Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Information-Availability-Driven Dynamic Probabilistic Updating of Post- Earthquake Building Repair Time

Haihui Zhao
China Earthquake Administration
Jianming Wang
China Earthquake Administration
Zifa Wang
zifa@iem.ac.cn

Key Laboratory of Earthquake Disaster Mitigation, Ministry of Emergency Management, Harbin, 150080, China
Dengke Zhao
China Earthquake Administration
Siran Yang
China Earthquake Administration

Research Article

Keywords: Post-earthquake building repair time, Information availability, Dynamic probabilistic updating, Probabilistic prediction

Posted Date: September 28th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10770435/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Information-Availability-Driven Dynamic Probabilistic Updating of Post-Earthquake Building Repair Time

*Haihui Zhao[3], Jianming Wang[3], Zifa Wang[1, 2, *], Dengke Zhao [3], Siran Yang[3]*

[1] Key Laboratory of Earthquake Engineering and Engineering Vibration, Institute of Engineering Mechanics, China Earthquake Administration, Harbin, 150080, China

[2] Key Laboratory of Earthquake Disaster Mitigation, Ministry of Emergency Management, Harbin, 150080, China

[*] Corresponding author: zifa@iem.ac.cn

[3] Institute of Engineering Mechanics, China Earthquake Administration, Harbin, Heilongjiang, 150080, China

Abstract: Post-earthquake building repair time is an important indicator of urban recovery capacity and supports post-disaster recovery planning and resource allocation. However, information required for repair-time prediction becomes available progressively rather than simultaneously during real recovery processes. Existing approaches are generally developed under relatively complete and fixed information conditions, limiting their ability to accommodate evolving information availability. This study proposes an information-availability-driven dynamic probabilistic updating framework for post-earthquake building repair time. The framework organizes information into four sequential stages according to its practical availability before and after an earthquake. Stage 0 uses historical repair information to establish a baseline probabilistic state. Stage 1 incorporates building attributes. Stage 2 further introduces earthquake-specific ground-motion and site-condition information. Stage 3 additionally incorporates repair-cost and assessment-lag information obtained from post-earthquake assessments. At each updating stage, the preceding repair-completion hazard trajectory, cumulative available information, and newly acquired information are jointly used to update the probabilistic prediction. Repair completion is formulated as a discrete-time event process, with LightGBM

estimating interval-specific conditional repair-completion hazards over a fine-resolution time grid. The predicted hazards are transformed into survival and cumulative repair-completion probabilities, from which four broader engineering reporting ranges are derived rather than directly predicted as categorical outcomes. Using repair records from the 2010–2011 Canterbury earthquake sequence, the framework shows progressively improved point-summary and probabilistic prediction performance and reduced predictive uncertainty as information becomes available. The framework provides stage-wise probabilistic support for post-earthquake repair-progress assessment and recovery resource allocation under evolving information conditions.

Keywords: Post-earthquake building repair time; Information availability; Dynamic probabilistic updating; Probabilistic prediction

1 Introduction

Earthquakes can cause not only structural and nonstructural damage to buildings but also cascading consequences such as temporary population displacement, business interruption, infrastructure-service disruption, and regional economic losses (Miles and Chang 2006; Cook et al. 2022). Restoring building functionality is therefore an important component of post-earthquake recovery, and reducing recovery time is a central objective of seismic-resilience enhancement (Bruneau et al. 2003; Burton et al. 2016). Unlike conventional seismic damage assessment, which primarily characterizes the physical condition of a building immediately after an earthquake, building repair time describes the temporal progression from damage occurrence and assessment through repair implementation to eventual repair completion. Reliable prediction of this time can support re-occupancy planning, business resumption, repair-resource allocation, construction prioritization, and long-term recovery management.

Post-earthquake building repair time is affected by multiple interacting factors. Building characteristics such as floor area, number of storeys, construction age, structural and material types, and occupancy can influence both seismic performance

and repair complexity. Ground-motion intensity, source-to-site distance, local site conditions, liquefaction, and ground deformation characterize the seismic demand experienced by individual buildings, while damage severity, economic loss, financing conditions, approval processes, contractor and labor availability, and repair-management arrangements can further affect the initiation and completion of repair activities (Comerio 2006; Di Ludovico et al. 2017a, b; FEMA 2018; Basaglia et al. 2020; Xiong et al. 2020; Sattar et al. 2021; Cook et al. 2022). Previous post-earthquake studies have demonstrated the predictive value of building characteristics, geographic information, source-distance measures, and ground-motion parameters for building-level seismic assessment. Consequently, repair time is not governed by a single physical or socioeconomic indicator but by the combined effects of building characteristics, seismic demand, damage consequences, and recovery-management conditions.

