. Farr, T. G. et al. The shuttle radar topography mission. *Rev. Geophys.* **45** (2007).
53. Beven, K. J. & Kirkby, M. J. A physically based, variable contributing area model of basin hydrology/un modèle à base physique de zone d'appel variable de l'hydrologie du bassin versant. *Hydrol. Sci. J.* **24**, 43–69 (1979).
54. Lehner, B., Verdin, K. & Jarvis, A. Hydrosheds technical documentation. *World Wildlife Fund US, Washington, DC* **5** (2006).
55. Hartmann, J. & Moosdorf, N. The new global lithological map database GLIM: a representation of rock properties at the Earth surface. *Geochem. Geophys. Geosyst.* **13** (2012).
56. Matérn, B. *Spatial Variation* Vol. 36 (Springer Science & Business Media, 2013).
57. Ma, P. & Bhadra, A. Beyond matérn: on a class of interpretable confluent hypergeometric covariance functions. *J. Am. Stat. Assoc.* **118**, 2045–2058 (2023).
58. Muir, J. B. & Ross, Z. E. A deep Gaussian process model for seismicity background rates. *Geophys. J. Int.* **234**, 427–438 (2023).
59. Federal Emergency Management Agency. Hazus-mh technical manual. Technical Manual FEMA-433, Department of Homeland

Security Emergency Preparedness and Response Directorate, Washington, D.C. (2003).

60. U.S. Geological Survey. M 7.8—26 km ene of Nurdagi, Turkey—ground failure. Online <https://earthquake.usgs.gov/earthquakes/eventpage/us6000jllz/ground-failure/summary> (2023).

Acknowledgements

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: X.L., S.G., R.G., and S.X. are supported by NSF DRRG CMMI-2242590 and CAREER 2442712. Any mention of commercial products is for informational purposes and does not constitute an endorsement by the US government.

Author contributions

X.L., S.G., and S.X. conceptualized the research, developed the framework. R.G. conducted the data cleaning. X.L., S.G., R.G. implemented the code. X.L. and S.G. conducted the experiments, analyzed the results. S.X. acquired the funding to support this research. All authors wrote, read, and approved the final manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Correspondence and requests for materials should be addressed to Susu Xu.

Reprints and permissions information is available at <http://www.nature.com/reprints>

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

© The Author(s) 2025