

Seismic multi-hazard and impact estimation via causal inference from satellite imagery

Received: 10 January 2022

Susu Xu ^{1,2}✉, Joshua Dimasaka ³, David J. Wald ⁴ & Hae Young Noh³

Accepted: 1 December 2022

Published online: 17 December 2022

 Check for updates

Rapid post-earthquake reconnaissance is important for emergency responses and rehabilitation by providing accurate and timely information about secondary hazards and impacts, including landslide, liquefaction, and building damage. Despite the extensive collection of geospatial data and satellite images, existing physics-based and data-driven methods suffer from low estimation performance due to the complex and event-specific causal dependencies underlying the cascading processes of earthquake-triggered hazards and impacts. Herein, we present a rapid seismic multi-hazard and impact estimation system that leverages advanced statistical causal inference and remote sensing techniques. The unique feature of this system is that it provides accurate and high-resolution estimations on a regional scale by jointly inferring multiple hazards and building damage from satellite images through modeling their causal dependencies. We evaluate our system on multiple seismic events from diverse countries around the globe. Our results corroborate that incorporating causal dependencies significantly improves large-scale estimation accuracy for multiple hazards and impacts compared to existing systems. The results also reveal quantitative causal mechanisms among earthquake-triggered multi-hazard and impact for multiple seismic events. Our system establishes a new way to extract and utilize the complex interactions of multiple hazards and impacts for effective disaster responses and advancing understanding of seismic geological processes.

Moderate-to-large earthquakes are often followed by a series of ground failures and subsequent impacts, such as landslides, liquefaction, and building damage. These cascading hazards and impacts exacerbate seismic losses, including fatalities and disruption to lifelines¹. For example, the 2018 Hokkaido earthquake in Japan triggered more than 6000 landslides and 1000 associated debris deposits, resulting in more than 80% of the casualties². The series of earthquakes in Ridgecrest, California, in 2019, induced liquefaction in the town of Trona, where one large lateral spread caused severe damage to around 30 square blocks³.

Earthquake rapid response systems play an important role immediately after earthquakes in rapidly and accurately locating and

estimating the spatial occurrences of seismic hazards (e.g., landslides, liquefaction) and seismic impacts (e.g., building damage). This information serves to direct response and recovery to those areas most in need of critical assistance. Some existing global or regional rapid seismic hazard and impact estimation efforts include GDACS—Global Disaster Alert and Coordination System⁴, PAGER—Prompt Assessment of Global Earthquakes for Response⁵ from the U.S. Geological Survey (USGS), and NERIES-ELER—Network of Research Infrastructures for European Seismology⁶.

Despite the deployment of sophisticated earthquake rapid response systems, the task of accurate and timely localization and estimation of earthquake impacts remains challenging due to

¹Department of Civil Engineering, Stony Brook University, Stony Brook, NY 11790, USA. ²Department of Computer Science, Stony Brook University, Stony Brook, NY 11790, USA. ³Department of Civil and Environmental Engineering, Stanford University, Stanford, CA 94305, USA. ⁴U.S. Geological Survey, Golden, CO 80401, USA. ✉e-mail: susu.xu@stonybrook.edu

uncertainties in data and models, event-to-event variability, and co-occurrence of secondary hazards and impacts. One main difficulty lies in complex causal dependencies underlying the process of earthquake-triggered multi-hazard and impacts. To provide regional earthquake-induced ground failure estimation, previous data-driven approaches^{7,8} utilize statistical or black-box machine learning models to directly fit the complex mapping from designated features to a single hazard to historical inventories. Many of these methods are constrained by the limited availability of ground truth data as well as outdated and coarse-grained geospatial proxies (around 230 m) that entail large uncertainties and low spatial resolutions. Regarding rapid estimations of seismic impacts (e.g., building damage) in a large scale, existing systems often apply fragility functions from a pre-built database (e.g., PAGER^{9,10} and HAZUS¹¹). However, it is difficult to gain accurate predictions by applying fragility functions to given ground motion maps. This is mainly because there is no accurate local field information available, such as building type, for most seismic zones and many developing countries that do not have a complete fragility function database. Moreover, these methods mainly focus on modeling a single type of hazard without considering the causal interdependencies among multiple hazards and consequent impacts such as building damage, which led to limited performance and difficulty in physical interpretations.

In recent years, remote sensing observations, such as synthetic aperture radar (SAR) images, have become available within hours to days of an event. SAR can penetrate cloud cover compared to optical satellite images, and provide more complete spatial coverage and higher resolution information about seismic hazards or building damage in various weather conditions^{12,13}. Most notably, damage proxy maps (DPMs), developed by researchers of the NASA Advanced Rapid Imaging and Analysis (ARIA) team, extract correlation changes between pre- and post-event images, showing the potential to provide supplementary information for rapid seismic hazard and impact estimation¹⁴. Recent studies have explored integrating DPMs or other InSAR images with prior geospatial models to estimate or analyze single-type hazard^{15–17} or building damage^{18,19}. However, these models (1) all focus on single-type hazards, and (2) they are often trained in a supervised way and need a reasonable amount of ground truth labels for model training—which is not immediately available after earthquakes. One of the key challenges associated with seismic hazard and impact estimation from these remote sensing data is spatially overlapping multiple cascading hazards, building damage, and a large amount of event-specific environmental noise (e.g., changes induced by agricultural activities). This challenge made it difficult for current studies to effectively extract and utilize the information in DPMs for a fast and accurate earthquake-triggered landslide or building damage estimation, let alone predicting multiple hazards and impact simultaneously¹⁸. For example, direct correlation analysis shows a weak correlation between DPMs and building damage in the 2015 Nepal earthquake¹⁸. Without considering physical causalities in seismic hazard processes, it is difficult to categorize the types of changes observed from these imagery data, such as ground failures, building damage, and noise from vegetation growth and anthropogenic modifications, especially when these changes co-occur.

To this end, we introduce a new rapid seismic multi-hazard estimation system that effectively integrates single-type hazard geospatial models with rich but noisy information from DPMs, through quantitatively modeling causal dependencies underlying the geological processes and ground surface changes in satellite images. This causal inference model enables the system to jointly estimate multiple hazards and impacts, which then improves the overall accuracy and resolution in regional-scale estimation. We encode the causal dependencies among ground shaking, seismic hazards and impact, as well as ground surface changes captured by satellite images in a causal graph. The causal graph connects and fuses the rich information of event-

specific surface change patterns in high-resolution DPMs with event-shared physical insights in geospatial models, through explicitly modeling complex and nonlinear relationships among multiple hazards, building damage, environmental factors, and DPMs. The incorporation of causal dependencies enables comprehensive physical reasoning of the image changes in DPMs as the co-occurrence of multiple seismic hazards and building damage.

Results

Causal Bayesian network for modeling seismic multi-hazard and impacts

The causal graph is formulated as a Bayesian network for encoding a set of conditional dependency relationships, i.e., physical causal relationships, among multiple random variables. Notably, recent advances have shown that causal Bayesian network is a powerful tool for deciphering complex causation among multiple variables from a group of data^{20,21}. In our problem, the random variables include unobserved seismic hazards and environmental noise, observed satellite images product (DPMs), each hazard's geospatial model output, and ground shaking, as well as bias and uncertainties contained in all these observations. The random variables are represented as nodes, while their causal relationships are modeled through edges with directions. An example of our causal graph is shown in Fig. 1. Our model takes prior geospatial models, building footprints, and Damage Proxy Maps for the given earthquake as inputs. The final output of our model includes landslide, liquefaction, and building damage probability maps, as well as causal coefficients between them. The model is further trained in an unsupervised way through Bayesian updating to ensure the model generalization to different countries and regions with a variety of data availability.

The main challenges associated with seismic multi-hazard estimation through the causal Bayesian network are that (1) the ground truth data of multi-hazard and impact are typically not available within weeks to months after an earthquake event; (2) the causal dependencies between multiple hazards, impacts, and remote sensing observations are complex and thus hard to quantify; and (3) the introduced causal Bayesian network has a unique mixture of multiple layers with multiple forms of variables, including binary, log-normal, and logit-normal variables, which makes previous Bayesian inference algorithms no longer applicable. Consequently, it is difficult, if not impossible, to infer the posterior distributions of multiple unobserved variables, especially for events in which multiple seismic hazards and building damage are compounded in a complicated way.

To address these challenges, we derive a novel variational inference algorithm to estimate the posterior distributions of unobserved intermediate seismic hazards, building damage, and their quantitative causal dependencies. Conventional variational inference approximates the posterior distributions of unobserved variables by maximizing the marginal likelihood of observed variables. This has been proven to be a powerful statistical machine learning technique for inferring complex probabilistic models with many unobserved variables²². However, existing inference methods assume that intermediate “causal” variables are conditionally independent, which our causal Bayesian network structure violates. For example, both landslide and building damage can cause changes in DPMs, but they are not conditionally independent given DPMs because landslides might be one of the major causes of building damage.

To develop a new variational inference algorithm that works for our problem, we first approximate the likelihood of observed variables by factorizing over the Bayesian network with posteriors of unobserved variables (see Methods). We then derive the lower bound for this approximated likelihood, referred to as variational lower bound, that provides a rigorous theoretical guarantee for the optimality of the joint posterior inference (see Methods). An expectation-maximization algorithm alternatively optimizes the posterior of location-specific