A practical difficulty is that these sources of information do not become available simultaneously. Basic building attributes and geographic information may already exist before an earthquake or can be rapidly retrieved from existing databases. In contrast, earthquake-specific ground-motion products and spatial intensity measures become available as monitoring and spatial analyses are completed, while detailed damage, economic-loss, assessment, and repair-related information generally requires subsequent inspection and verification. Accordingly, the information available for repair-time prediction evolves throughout the earthquake and recovery process. Predictions made shortly after an event must therefore rely on more limited information than predictions made after detailed assessment data become available (Giordano et al. 2022).

Machine-learning methods, including random forests, gradient-boosted trees, support vector machines, and deep-learning approaches, have increasingly been applied to seismic damage identification, vulnerability assessment, and loss prediction because of their ability to represent nonlinear relationships and interactions among heterogeneous predictors. Studies by Di Michele et al. (2023), Ghimire and Guéguen

(2024), Doğan et al. (2025), and Bektaş and Kegyes-Brassai (2025) demonstrate the potential of data-driven methods to integrate building, ground-motion, site, and post-earthquake survey information. However, most such approaches are developed using a predefined and relatively fixed feature set and focus primarily on prediction performance under a given information condition.

Some recent studies have begun to examine prediction revision as new observations become available (Gehl et al. 2022). Negulescu et al. (2020) updated exposure and vulnerability states after preceding damaging events, while Tomar et al. (2022) used Bayesian inference and discrete-event simulation to update restoration forecasts for water-distribution systems as damage and repair information accumulated. Haindl et al. (2024) further showed that newly acquired post-earthquake inspection information can constrain predictions and reduce uncertainty for damaged buildings. These studies demonstrate the value of incorporating newly available evidence into post-earthquake prediction. Nevertheless, an explicit mechanism for progressively updating building repair-time probability distributions according to changing information availability remains limited.

Existing building repair-time prediction approaches generally assume that the predictors required by the model have already been obtained. This assumption is appropriate when relatively complete post-event datasets are available for retrospective model development, but it does not directly represent the transition from limited to increasingly complete information encountered during actual recovery. If part of the required information is not yet available, a model developed using the complete feature set cannot be applied under the same conditions; conversely, when additional information subsequently becomes available, previous predictions generally lack a systematic mechanism for incorporating the new evidence. Therefore, a repair-time prediction framework is needed that can establish a probabilistic assessment under limited information and progressively revise that assessment as the information state changes.

To address this gap, this study proposes an information-availability-driven dynamic probabilistic updating framework for post-earthquake building repair time**.** The available information is organized into four sequential and nested stages according to its practical temporal availability before and after an earthquake.

Historical repair information is first used to establish a baseline probabilistic state. Building attributes are subsequently incorporated, followed by earthquake-specific ground-motion and site-condition information and, finally, repair-cost and assessment-lag information obtained through later post-earthquake assessments. These stages represent changes in information availability rather than fixed post-earthquake calendar times, and each subsequent stage retains the information accumulated previously while introducing newly available information.

Different from directly assigning buildings to predefined repair-time categories, the proposed framework formulates repair completion as a discrete-time event process. The repair-time axis is divided into fine temporal intervals, and LightGBM is used as a probabilistic base learner to estimate the conditional probability that a building completes repair during a given interval, conditional on remaining unrepaired at the beginning of that interval. The resulting interval-specific repair-completion hazards are transformed into survival probabilities and cumulative repair-completion probabilities. Four broader ranges—within 1 year, $1-2$years, $2-3$ years, and more than 3 years—are subsequently derived from the cumulative distribution for engineering interpretation rather than used as direct categorical model outputs.

The key updating mechanism is the sequential propagation of the predicted repair-completion state. At each information stage, the hazard trajectory obtained from the preceding stage is retained and combined with the cumulative information already available and the information newly acquired at the current stage. The updated hazard trajectory is then propagated to the subsequent stage. Thus, the framework does not consist of four independent prediction models and does not

involve continuous online retraining. Instead, dynamic probabilistic updating refers to stage-wise probabilistic revision triggered by changes in information availability**.**

The framework is evaluated using building repair records from the 2010–2011 Canterbury earthquake sequence in New Zealand. The study focuses on three related objectives: establishing a probabilistic repair-completion prediction under limited information, quantifying how prediction accuracy and uncertainty evolve as additional information becomes available, and examining how different types of newly available information modify the predicted repair-completion process. Both building-level predictions and building-stock cumulative repair-completion trajectories are evaluated. The methodological contribution therefore lies not simply in achieving a final high-accuracy prediction under the most complete information condition, but in establishing a systematic retain–update–propagate mechanism for probabilistic repair-time prediction under evolving information conditions**.**

2 Data

2.1 Study background and data sources

This study focuses on the 4 September 2010 Darfield earthquake and the 22 February 2011 Christchurch earthquake, both within the Canterbury earthquake sequence. As shown in Fig. 1, stronger ground motions during the 2010 Darfield earthquake were mainly distributed west of Christchurch, whereas the high-PGA region of the 2011 Christchurch earthquake was concentrated within and around the Christchurch urban area. The red stars indicate earthquake epicenters and the circles denote strong-motion stations.