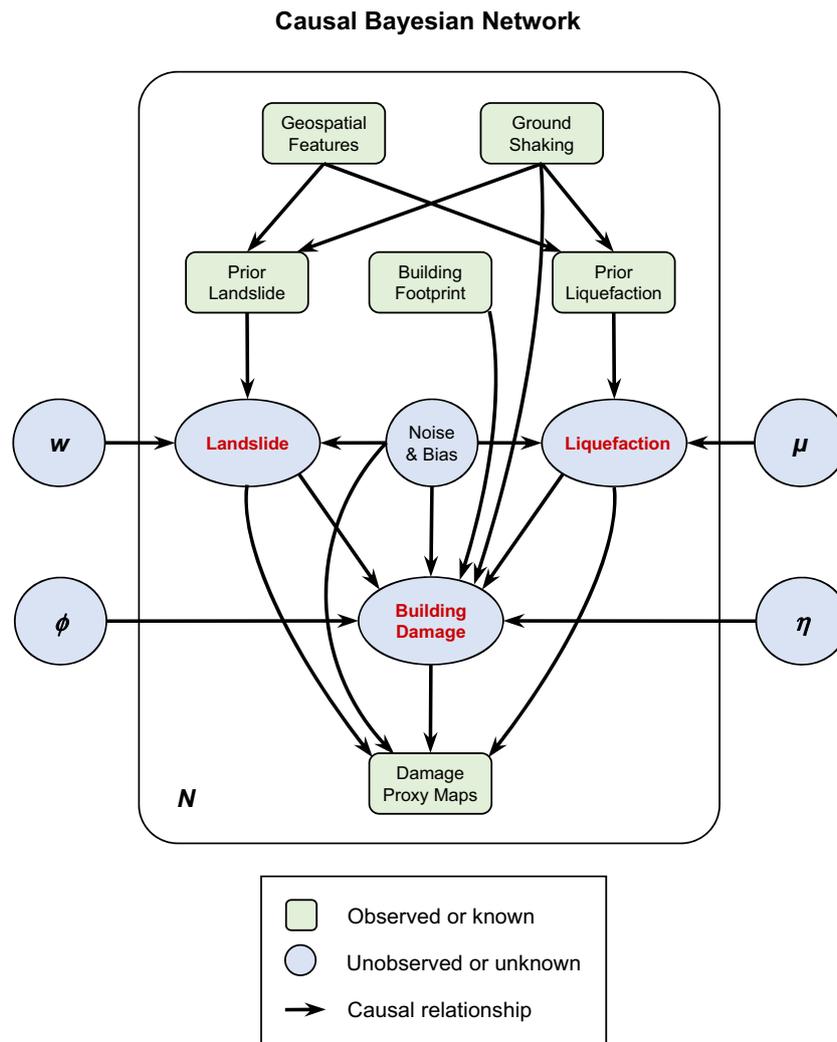


Fig. 1 | Causal Bayesian network that depicts causal dependencies among different seismic ground failures, building damage, Damage proxy maps (DPMs), and environmental noises. N refers to N locations/pixels in the target area. The posterior probability of landslide, liquefaction, and building damage at each location are the objectives that our system estimates. Green boxes refer to the variables

that have data constraints. Blue circles refer to nodes that are not observed or unknown. θ , ϕ , μ , η are the causal coefficients, which are unknown, that quantify the causal effects of parent nodes to landslide, liquefaction, building damage, and DPMs.

multi-hazard and impact together with causal dependencies among them by maximizing the variational bound (see Methods). Furthermore, the stochastic gradient descent and local pruning strategy ensure the scalability and computational efficiency for high-resolution regional multi-hazard and impact estimation. Our model is flexible to incorporate available building footprints, such as those developed by OpenStreetMap²³ or Microsoft²⁴, to further enhance the inference performance over the entire graph by setting the posterior of building damage in those areas without any building footprints to zero. The variational lower bound is generalizable to diverse applications involving discovering causalities among compound hazards.

More importantly, a wealth of experience has been accumulated in recent years in understanding the cascading hazards induced by earthquakes, but our knowledge is still fragmentary. Previously, the causal coefficients and posterior probability could not be learned/optimized simultaneously, owing to the complex causal dependencies. Therefore, most causal coefficients are recovered from correlation analysis using ground truth data of two related variables (e.g., landslide and building damage severity) with the underlying causal direction assumptions^{25–27} (e.g., from landslide to building damage). However, these correlation analyses are often inaccurate because they do not

decouple the direct impacts and indirect impacts from shared parents and noise. By jointly optimizing the quantitative causal dependencies over the entire causal Bayesian network, our system provides a systematic way to reveal insights into the event-specific causal mechanisms of compound seismic hazards.

Evaluation of seismic multi-hazard estimation over multiple seismic events

We evaluate the performance of our rapid multi-hazard and impact estimation strategy using four earthquakes in Puerto Rico (2020)²⁸, Hokkaido, Japan (2018)²⁹, Ridgecrest, California (2019)³⁰, and Central Italy (2016)³¹. These earthquakes occurred in regions with a variety of building structures and seismic, geological, and environmental characteristics. NASA ARIA team generated DPMs using SAR images captured by Japan Aerospace Exploration Agency's ALOS-2 satellites and European Space Agency's Sentinel-1 satellites. These DPMs are available within 6–14 days after the earthquakes. We evaluate the predictions of our system using the ground truth data collected by post-event reconnaissance groups, and compare the prediction performance with existing USGS earthquake ground failure products^{7,32} (prior models, see Methods), a regression model that only uses DPMs

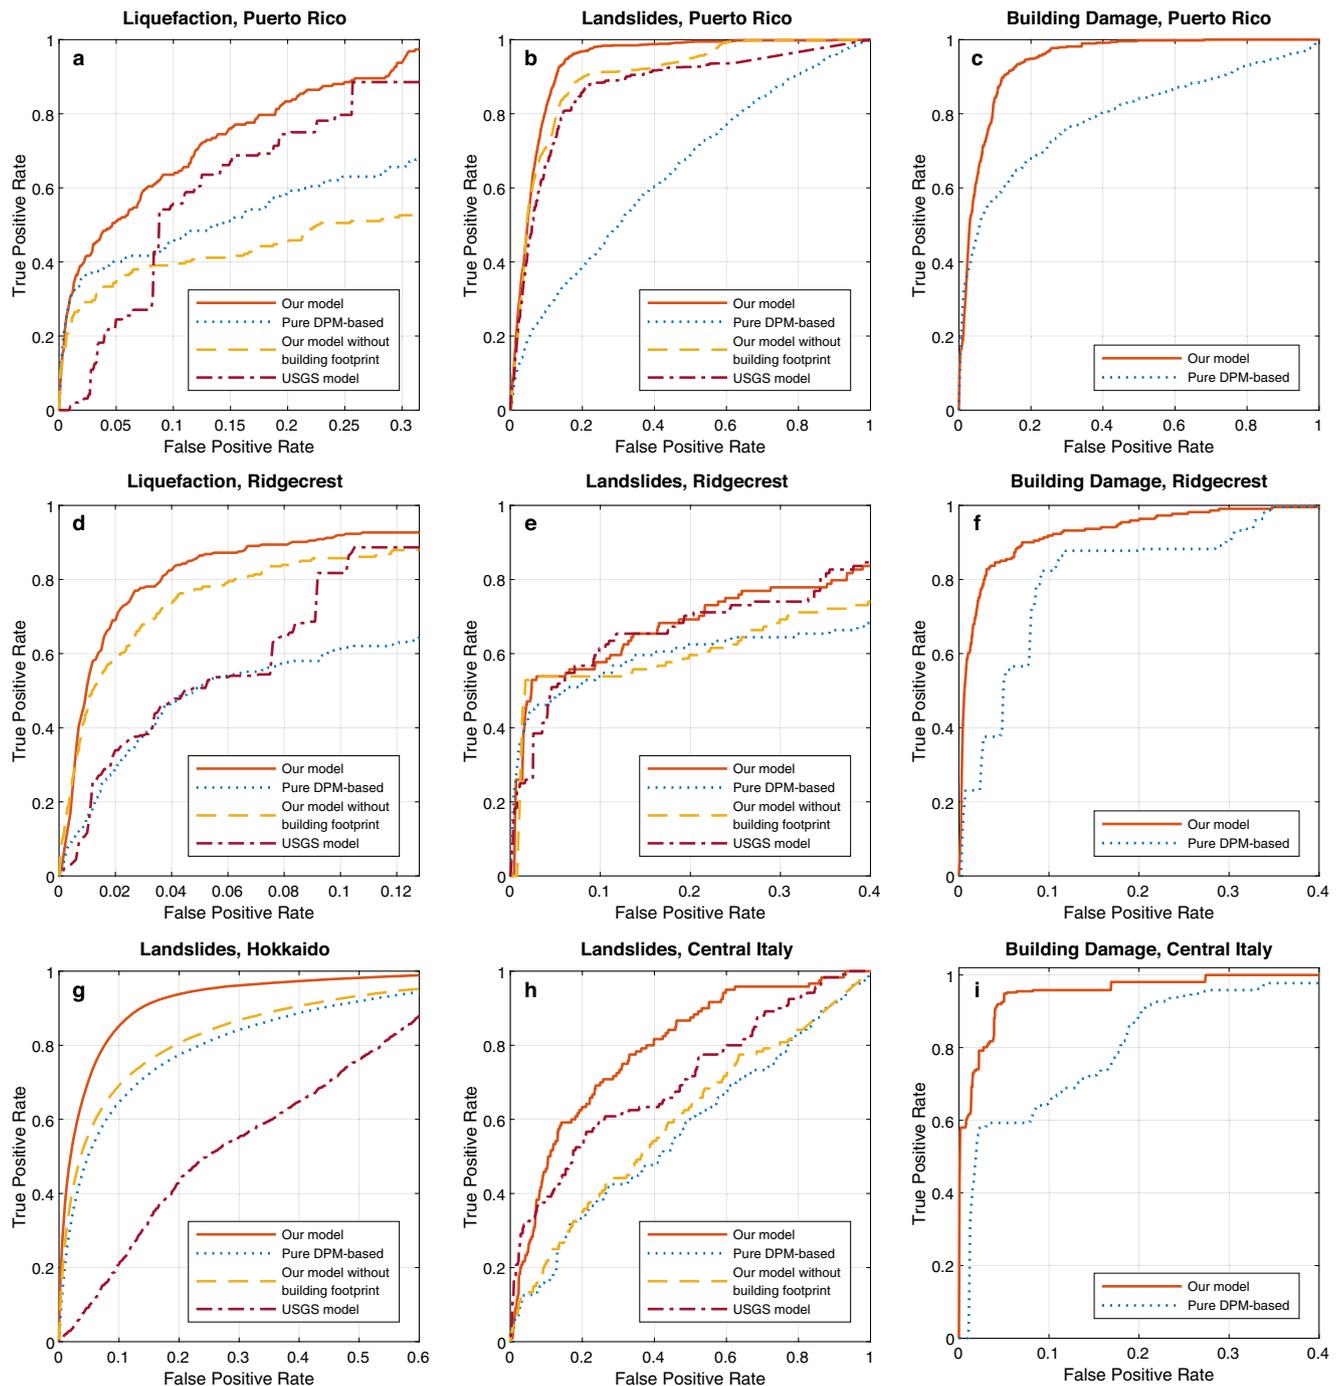


Fig. 2 | Comparisons of receiver operating characteristic (ROC) curves of four different earthquake events. a–c Liquefaction, landslides, and building damage ROC curves of the 2020 Puerto Rico earthquake. **d–f** Liquefaction, landslides, and

building damage ROC curves of the 2019 Ridgecrest earthquake. **g** Landslides ROC curve of the 2018 Hokkaido, Japan earthquake. **h, i** Landslides and building damage ROC curve of the 2016 Central Italy earthquake.

for ground failure/damage estimation in four different earthquakes, and our model without using building footprint. The ground truth availability is presented in Supplementary Table 3. Furthermore, we show that the introduced system reveals global and regional causality patterns among compound hazards and building damage in post-earthquake settings.