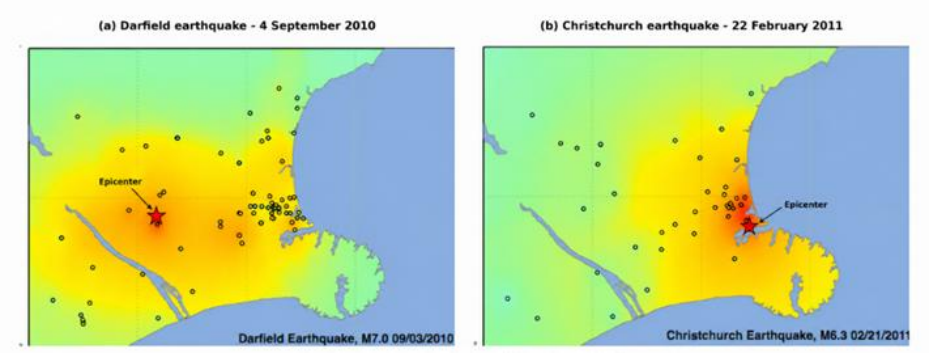


Fig. 1 PGA distributions and strong-motion station locations for the 2010

Darfield and 2011 Christchurch earthquakes

The primary building dataset was derived from a New Zealand post-earthquake insurance database containing building attributes, geographic information, loss-assessment results, and managed-repair records. Ground-motion and site-condition information was obtained from the GeoNet strong-motion network and the New Zealand Strong Motion Database (NZSMDB). Building coordinates and earthquake-event information were integrated with these sources to derive building-level ground-motion intensity measures, source-to-site distances, and site-condition variables.

The two earthquake datasets serve different roles in the proposed framework. The 2010 Darfield repair records are used to establish the historical baseline repair-completion model, whereas the 2011 Christchurch earthquake is treated as the target event for stage-wise probabilistic updating. Detailed sample sizes and data partitioning are described in Sect. 4.1.

2.2 Feature engineering

The raw dataset contains variables that are irrelevant to the research objective, redundant, severely incomplete, identifier-related, or potentially associated with information leakage. Variables were therefore screened according to their physical relevance, practical time of availability, and relationship with the prediction target. Identifier fields, duplicated variables, irrelevant variables, and variables that could reveal the repair outcome before prediction were excluded. Additional variables were subsequently constructed to characterize building geometry, seismic demand, source-to-site effects, local site conditions, and post-earthquake assessment and loss characteristics.

For building i, repair time T_i is defined as the calendar duration between the earthquake-event date and the recorded managed-repair completion date and is expressed directly in years:

$$T_i = \frac{D_{i,\text{complete}} - D_{i,\text{event}}}{365.25} \quad (1)$$

where $D_{i,\text{complete}}$ denotes the managed-repair completion date and $D_{i,\text{event}}$ denotes the corresponding earthquake-event date.

This definition represents the total elapsed time from earthquake occurrence to recorded completion of managed building repair. It may therefore include damage assessment, insurance-claim processing, engineering design and approval, waiting periods, on-site repair work, and completion registration and should not be interpreted solely as construction duration (Molina Hutt et al. 2022).

To characterize building geometry, the building area ratio is defined as

$$R_{A,i} = \frac{A_{i,\text{plan}}}{A_{i,\text{floor}}} \tag{2}$$

where $R_{A,i}$ denotes the building footprint area and $A_{i,\text{floor}}$ denotes the total floor area. This study-specific descriptor approximately represents the vertical distribution of floor area and provides complementary information to the original building attributes.

For records without directly available building height, an approximate height was obtained from the number of storeys:

$$H_i = 3.0N_i \tag{3}$$

where N_i is the number of storeys and H_i is the estimated building height in metres. The nominal storey height of 3.0 m is used only to obtain an approximate geometric quantity required for estimating the building fundamental period.

The corresponding fundamental period is estimated as

$$T_{1,i} = 0.0625H_i^{0.75} \tag{4}$$

where $T_{1,i}$ denotes the estimated fundamental period of building in seconds.

The estimated height and fundamental period are primarily used to determine the period-specific seismic demand of each building. Because these quantities are derived from the original geometric information, no additional composite vertical-scale index is introduced, thereby avoiding unnecessary deterministic redundancy among engineered predictors.