We summarize the performance of our system and baseline methods using the receiver operating characteristic (ROC) curve, which visualizes how the true positive rate varies with false positive rate (see Methods). We also calculate the area under the ROC curve (AUC)³³, representing the ability of a system to distinguish the positive

and negative cases, and cross-entropy loss³⁴, measuring the distributional dissimilarity between predictions and ground truth data (see Methods). Both evaluation metrics are used to evaluate the disaster impact estimation system's ability to identify ground failures and to accurately estimate their probabilities compared to the binary ground truth. Figure 2 presents the performance evaluation results, including the ROC curves of our system (orange solid line), existing USGS model (red dotted line), pure DPM-based indicator (blue dotted line), and our model without using building footprint (yellow dotted line). In Fig. 2a, d, we compare the binary liquefaction estimation performance of our system with baseline methods for the Puerto Rico and the Ridgecrest

earthquakes, respectively. These two earthquakes induced substantial liquefaction and subsequent building damage in the target estimation areas. Figure 2b, e, g, h compares the landslides occurrence estimation performance of our system with baseline methods for the Puerto Rico earthquake, the Ridgecrest earthquake, the Hokkaido earthquake, and the Central Italy earthquake, respectively. Figure 2c, f, i presents the ROC curves for binary building damage estimates (no/slight damage versus moderate/severe/collapse damage) in the Puerto Rico earthquake, the Ridgecrest earthquake, and the Central Italy earthquake. The Hokkaido earthquake does not have building damage ground truth available.

As Fig. 2 shows, our system significantly improves the landslide and liquefaction prediction ability compared to the USGS models. It can be seen that our model effectively integrates the information from building footprints, prior models, and DPMs, and significantly improves the ground failure estimation performance compared to traditional geospatial models and pure DPMs-based models. One most recent study explored updating the empirical landslide models progressively using landslide indicators derived from InSAR images to this model in the order they were acquired following the earthquake¹⁵. But when comparing the landslide estimation performance in the 2018 Hokkaido earthquake, our model achieves the AUC value of 0.93, which significantly improves the landslide estimation performance compared to their DPM-based model¹⁵ which achieves 0.78 AUC. The comparison also shows that the integration of building footprints contributes to the performance improvement of our system significantly, especially for predicting landslides in the 2016 Central Italy earthquake and predicting liquefaction in the 2020 Puerto Rico earthquake. This is mainly because they mainly occurred in the areas with many buildings. The building footprint information directly improves building damage estimation from DPMs, and indirectly improves liquefaction and landslide estimations. Our model significantly improves the predictions compared to USGS models or DPMs alone. In our system, the building footprint provides information about buildings' existence in different locations, to further reduce the uncertainty in the process of attributing DPM data to landslide/liquefaction/building damage.

We also show that even without the prior fragility function, the model can still provide meaningful predictions about the building damage distribution. We compare the performance for building damage detection (no/slight building damage versus moderate/severe/collapse building damage), between our model and the pure DPM-based method in Fig. 2c, f, i. The results show that our system improves the binary building damage detection accuracy by up to 16.92%. Another recent work G-DIF proposes to train a kriging regressor that combines geospatial models with InSAR images for building damage estimation using inventory data and ground truth, but the supervised method performance is constrained on ground truth data availability¹⁸. Our results show that by integrating information from building footprints, DPMs, and causal ground failure factors, it is possible to estimate large-scale building damage distributions in an accurate and rapid manner, even without prior fragility function and ground truth data.

More detailed information about the ground truth data availability can be found in Supplementary Table 3. Note that the ultimate goal of performance comparison is to show our model's capability to improve the estimation performance of multi-hazards and damage by jointly modeling them and integrating geospatial models with DPMs. Therefore, even though there are only one or two types of ground truth observations, the performance improvement can still indicate the capability of our model. In the meantime, in most cases, if the ground truth data of one ground failure is not available, it usually indicates that that ground failure is not very critical so the survey agency deprioritizes the investigation of that ground failure. For example, in the Hokkaido earthquake, although liquefaction ground

truth is not available, we still show that by integrating multiple hazards and damage with DPMs, we can significantly improve the performance of landslide prediction—which is the most critical ground failure in the covered area³⁵.

In the following sections, we present detailed results in the 2020 Puerto Rico earthquake and the 2018 Hokkaido earthquake, which have different ground failure and damage patterns. The Puerto Rico earthquake-induced landslides and severe liquefaction, resulting in damage to residential and commercial structures and human fatalities, according to the reconnaissance reports from multiple groups³⁶. The Hokkaido earthquake mainly triggered numerous shallow landslides, especially in the towns of Atsuma and Abira, and subsequent building damage^{37,38}. We present the detailed evaluation results of our system on these two distinct types of earthquake events. More evaluation results about the 2019 Ridgecrest earthquake and the 2016 Central Italy earthquake are presented in the “Methods” section.

The 2020 Puerto Rico Earthquake. On January 7, 2020, at 4:24 a.m. (AST), a magnitude 6.4 earthquake shook the southwest part of Puerto Rico²⁸. Following the disaster, post-disaster reconnaissance efforts revealed widespread casualties with over 775 affected buildings³⁹ and 800 ground failure observations^{40,41}. To locate potentially damaged areas, the ARIA team created DPMs using SAR images from the Sentinel-1 satellite⁴². Ground truth observations were later collected in field reconnaissance conducted by researchers from the USGS, the University of Puerto Rico Mayagüez, the GEER team and the Structural Extreme Events Reconnaissance team^{36,43,44}.

We compared our system performance with benchmark methods for estimating ground failures and building damage occurrence in the Puerto Rico earthquake in Figs. 2a–c and 3. The AUC for our liquefaction ROC curve is improved by 15.64% and cross-entropy loss is reduced by 52.45% compared to the existing model. In landslide estimation, our model improves the AUC by 7.40% and reduces the cross-entropy loss by 51.57%. In building damage estimation, our model improves the AUC by 33.14% and reduces the cross-entropy loss by 44.34%. As an example, when our system performance converges as false positive rate increases to 31.80%, our system achieved a 97.40% true positive rate for liquefaction estimates compared to an 88.54% true positive rate of the USGS liquefaction model, as shown in Fig. 2a. Figure 3 visualizes our model performance and ground truth observations in the shaken region for joint landslide (Fig. 3c), liquefaction (Fig. 3e), and building damage estimation (Fig. 3g), compared with the prior landslide model (Fig. 3b) and prior liquefaction model (Fig. 3d), and building footprint (Fig. 3f). In Fig. 3a, the rectangle (i) covers the western areas of Ponce, which is the second largest city in Puerto Rico and suffered from severe liquefaction and lateral spreading during the earthquake³⁶ (magenta squares in Fig. 3d, e). Our system predicts a high probability of liquefaction in these areas (Fig. 3e(i)) compared to the existing USGS liquefaction model (Fig. 3d(i)), showing that integrating DPMs through our causal Bayesian network helped improve the accuracy of rapid liquefaction identification.

The 2018 Hokkaido, Japan Earthquake. On September 6, 2018, at 3:08 a.m. (JST), two days after Typhoon Jebi made landfall with heavy rain, a Mw 6.6 earthquake caused catastrophic landslides in the southern area of Hokkaido, Japan^{29,45}. Using the SAR images from the Japan Aerospace Exploration Agency's ALOS-2 satellites, the ARIA team produced DPMs⁴⁶ that includes the town of Atsuma, which was located near the sites of large-scale landslides⁴⁷. Figure 4a presents the target seismic zone and corresponding DPM. This DPM covered the towns of Atsuma and Abira. Ground truth observations of landslides were later collected by the Geospatial Institute of Japan⁴⁷. The liquefaction investigation and building damage reports do not cover the extent of

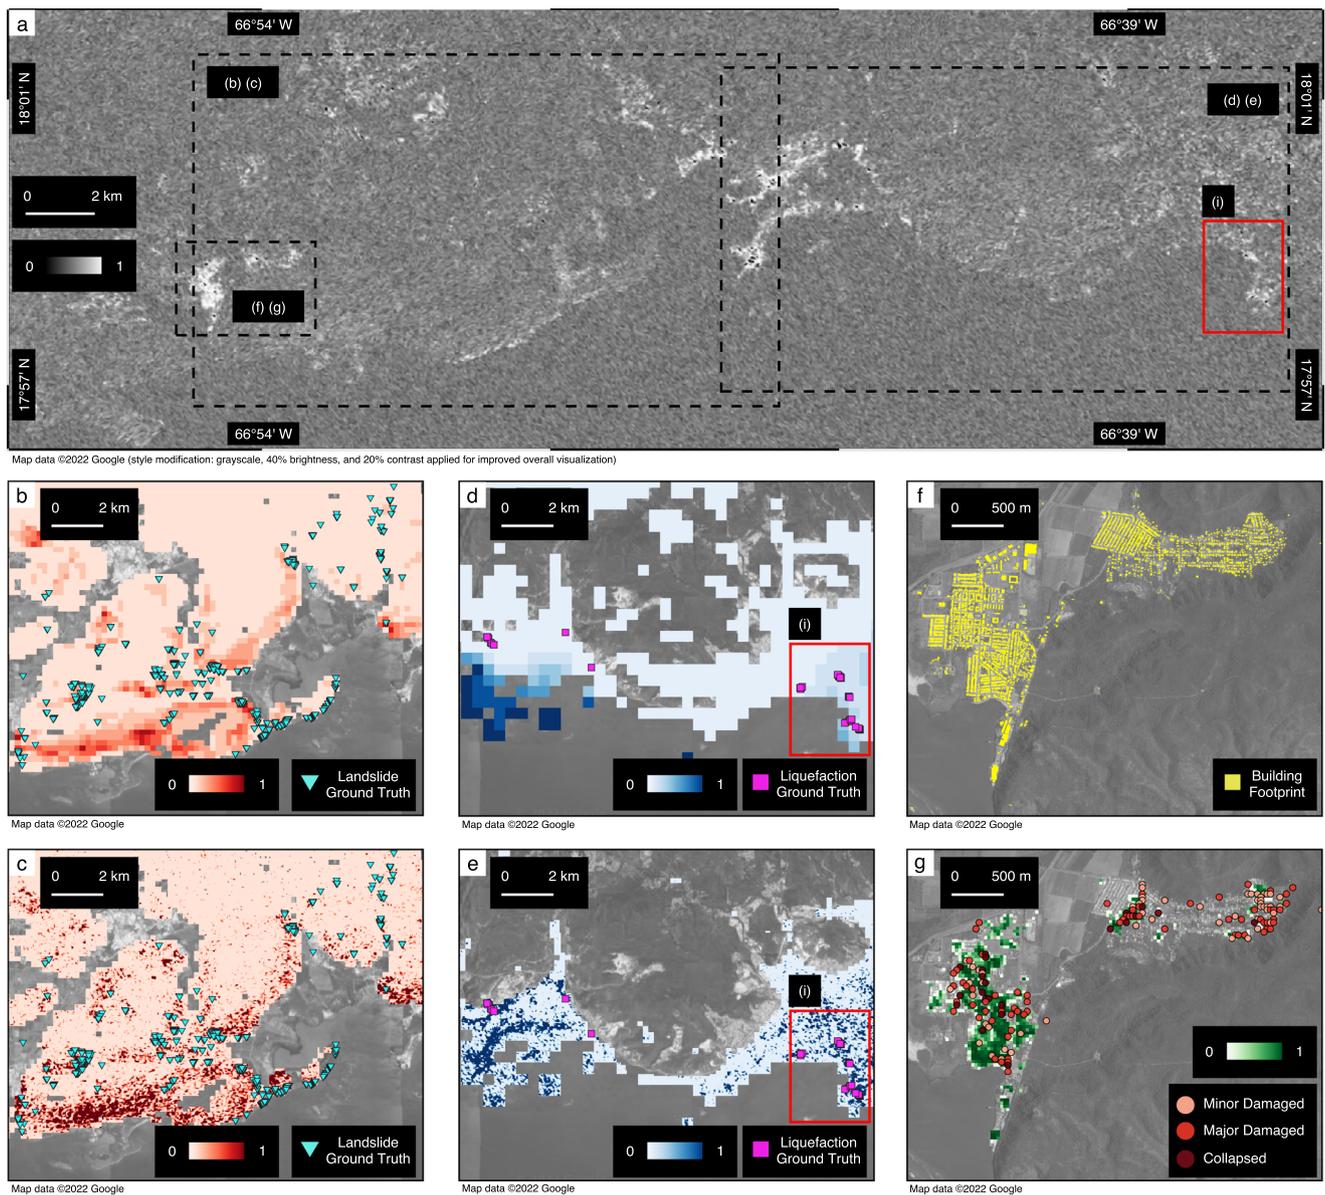


Fig. 3 | Geospatial prior and posterior estimation models of the 2020 Puerto Rico earthquake. **a** Damage proxy map (30-m resolution) with rectangular extents of each focused area. **b** Prior landslide estimates in red and ground truth (triangles). **c** Posterior landslide estimates in red and ground truth (triangles). **d** Prior

liquefaction estimates in blue and ground truth observations (squares). **e** Posterior liquefaction estimates in blue and ground truth (squares). **f** Building footprint map. **g** Posterior building damage estimates in green and ground truth (rounded points). Rectangle **i** highlights the area of western Ponce.

the DPM. With the current ground failure hazard system of the USGS, the prior estimations of landslide and liquefaction probabilities were generated using the corresponding ShakeMap in Atlas V4²⁹.

We present our model performance for joint landslide, liquefaction, and building damage estimation in Fig. 4, compared with the USGS landslide models and USGS liquefaction models, as well as ground truth inventories^{7,23,24,32,48}. The results show that our landslide model (Fig. 4c) resembles the spatial distribution of ground truth observations (Fig. 4e) more accurately than the USGS models (Fig. 4a). The posterior model also has a higher resolution (30 m) than the USGS model (230 m)²⁹. Our model, by integrating the DPM and prior models (Fig. 4a), identifies more landslide occurrences that are later validated by ground truth observations, compared to existing USGS models. The AUC is improved by 42.02%, achieving 0.9331, and the cross-entropy is reduced by 47.75%. In addition, the liquefaction and landslide estimation map resolutions are improved by 7.7 and 15 times, respectively. The results demonstrate that the joint estimation of landslide and

liquefaction using high-resolution satellite images significantly improves the landslide estimates.

Revealed causality relationships

The inferred causal graphs for each earthquake site revealed the causal dependencies between ground shaking, landslide, liquefaction, building damage, and environmental noise. In Puerto Rico and Ridgecrest, California (see Methods), the building damage is mainly caused by liquefaction rather than landslides. The quantified causal coefficient from liquefaction to building damage for Puerto Rico and Ridgecrest, California are 1.718 (± 0.116) and 2.059 (± 0.151), respectively, which is significantly higher than the coefficient from landslide to building damage (0.527 ± 0.305 and 1.560 ± 0.374 , respectively). In the regions covered by DPMs in the Hokkaido, Japan earthquake and the Central Italy earthquake (see Methods), the building damage is mainly dominated by landslides rather than liquefaction. The quantified causal coefficient from landslides to building damage for the

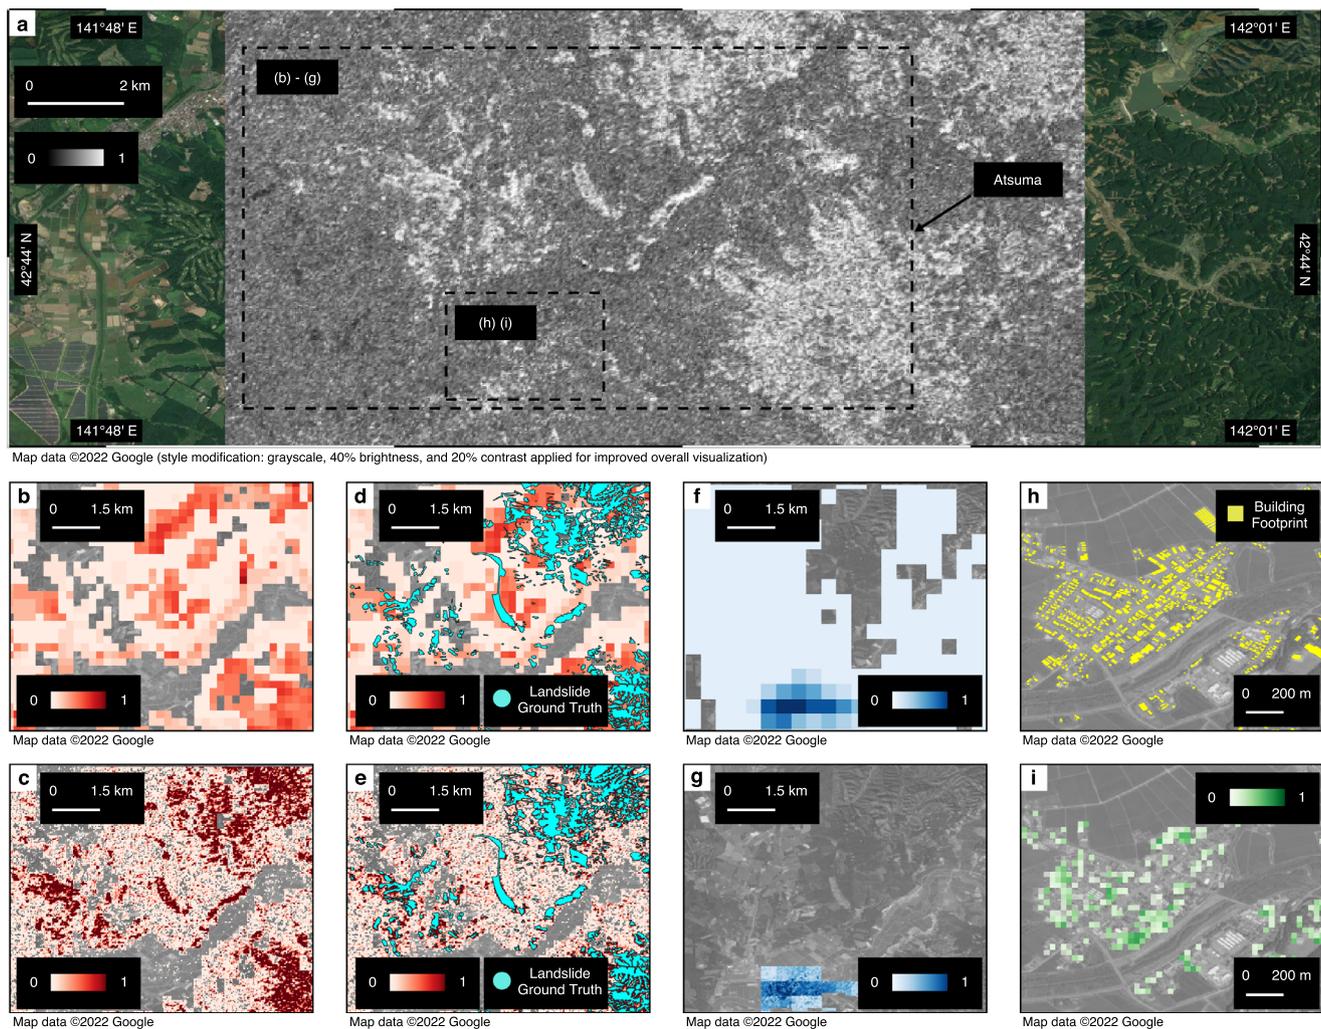


Fig. 4 | Geospatial prior and posterior estimation models of the 2018 Hokkaido, Japan earthquake. **a** Damage proxy map (30-m resolution) with rectangular extents of each focused area. **b** Prior landslide estimates. **c** Our landslide estimates. **d** Prior landslide estimates (in red) with ground truth observations (in blue

polygon). **e** Our landslide estimates (in red) with ground truth observations (in blue polygon). **f** Prior liquefaction estimates. **g** Our liquefaction estimates. **h** Building footprint map. **i** Our building damage estimates. The quantitative performance comparison results could be found in Fig. 2.

Hokkaido, Japan earthquake and the Central Italy earthquake are 0.848 (± 0.128) and 0.966 (± 0.575), respectively, which are significantly higher than the coefficient from liquefaction to building damage (0.432 ± 0.161 and 0.345 ± 0.157) in pairwise comparisons. We further enhance the causal graph using the absolute value of coefficients as the linewidth to depict different importance levels in the causal relationships in Fig. 5a, b. This enhanced causal graph visualization reveals that most building damages are directly dominated by the landslide and seismic ground shaking in the 2018 Hokkaido, Japan earthquake, whereas by liquefaction in the 2020 Puerto Rico earthquake. Moreover, Fig. 5 shows that the weight ratio of bias to landslide divided by prior landslide to landslide is higher in the Hokkaido event compared to the Puerto Rico event. This indicates that the prior landslide model in the Puerto Rico event contributes more information to the posterior landslide estimation than the noise/bias term, compared to Hokkaido event, which could be further validated by comparing the AUC of ROC curves of prior landslide models for the Hokkaido event (0.69) and the Puerto Rico event (0.90).

The quantified causal relationships, in conjunction with spatial distribution of buildings, determine the building damage patterns in seismic zones. By investigating the spatial distributions of building footprints, we found that in Ridgecrest, California, more than 80.17% of buildings are located in liquefaction zones, whereas

in Puerto Rico, Hokkaido, and Central Italy, 61.87%, 68.39%, and 90.25% of buildings are located in the potential landslide zones, respectively. Interestingly, in the Puerto Rico earthquake, though most buildings are located in the landslide zones, the major building damage are caused by liquefaction. From the distributions in Fig. 5d, we found that building damage distribution in liquefaction zones are skewed towards higher probability than in landslide zones, which indicates more severe building damage in liquefaction zones. Compared to Puerto Rico, as shown by Fig. 5c, in Hokkaido most severe building damage happened in the landslide zones, which matches our final causal graph that landslide is the major cause of building damage in the towns of Atsuma and Abira³⁷. Moreover, the revealed causalities in the 2019 Ridgecrest, California earthquake are consistent with the ground truth observations³. The significant effects of liquefaction on building damage are also shown in Trona-Argus area of Ridgecrest where heavily damaged buildings were located on or near the alluvial fan and exhibited evidence of lateral spreading in the form of tension cracks. In addition, the building damage observed in Argus had the form of chimney separations from buildings, masonry cracks from extensional movements, toppled walls, and punch-through in retention walls, which are commonly observed types of damage induced by liquefaction³.

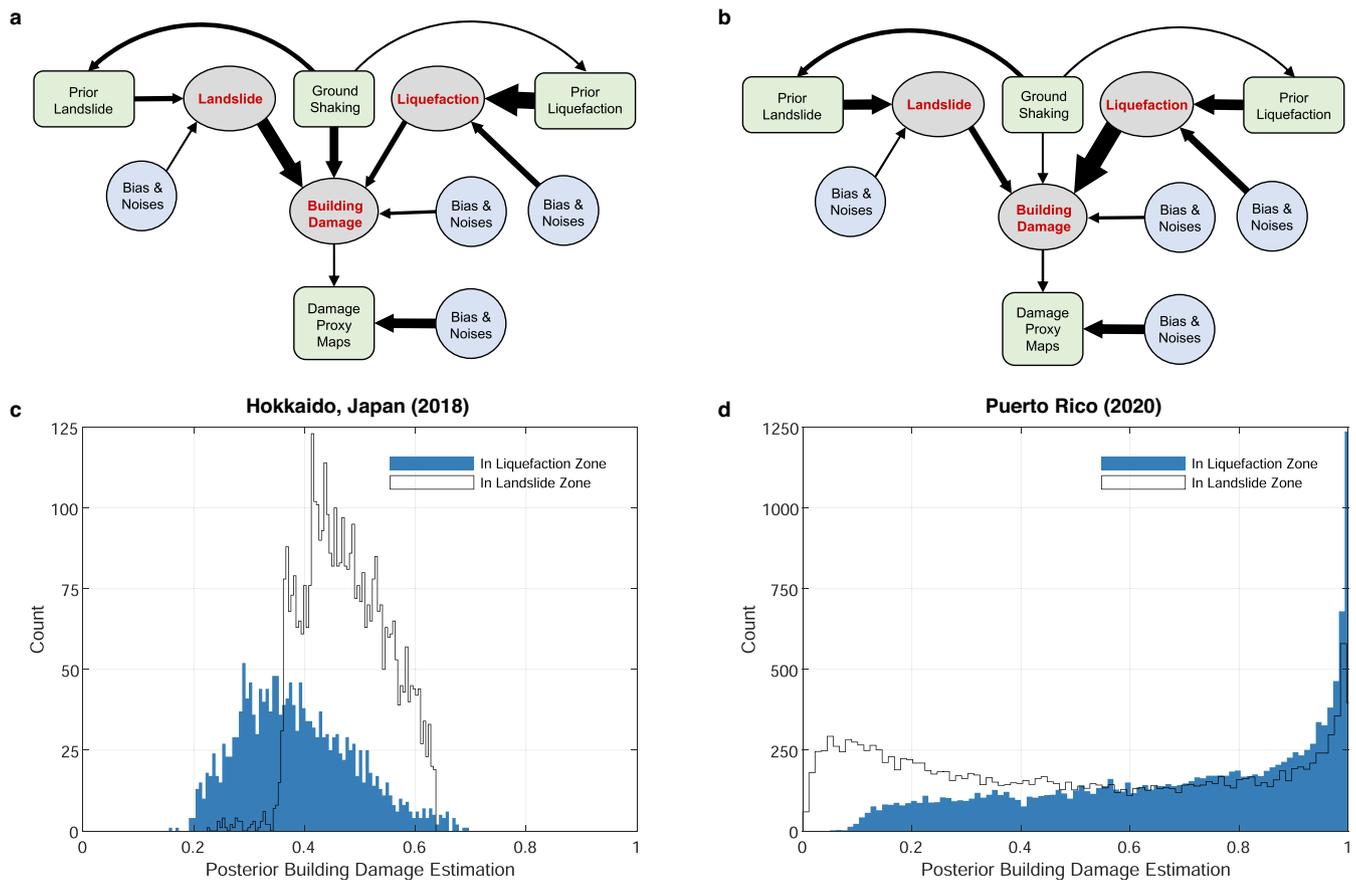


Fig. 5 | Analysis of causal mechanisms revealed by the graph. **a** Visualization of quantified causal graph for the 2018 Hokkaido, Japan earthquake. The thickness of each edge represents quantified importance of causal effect over the edge. **b** Visualization of quantified causal graph for the 2020 Puerto Rico earthquake. The thickness of each edge represents quantified importance of causal effect over the

edge. **c** Estimated building damage probability distribution under the landslide zone (transparent white) and the liquefaction zone (colored blue) for the 2018 Hokkaido, Japan earthquake. **d** Estimated building damage probability distribution under the landslide zone (transparent white) and the liquefaction zone (colored blue) for the 2020 Puerto Rico earthquake.

The causal dependencies revealed in our system contribute to the quantitative explanation of the underlying mechanisms in cascading seismic hazards and impacts, providing a rapid and quantitative idea of potential dominant causations of building damage in different earthquake events from data perspectives. When modeling the causal graph, we incorporated prior understanding of causal dependencies in underlying hazards processes. For example, landslides and liquefaction may induce building damage. The probabilistic inference based on DPMs helps quantify this causal dependencies. This is similar as using causal Bayesian network quantify the causal dependencies between diseases and different signs and symptoms, which is commonly used for human disease modeling and diagnosis^{49,50}.

Discussion

In this study, we show that by jointly estimating multiple seismic hazards and losses considering their causal dependencies, the accuracy and resolution of rapid seismic multi-hazard and impact estimations are significantly improved. In addition, variational inference allows us to guarantee scalability and computational efficiency. This is a holistic work enabling joint estimation of earthquake-induced landslide, liquefaction, and building damage, without the need of any ground truth observations. We demonstrated that DPMs provide effective information related to various types of seismic hazards and impacts, and developed a systematic approach to extract and utilize such information. Meanwhile, causal Bayesian network we developed is a novel multi-layer network that can incorporate a mixture of

different data types, including binary log-normal, and logit-normal variables. We derived a tight variational lower bound for this type of Bayesian network, which can be further generalized to different applications involving discovering causalities among other compound hazard and impact distributions.

In addition, our system quantified the causal effects from multiple triggering factors to building damage. We found that the earthquake-induced building damage was mainly caused by landslides in the epicentral areas of 2018 Hokkaido, Japan earthquake and 2016 Central Italy earthquake, while liquefaction were the main contributors in the 2020 Puerto Rico earthquake and 2019 Ridgecrest earthquake. Among four earthquake events, ground shaking is always the most critical causal factor, directly or indirectly, on the occurrences of building damage. Furthermore, we inspected the estimates of building damage distributions in seismic zones with potential landslides and liquefaction separately. The building damage distribution patterns differ significantly, because the underlying earthquake-triggered geophysical impacts vary from event to event. The discovered causal dependencies need to be further combined with the domain knowledge and previous geological survey data to derive rigorous conclusions about local geohazards and impacts causal mechanisms. The geologists and engineers could use these quantitative causal measurements as auxiliary information to further guide their field survey and geological analysis.

The solution we developed establishes a new way to understand the complex interaction of multiple hazards from high-resolution satellite signals. The causal Bayesian network has the potential to be

applied to other extreme disasters such as hurricanes and flooding through modeling the causalities among their cascading hazards and impacts. One major advantage of our proposed causal Bayesian network is to formulate heterogeneous variables, such as hazards, impacts, and sensing data, into structured Bayesian representations and integrate them into one holistic model informed by physical causality. Therefore, it is flexible to add nodes to or delete nodes from the causal graph according to disaster types. Recently, emerging technology has enabled the use of satellite imagery to identify the conditions of remote areas not easily accessible by disaster responders. The revisit time frequency of satellites has been reduced in recent years from weeks to several days and even to hours⁵¹. Flexible SAR systems, such as airborne SAR⁵², are being deployed to acquire high-resolution images in a much more timely and controllable way. The developments in the SAR arena make the acquisition of large-scale high-resolution SAR images faster (within hours) and cheaper in the future⁵³, offering even greater potential for our computational strategies as these advancements come online. Depending on the timeline of emergency responses and rehabilitation of specific events, the outcome of our system can be combined with the initial loss estimation, loss updating, and/or rehabilitation planning tools to provide situational awareness to broader users in the critical early hours to days, prior to complete observations can be collected. Causality-informed multi-hazard and impact modeling like our variational Bayesian inference system will address the need for rapid and accurate disaster-induced multi-hazard and impact estimation, even without prior models for all subsequent hazards and impact.

Methods

Additional data

The 2016 Central Italy Earthquake. On August 24, 2016, at 3:36 a.m. (CEST), a Mw 6.2 earthquake resulted in over 298 deaths, 386 injured, 5000 displaced individuals, and many damaged buildings in the Arquata del Tronto-Accumoli-Amatrice area^{31,54}. To map the damage over the large affected area, the ARIA team produced DPMs⁵⁵ using the SAR images from the Italian Space Agency’s COSMO-SkyMed satellites and Japan Aerospace Exploration Agency’s ALOS-2 satellites within two days. The landslide inventories indicated a limited area exposed to landslides and no significant liquefaction^{41,48}.

The 2019 Ridgecrest Earthquake. On July 5, 2019, at 8:19 p.m. (PST), the largest event in the sequence, with a magnitude of 7.1, caused significant ground surface displacements⁵⁶ in the Indian Wells and Searles Valleys. In the towns of Ridgecrest and Trona³⁰ extensive field mapping revealed many damaged buildings, landslides, and liquefaction^{17,57–59}. After seven days, the ARIA team generated DPMs⁶⁰ using the Sentinel-1 SAR images to identify affected areas.

The joint natural hazards estimation workflow

The variational Bayesian causal inference system we introduced integrates geospatial models, multiple unobserved ground failures, unobserved building damage, Damage Proxy Maps (DPMs), and their causal dependencies for the joint estimation of multiple unobserved seismic hazards and impacts. Our system first formulates a causal Bayesian network based on the understanding of potential causal relationships underlying the geological processes, as shown in Fig. 1. In Fig. 1, all green rectangle nodes, including geospatial models, building footprints, and DPMs are already observed (known). Our objective is to use this information and this causal dependency relationships represented by the causal network to infer the unobserved variables—mainly landslide, liquefaction, and building damage maps, as well as the causal coefficients between them. Note that each edge in the graphical model

indicates a causal relationship but not the input-output data flow. A variational bound is further derived to approximate the intractable (i.e., cannot be directly modeled) posteriors of unobserved ground failure and building damage using the likelihood of observed DPMs. Finally, a stochastic coordinate descent method is employed to find the optimal coefficients representing the causal dependencies among different predictors, ground failure, building damage, environmental and anthropogenic noise, and DPMs. The resulting combinations of causal coefficients and posteriors of unobserved variables are near-optimal to maximize the marginal likelihood of observed DPMs given prior geospatial models. Our system is built upon this causal graphical model to output the probability of intermediate hazards and impacts at each high-resolution grid cell using the geospatial models and DPMs.

Causal graph-based Bayesian network. We model a causal graph-based Bayesian network to represent the causal dependencies among landslide (LS), liquefaction (LF), building damage (BD), and DPMs, as shown in Supplementary Table 1. We denote each location in the shaken area as l . y^l refers to the DPM pixel captured at l . Given a node i , we use $\mathcal{P}(i)$ to define the parents of i . For notational simplicity, we define unobserved seismic ground failure and impact nodes as x_i , where $i \in \{1, 2, 3\}$. x_1 refers to landslide (LS), x_2 refers to liquefaction (LF), and x_3 refers to building damage (BD). These unobserved nodes have binary variables $x_i \in \{0, 1\}$. We define a bias node x_0 , which is always active ($x_0 = 1$), which allows its child nodes to be active even when other parent nodes are inactive. For example, even though landslide and liquefaction do not occur, it is still possible to have building damage. All nodes are linked by an arbitrary directed acyclic graph as shown in Supplementary Fig. 1. We further give quantitative definitions of these links, i.e., causal dependencies, among different random variables. To simplify the notation, we define and w_{ki} (k is any parent node of i) to quantify the causal effects from the noise term to a node i as well as from a parent node k to a child node i . When i refers to LS, LF, BD, DPM, there is $w_{ki} = \theta_k, \mu_k, \phi_k, \eta_k$, respectively. All above weight nodes w_{ki} are formulated as deterministic variables. We assume the formulation of underlying causal dependencies from parent nodes to unobserved x_i and DPM (y) as follows:

$$\log \frac{p(x_i = 1 | x_{\mathcal{P}(i)}, \epsilon_i)}{1 - p(x_i = 1 | x_{\mathcal{P}(i)}, \epsilon_i)} = w_{0i} + w_{\epsilon_i} \epsilon_i + \sum_{k \in \mathcal{P}(i)} w_{ky} x_k \quad (1)$$

$$\log y = w_{\epsilon} \epsilon_y + w_{0y} + \sum_{i \in \mathcal{P}(y)} w_{iy} x_i \quad (2)$$

As ϵ_y is a normal distribution, $y | \mathcal{P}(y)$ is subject to a log-normal distribution.

The above logit relationship between LS/LF/BD and their parent nodes follows the assumption of a logistic regression model, which is used in the statistical models of LS and LF used by the USGS^{7,32}. With the above distribution and conditional distribution assumptions, we construct a Bayesian network based on the causal graph that effectively captures the dependencies between different ground failure types, building damage, and remote sensing observations. This network contains multiple unobserved random variables and unknown causal dependencies between random variables. The complex causal dependencies make the posterior of unobserved random variables intractable. Therefore, we develop a stochastic variational inference method to approximate the intractable posterior of unobserved ground failure and building damage.

Variational causal Bayesian inference. With the Bayesian network constructed, we further infer the posterior of ground failure and building damage. However, the problem is that (1) both causal dependencies and distributions of ground failures and building damages are unknown, and (2) the target affected area could be large, which makes it computationally expensive to jointly update the posteriors of ground failure/building damage over the entire map with the size ranging from 12.72 km × 17.34 km (2018 Hokkaido, Japan earthquake) to 78.15 km × 177.06 km (2019 Ridgecrest, CA earthquake). To address this problem, the variational inference is introduced to factorize the Bayesian network first and then approximate the posterior distributions of unobserved variables through maximizing the log-likelihood of observed variables. To ensure the scalability of our method, the variational inference is conducted on a small batch of randomly sampled locations in each iteration. For each location l , we define a variational distribution $q(X^l)$, which further factorizes over unobserved nodes:

$$q(X^l) = \prod_i q(x_i^l) = \prod_i (q_i^l)^{x_i^l} (1 - q_i^l)^{1-x_i^l}. \tag{3}$$

Given a map containing a set of locations N , we further derive a tight lower bound for the log-likelihood as follows:

$$\begin{aligned} \log P(Y) &= \sum_{l \in N} \log P(y^l) \\ &\geq \sum_{l \in N} \left\{ -\log y^l - \log w_{\epsilon_y} - \frac{(\log y^l)^2 + w_{0y}^2 + \sum_{k \in X_{py}^l} w_{ky}^2 q_k^l}{2w_{\epsilon_y}^2} \right. \\ &\quad \left. - \frac{\sum_{i,j \in X_{py}^l} 2w_{iy}w_{jy}q_i^lq_j^l - 2w_{0y} \log y^l - 2(\log y^l)(\sum_{k \in X_{py}^l} w_{ky}q_k^l)}{2w_{\epsilon_y}^2} \right. \\ &\quad \left. - \frac{2w_{0y} \sum_{k \in X_{py}^l} w_{ky}q_k^l}{2w_{\epsilon_y}^2} \right. \\ &\quad \left. - \sum_{\substack{i \in \{LS, LF\} \\ \nu \in \{0,1\}}} (q_i^l)^\nu (1 - q_i^l)^{1-\nu} \log \left(1 + \exp \left((-1)^\nu (w_{0i} + w_{\alpha_i} + \frac{w_{\epsilon_i}^2}{2}) \right) \right) \right. \\ &\quad \left. - \sum_{\substack{\nu_i, \nu_j \in \{0,1\} \\ i \in \{BD\}, j \in \mathcal{P}(i)}} \log \left\{ 1 + \exp \left[(-1)^{\nu_i} \cdot \left(w_{0i} + \sum_{j \in \mathcal{P}(i)} \frac{1 - (-1)^{\nu_j}}{2} w_{ji} \right) + \frac{w_{\epsilon_i}^2}{2} \right] \right\} \right. \\ &\quad \left. \prod_{k \in \{i,j\}} (q_k^l)^{\nu_k} (1 - q_k^l)^{1-\nu_k} \right. \\ &\quad \left. - \sum_{i \in \{LS, LF, BD\}} [q_i^l \log q_i^l + (1 - q_i^l) \log (1 - q_i^l)] \right\}. \tag{4} \end{aligned}$$

With the tight lower bound of log-likelihood of DPM observations, we can further maximize the lower bound to find optimal posteriors of unobserved variables, i.e., LS, LF, and BD.

Our final objective is to maximize the bound to find optimal combinations of posteriors and causal dependencies estimations. As both posteriors and weights are unknown, we develop an expectation-maximization approach to achieve this. In the step of expectation, we derive closed-form update equations for local posteriors of LS, LF, and BD, i.e., $q_i^l, i \in \{1,2,3\}, l \in N$, by maximizing the lower bound and setting the gradients of the lower bound as 0, i.e., $\partial L(q^l, w) / \partial q_i^l = 0$. This gradient is obtained via the chain rule, and the optimal posterior follows the form below:

$$q_i^l = \frac{1}{1 + \exp \left(-T(q_{\mathcal{P}(i)}, q_{S(i, \mathcal{C}(i))}, q_{\mathcal{C}(i)}, y^l, w) \right)}, \tag{5}$$

where $\mathcal{P}(i)$ refers to the set of parent nodes, $\mathcal{C}(i)$ refers to the set of child nodes, and $S(i)$ refers to the set of spouse nodes that share the same child nodes with i . T is a nonlinear function that is determined by

the prior of i , weights of edges associated with node i , and posteriors of the parent, child, and spouse nodes of i as follows:

$$\begin{aligned} T(\cdot) &= \sum_{\substack{k \in \mathcal{P}(i) \\ \nu \in \{0,1\}}} \left\{ (-1)^\nu q_k f \left((-1)^\nu (w_{0i} + w_{ki}) + \frac{w_{\epsilon_i}^2}{2} \right) + (-1)^\nu (1 - q_k) f \left((-1)^\nu w_{0i} + \frac{w_{\epsilon_i}^2}{2} \right) \right\} \\ &+ \sum_{\substack{k \in \mathcal{C}(i) \\ k \neq y \\ \nu \in \{0,1\}}} \left\{ (-1)^\nu q_k f \left((-1)^\nu (w_{0i} + w_{ik}) + \frac{w_{\epsilon_k}^2}{2} \right) + (-1)^\nu (1 - q_k) f \left((-1)^\nu w_{0i} + \frac{w_{\epsilon_k}^2}{2} \right) \right\} \\ &- \frac{2 \sum_{m \in S(i,y)} w_{iy} w_{my} q_m + (2 \log y - 2w_{0y} - w_{\epsilon_y}^2) w_{iy}}{2w_{\epsilon_y}^2} + \frac{(2 \log y - 2w_{0y} - w_{\epsilon_y}^2) w_{iy}}{2w_{\epsilon_y}^2}, \tag{6} \end{aligned}$$

where $f(x) = \log(1 + \exp(x))$. The logistic function ensures that the posterior is between 0 and 1. The update for unobserved node i depends only on the states of its parents, children, and spouses.

In the maximization step, we conduct stochastic gradient descent updates to estimate the optimal weights using a mini-batch of data randomly sampled from different locations. Stochastic variational inference is designed here to accelerate the computational process over a large-scale high-resolution map. The edge weights at the iteration $t + 1$ are therefore updated as:

$$w^{(t+1)} = w^{(t)} + \rho \mathcal{A} \nabla L^{(t)}(w) \tag{7}$$

where ρ controls the learning rate, and \mathcal{A} is a preconditioner. Here, \mathcal{A} is set up as the identity matrix to accelerate convergence to high-likelihood models. In each iteration, we first randomly sample a mini-batch of locations from the given map. Then, the expectation step and maximization step are implemented to update the posterior estimations and the global weight parameters. As the model converges, the optimal posteriors of landslide, liquefaction, and building damage at each location are estimated.

Local pruning for accelerating the computation. Computational efficiency is another key element determining the applicability of multi-hazard estimation in real-world scenarios. We observed that real-world causal graph tends to be more sparse in many locations: only a subset of nodes are active for each location. For example, if there is no building footprint, the building damage node will be inactive. We also observed that when the slope is steep, liquefaction has a low probability to co-occur with landslide⁶¹. Therefore, when the prior landslide probability is significantly higher than the prior liquefaction estimation (e.g., the difference between prior landslide and liquefaction is higher than the median of the overall difference distribution between prior landslide and liquefaction estimations), we can treat the liquefaction node as inactive. For inactive nodes, the posterior probability of their ancestor nodes is typically very small. Moreover, these inactive nodes have little influence on parameter updates on the graph as the gradients of edge weights during updating process are proportional to the activation probabilities. According to the above intuition, we developed a local pruning strategy to remove these irrelevant subsets of the graph, while still retaining the nodes that contain information crucial to the subsequent parameter update. Specifically, we construct a location-specific local model as follows:

1. Select \mathcal{O}_l^+ , the set of active nodes for the location l .
2. Select \mathcal{H}_l , the parents of nodes in \mathcal{O}_l^+ . We do explicit variational inference updates only for this subset of nodes.
3. Link the selected node with causal dependencies.

Given each location, we conduct local pruning first to obtain a local model and then adapt stochastic variational inference over the local model. Our derived theoretically sound variational bound holds on any similar network structure and variable assumptions as our

network. Therefore, the stochastic variational inference algorithm can be easily adapted by setting the posterior q_i of inactive nodes i as zero. The combination of stochastic variational inference and local pruning dramatically reduces the computational cost and memory demands.

Prior geospatial models. The prior landslide and liquefaction models we used for comparison in the results are geospatial models from existing USGS ground failure products. As part of the Earthquake Hazards Program, the USGS developed a ground-failure earthquake product to augment the Prompt Assessment of Global Earthquakes for Response system⁵. Following major earthquakes worldwide, this product provides near real-time maps of earthquake-induced landslide and liquefaction probabilities. Both the liquefaction and landslide maps are derived from models that utilize ground-shaking intensity as an input, which allows the maps to be rapidly generated, but do not take into account information from remotely sensed images. Ground-motion inputs are taken from ShakeMaps, which, in turn, are derived from instrumental recordings, ground-motion models, and site conditions estimated from the shallow shear velocity. Additional inputs related to liquefaction vulnerability include mean annual precipitation, distance from the coast, distance from rivers, and water table depth.

Evaluation metrics. We mainly utilize the ROC (receiver operating characteristic) curve and Cross-entropy loss to evaluate and compare the performance of our system.

ROC curves. are a visualization method to represent the ability of a binary classifier system as its discrimination threshold is varied⁶². Both our system and prior models mainly output the probability estimation of landslide, liquefaction, or building damage. In real-world applications, a threshold is often applied to the probability to divide the estimates into “there exists damage” and “no damage”. As the threshold varies from event to event, it would be unfair to directly set up a fixed threshold to evaluate the system performance. In contrast, the ROC curve visualizes how the system performance changes under a varying threshold ranging between 0 and 1, which is fairer. The ROC curve is created by plotting the true positive rate (TPR) against the false positive rate (FPR) at various threshold settings. Given a binary classification problem, label 1 represents that a damage has occurred in the location. We define the true positive (TP) as that both predicted label and true label are 1, i.e., both prediction model and ground truth observations that the damage has occurred. Similarly, true negative (TN) refers to that both prediction model and ground truth observations indicate that no damage has occurred. In contrast, if the predicted label is 1 but true label is 0, i.e., the prediction model estimates the damage has occurred, but the damage is not observed in the ground truth data, we define it as a false positive (FP) prediction and otherwise false negative (FN). The true positive rate (TPR) is also known as sensitivity, recall, or probability of detection, and can be calculated as $\frac{\#TP}{\#TP + \#FN}$. The false positive rate (FPR) is also known as the probability of false alarm and can be calculated as $\frac{\#FP}{\#FP + \#TN}$.

Cross-entropy loss. is one of the most widely used loss functions⁶³ to measure the performance of a classification model whose output is a probability value between 0 and 1. It is designed as the mutual information between two probability distributions, mostly one discrete distribution and one continuous distribution, over the same underlying set of events. Mathematically, for binary classification models, a cross-entropy loss (CEL) is defined as below:

$$CEL = -\frac{1}{N} \sum_{n=1}^N [z_n \log P_n + (1 - z_n) \log(1 - P_n)], \quad (8)$$

where $z_n \in \{0, 1\}$ represents the true label, $z_n=1$ represents that damage occurred, and P_n represents the predicted probability of the

damage. As the predicted probability diverges from the true label, the cross-entropy loss increases. For example, predicting one pixel's landslide probability of 0.001 while actual landslides are observed in the corresponding location will yield a high loss value, meaning that the prediction model has high uncertainties about the landslides occurrence variable. Ideally, a perfect model would have a cross-entropy loss of 0. Therefore, smaller cross-entropy loss often means a more accurate model that can better predict the true label distribution.

Data availability

For the four earthquake events presented, the building damage, posterior landslide, and posterior liquefaction models generated in this study have been deposited in the Zenodo database under accession code <https://doi.org/10.5281/zenodo.7319726>⁶⁴. For the input datasets, the original sources of prior models and building footprints are publicly available at <https://earthquake.usgs.gov/>^{28–31} and <http://download.geofabrik.de/>²³, respectively. The DPM for the 2020 Puerto Rico earthquake is also publicly available at https://aria-share.jpl.nasa.gov/20200106-Puerto_Rico_EQ/DPM/⁴². The updated DPMs for other earthquake events are under restricted access because these are not made publicly available at ARIA website, access can be obtained by contacting S.X. Post-processed DPMs, prior models, and building footprints to be used as inputs to our code have also been deposited in the Zenodo database under accession code <https://doi.org/10.5281/zenodo.7319726>⁶⁴. Some of the original sources of ground truth datasets are publicly available through the following links: (1) 2016 Central Italy earthquake: landslide (<https://www.sciencebase.gov/catalog/item/5b6b2201e4b006a11f7795c3>)^{41,48}; (2) 2018 Hokkaido, Japan earthquake: landslide (<https://www.gsi.go.jp/common/000204728.zip>)⁴⁷; (3) 2019 Ridgecrest, California earthquake: landslide, liquefaction, and building damage (<https://www.designsafe-ci.org/data/browser/public/designsafe.storage.published/PRJ-2440>)^{17,57,59}; and (4) 2020 Puerto Rico earthquake: landslide, liquefaction, and building damage (<https://www.sciencebase.gov/catalog/item/5eb5b9dc82ce25b5135ae83a>)⁴⁰. The building damage ground truth dataset⁵⁴ after the 2016 Central Italy earthquake and the some landslide and liquefaction ground truth data⁵⁸ are under restricted access because these are not made publicly available by the original authors, access can be obtained by contacting S.X.

Code availability

The source codes for the implementation of the joint natural hazards estimation workflow including the scripts for evaluation metrics are publicly available online under accession code <https://doi.org/10.5281/zenodo.7319726>⁶⁴.

References

- Bird, J. F. & Bommer, J. J. Earthquake losses due to ground failure. *Eng. Geol.* **75**, 147–179 (2004).
- Wang, F. et al. Coseismic landslides triggered by the 2018 Hokkaido, Japan (Mw 6.6), earthquake: spatial distribution, controlling factors, and possible failure mechanism. *Landslides* **16**, 1551–1566 (2019).
- Jibson, R. W. Types and areal distribution of ground failure associated with the 2019 Ridgecrest, California, earthquake sequence. *Bull. Seismological Soc. Am.* **110**, 1567–1578 (2020).
- De Groeve, T., Vernaccini, L., Annunziato, A., Van de Walle, B. & Turoff, M. Global disaster alert and coordination system. In *Proc. 3rd International ISCRAM Conference* (eds. Van de Walle, B. & Turoff, M.) 1–10 (Newark, 2006).
- Earle, P. S. et al. Prompt assessment of global earthquakes for response (PAGER): A system for rapidly determining the impact of earthquakes worldwide. *US Geol. Surv. Open-File Rep.* **2009-1131**, 15 (2009).

6. Spinuso, A. et al. Network of Research Infrastructures for European Seismology (NERIES)—Web Portal Developments for Interactive Access to Earthquake Data on a European Scale. *Geoinformatics* **11**–13 (2008).
7. 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).
8. Ghorbanzadeh, O. et al. Evaluation of different machine learning methods and deep-learning convolutional neural networks for landslide detection. *Remote Sens.* **11**, 196 (2019).
9. So, E. & Spence, R. Estimating shaking-induced casualties and building damage for global earthquake events: a proposed modelling approach. *Bull. Earthq. Eng.* **11**, 347–363 (2013).
10. Jaiswal, K., Wald, D. & Porter, K. A global building inventory for earthquake loss estimation and risk management. *Earthq. Spectra* **26**, 731–748 (2010).
11. Nielson, B. G. & DesRoches, R. Analytical seismic fragility curves for typical bridges in the central and southeastern united states. *Earthq. spectra* **23**, 615–633 (2007).
12. Lee, S. Application of logistic regression model and its validation for landslide susceptibility mapping using GIS and remote sensing data. *Int. J. remote Sens.* **26**, 1477–1491 (2005).
13. Zhao, C. & Lu, Z. Remote sensing of landslides—A review. *Remote Sens.* **10**, 279 (2018).
14. Yun, S.-H., Fielding, E. J., Webb, F. H. & Simons, M. Damage proxy map from interferometric synthetic aperture radar coherence. US Patent 9,207,318 (2015).
15. Burrows, K., Milledge, D., Walters, R. J. & Bellugi, D. Improved rapid landslide detection from integration of empirical models and satellite radar. *Nat. Hazards Earth Syst. Sci. Discuss.* Preprint at <https://doi.org/10.5194/nhess-2021-148> (2021).
16. Li, Q., Wang, W., Wang, J., Zhang, J. & Geng, D. Exploring the relationship between insar coseismic deformation and earthquake-damaged buildings. *Remote Sens. Environ.* **262**, 112508 (2021).
17. Zimmaro, P. et al. Liquefaction and related ground failure from July 2019 ridgecrest earthquake sequence. *Bull. Seismological Soc. Am.* **110**, 1549–1566 (2020).
18. Loos, S. et al. G-DIF: a geospatial data integration framework to rapidly estimate post-earthquake damage. *Earthq. Spectra* **36**, 1695–1718 (2020).
19. Sun, X. et al. Using insar and polsar to assess ground displacement and building damage after a seismic event: Case study of the 2021 baicheng earthquake. *Remote Sens.* **14**, 3009 (2022).
20. Zhang, N. L. & Poole, D. Exploiting causal independence in Bayesian network inference. *J. Artif. Intell. Res.* **5**, 301–328 (1996).
21. Nielsen, T. D. & Jensen, F. V. *Bayesian Networks and Decision Graphs* (Springer Science & Business Media, 2009).
22. Blei, D. M., Kucukelbir, A. & McAuliffe, J. D. Variational inference: a review for statisticians. *J. Am. Stat. Assoc.* **112**, 859–877 (2017).
23. Geofabrik. OpenStreetMap Data Extracts. <http://download.geofabrik.de/osm/> (2020).
24. Microsoft. US Building Footprints. <https://github.com/microsoft/USBuildingFootprints> (2020).
25. Budimir, M., Atkinson, P. & Lewis, H. Earthquake-and-landslide events are associated with more fatalities than earthquakes alone. *Nat. Hazards* **72**, 895–914 (2014).
26. Popescu, M. E. Landslide causal factors and landslide remedial options. In *3rd International Conference on Landslides, Slope Stability and Safety of Infra-Structures*, 61–81 (Citeseer, 2002).
27. Zhang, J. et al. How size and trigger matter: analyzing rainfall-and earthquake-triggered landslide inventories and their causal relation in the Koshi River basin, central Himalaya. *Nat. Hazards Earth Syst. Sci.* **19**, 1789–1805 (2019).
28. U.S. Geological Survey. M 6.4 - 8km S of Indios, Puerto Rico. earthquake.usgs.gov/earthquakes/eventpage/us70006vll (2021).
29. U.S. Geological Survey. M 6.6 - 27km ENE of Tomakomai, Japan. earthquake.usgs.gov/earthquakes/eventpage/us2000h8ty/ (2021).
30. U.S. Geological Survey. M 7.1 - 2019 Ridgecrest earthquake sequence. <https://earthquake.usgs.gov/earthquakes/eventpage/ci38457511> (2021).
31. U.S. Geological Survey. M 6.2 - 10km SE of Norcia, Italy. earthquake.usgs.gov/earthquakes/eventpage/us10006g7d (2021).
32. Zhu, J., Baise, L. G. & Thompson, E. M. An updated geospatial liquefaction model for global application. *Bull. Seismological Soc. Am.* **107**, 1365–1385 (2017).
33. Bradley, A. P. The use of the area under the ROC curve in the evaluation of machine learning algorithms. *Pattern Recognit.* **30**, 1145–1159 (1997).
34. Zhang, Z. & Sabuncu, M. R. Generalized cross entropy loss for training deep neural networks with noisy labels. In *32nd Conference on Neural Information Processing Systems (NeurIPS)* (2018).
35. Fire and Disaster Management Agency (FDMA). Damage caused by the 2018 Hokkaido Eastern Iburi earthquake and response status of firefighting agencies. *Tech. Rep.* <https://www.fdma.go.jp/disaster/info/2018/> (2021).
36. Allstadt, K. E. et al. Ground failure triggered by the 7 January 2020 m 6.4 puerto rico earthquake. *Seismological Soc. Am.* **93**, 594–608 (2022).
37. Osanai, N. et al. Characteristics of landslides caused by the 2018 hokkaido eastern iburi earthquake. *Landslides* **16**, 1517–1528 (2019).
38. Aoki, T., Katsura, S., Tanaka, Y., Koi, T. & Yamada, T. Hydraulic properties of and pressure-head dynamics in thick pyroclastic-fall deposits in Atsuma, Northern Japan: implications for the role of water in shallow landslides induced by the 2018 Hokkaido Eastern Iburi Earthquake. *Landslides* **19**, 1813–1824 (2022).
39. Federal Emergency Management Agency. Puerto Rico M 6.4 earthquake preliminary damage assessments dashboard. <https://fema.maps.arcgis.com/apps/opsdashboard/index.html> (2020).
40. Allstadt, K. et al. Field observations of ground failure triggered by the 2020 Puerto Rico earthquake sequence. U.S. Geological Survey data release. <https://doi.org/10.5066/P96QNFMB> (2021).
41. Schmitt, R. G. et al. An open repository of earthquake-triggered ground-failure inventories. *Tech. Rep.*, US Geological Survey (2017).
42. 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).
43. Miranda, E. et al. StEER - 07 Jan. 2020 Puerto Rico mw6.4 Earthquake: preliminary virtual reconnaissance report (PVRr). Designsafe-CI (2020).
44. Miranda, E. et al. StEER - Puerto Rico Earthquake Sequence December 2019 to January 2020: early access reconnaissance report (EARR). Designsafe-CI (2020).
45. Kayen, R. et al. Seismological, geological, and geotechnical engineering aspects of the 2018 MW 6.6 Hokkaido Eastern Iburi earthquake. *Tech. Rep.*, Geotechnical Extreme Events Reconnaissance Association (GEER) (2019).
46. Advanced Rapid Imaging and Analysis (ARIA) - Center for Natural Hazards. ARIA Data Share. aria-share.jpl.nasa.gov/20180905_Hokkaido_EQ/DPM/ (2020).
47. Geospatial Institute of Japan. Slope collapse and sedimentation distribution map associated with the 2018 Hokkaido Eastern Iburi Earthquake (around Atsuma Town, Hokkaido), Technical Data D1-No.991. www.gsi.go.jp/BOUSAI/H30-hokkaidoiburi-east-earthquake-index.html (2021).
48. Martino, S., Prestininzi, A. & Romeo, R. Earthquake-induced ground failures in Italy from a reviewed database. *Nat. Hazards Earth Syst. Sci.* **14**, 799–814 (2014).
49. Darwiche, A. Bayesian networks. *Commun. ACM* **53**, 80–90 (2010).

50. Fahmi, A. et al. Causal bayesian networks for medical diagnosis: A case study in rheumatoid arthritis. In *2020 IEEE International Conference on Healthcare Informatics (ICHI)*, 1–7 (IEEE, 2020).
51. Stringham, C. et al. The capella X-band SAR constellation for rapid imaging. In *IGARSS 2019-2019 IEEE International Geoscience and Remote Sensing Symposium*, 9248–9251 (IEEE, 2019).
52. Ekhtari, N. & Glennie, C. High-resolution mapping of near-field deformation with airborne Earth observation data, a comparison study. *IEEE Trans. Geosci. Remote Sens.* **56**, 1598–1614 (2017).
53. Synthetic-aperture radar is making the earth's surface watchable 24/7. *The EconomistTechnology Quarterly* (2022).
54. Fiorentino, G. et al. Damage patterns in the town of Amatrice after August 24th 2016 Central Italy earthquakes. *Bull. Earthq. Eng.* **16**, 1399–1423 (2018).
55. Advanced Rapid Imaging and Analysis (ARIA) - Center for Natural Hazards. ARIA Data Share. https://aria-share.jpl.nasa.gov/20160824-Italy_EQ/DPM/ (2020).
56. DuRoss, C. B. et al. Surface displacement distributions for the July 2019 Ridgecrest, California, earthquake ruptures. *Bull. Seismological Soc. Am.* **110**, 1400–1418 (2020).
57. Brandenburg, S. J. et al. Preliminary report on engineering and geological effects of the July 2019 Ridgecrest earthquake sequence. Tech. Rep., Geotechnical Extreme Event Reconnaissance Association (2019).
58. Jibson, R. W. Types and areal distribution of ground failure associated with the 2019 Ridgecrest, California, earthquake sequence. *Bull. Seismological Soc. Am.* **110**, 1567–1578 (2020).
59. Brandenburg, S. et al. GEER field reconnaissance, Ridgecrest, CA earthquake sequence, July 4 and 5, 2019. <https://doi.org/10.17603/DS2-VPMV-5B34> (2021).
60. Advanced Rapid Imaging and Analysis (ARIA) - Center for Natural Hazards. ARIA Data Share. aria-share.jpl.nasa.gov/20190704-0705-Searles_Valley_CA_EQs/DPM/ (2020).
61. Allstadt, K. E. et al. Integrating landslide and liquefaction hazard and loss estimates with existing USGS real-time earthquake information products (2017).
62. Gonçalves, L., Subtil, A., Oliveira, M. R. & de Zea Bermudez, P. Roc curve estimation: an overview. *REVSTAT-Stat. J.* **12**, 1–20 (2014).
63. Murphy, K. P. *Machine Learning: A Probabilistic Perspective* (MIT press, 2012).
64. Susu Xu, J. D. Seismic multi-hazard and impact estimation via causal inference from satellite imagery. <https://doi.org/10.5281/zenodo.7319726> (2022).

Acknowledgements

Any use of trade, firm, or product names is for descriptive purposes only and does not imply endorsement by the U.S. Government. S.X. was supported by USGS Earthquake Hazards Program G22AP00032 and Research Foundations at Stony Brook University. J.D. was supported by the Stanford Knight-Hennessy Scholars Program. J.D. was affiliated with the Department of Civil and Environmental Engineering at Stanford University at the time of work and is currently affiliated with the UKRI CDT in the application of artificial intelligence to the study of Environmental

Risks (AI4ER) at the University of Cambridge. H.N. was supported by USGS IPA. The authors would like to thank Dr. Sang-Ho Yun, Dr. Eric Fielding, and NASA ARIA team for providing their Damage Proxy Maps products. The authors would like to thank Dr. Kate Allstadt, Dr. Eric Thompson, Dr. Kishor Jaiswal, Dr. Sabine Loos, and Davis Engler from USGS for providing their prior geospatial models and valuable suggestions.

Author contributions

S.X., D.J.W., and H.N. developed the causal graph model formulation and variational Bayesian inference system. S.X. and J.D. performed the experiments and worked with D.J.W. to process the experimental data. S.X. and J.D. implemented the seismic multi-hazard and impact estimation experiments. All authors wrote, read, and approved the final manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41467-022-35418-8>.

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

Peer review information *Nature Communications* thanks Gemma Cremen and the other, anonymous, reviewer(s) for their contribution to the peer review of this work.

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 4.0 International License, which permits use, sharing, adaptation, 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 license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license 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 license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2022