To characterize building-specific seismic demand and local site conditions, earthquake-specific variables were constructed or matched at individual building

222 locations. These include peak ground acceleration (PGA), liquefaction-related
223 information, median vertical ground displacement, $V_{S30,i}$,$Z_{1.0,i}$, $R_{\mathrm{jb},i}$, $R_{\mathrm{rup},i}$, $S_{a,i}$ .The
224 variable $S_{a,i}$ represents the 5%-damped spectral acceleration corresponding to the
225 estimated fundamental period $T_{1,i}$ of each building:

226 $$S_{a,i} = S_a\left(T_{1,i}\right),$$

227 thereby providing a period-dependent representation of seismic demand. New
228 Zealand-specific ground-motion studies provide the empirical basis for estimating
229 spectral demand, while strong-motion observations and spatial information are
230 incorporated to obtain building-level ground-motion estimates (Bradley 2013; Bradley
231 and Hughes 2013; Bradley 2014).

232 $V_{S30,i}$ characterizes near-surface site stiffness, whereas $Z_{1.0,i}$ represents the depth
233 to the deeper shear-wave-velocity horizon and therefore provides information on
234 sedimentary structure. $R_{\mathrm{jb},i}$ and $R_{\mathrm{rup},i}$ describe source-to-site distance using different
235 rupture-distance definitions. Together with PGA, $S_{a,i}$, liquefaction observations, and
236 ground-deformation information, these variables provide building-specific
237 representations of seismic demand, source-to-site effects, and local site conditions.

238 Post-earthquake building repair is also influenced by the timing and outcome of
239 assessment. Previous studies of the Canterbury earthquake sequence have used
240 milestone dates in insurance and repair records to quantify claim, assessment, repair,
241 and settlement durations (Khakurel et al. 2021; Li et al. 2025).

242 For building i, assessment lag is defined as

243 $$L_{A,i} = D_{i,\mathrm{assessment}} - D_{i,\mathrm{event}},$$

244 where $D_{i,\mathrm{assessment}}$ denotes the assessment date. Assessment lag represents the elapsed
245 time from earthquake occurrence until assessment information becomes available. It
246 is therefore introduced only at the information stage at which the corresponding
247 assessment has been completed.

248 Economic loss severity is represented by the building loss ratio,

$$LR_i = \frac{C_{i,\text{repair}}}{V_{i,\text{building}}},$$

where $C_{i,\text{repair}}$ denotes the assessed building repair cost and $V_{i,\text{building}}$ denotes the rated building value. Larger values of LR_i indicate greater repair cost relative to building value and therefore higher normalized loss severity (De Martino et al. 2017). Repair-cost and assessment-related variables are used only after the corresponding post-earthquake assessment information becomes available.

2.3 Information-availability states

For building i, let $\mathcal{J}_i^{(s)}$ denote the information available at stage s, where $s \in \{0,1,2,3\}$,. The four information sets are nested as

$$\mathcal{J}_i^{(0)} \subset \mathcal{J}_i^{(1)} \subset \mathcal{J}_i^{(2)} \subset \mathcal{J}_i^{(3)} \quad (5)$$

Here, $\mathcal{J}_i^{(0)}$ represents the historical baseline and contains historical post-earthquake repair information. $\mathcal{J}_i^{(1)}$ additionally includes building attributes that are available before the earthquake or can be rapidly retrieved after the event. $\mathcal{J}_i^{(2)}$ further incorporates earthquake-specific ground-motion and site-condition information, including intensity measures, source-to-site distances, site parameters, and liquefaction-related information. Finally, $\mathcal{J}_i^{(3)}$ includes repair-cost and assessment-lag information obtained from subsequent post-earthquake assessments.

The four information sets therefore form a progressively expanding structure in which each stage retains all information available previously and adds the information newly available at the current stage. These stages represent information-availability conditions rather than fixed post-earthquake calendar times. The complete list of variables and their allocation to the four information states is provided in Appendix Table A1.

3 Methods

3.1 Information-availability-driven dynamic probabilistic updating framework

Post-earthquake building repair-time prediction is performed under evolving information conditions. Based on the four information states defined in Sect. 2.3, this study develops a stage-wise probabilistic updating framework in which the current probabilistic state is retained and revised when additional information becomes available.

For each updating stage s = 1,2,3, the newly available information is defined as

$$\Delta \mathcal{J}_i^{(s)} = \mathcal{J}_i^{(s)} \backslash \mathcal{J}_i^{(s-1)} \tag{6}$$

and the cumulative information state is updated as

$$\mathcal{J}_i^{(s)} = \mathcal{J}_i^{(s-1)} \cup \Delta \mathcal{J}_i^{(s)} \tag{7}$$

This formulation distinguishes previously accumulated information from newly introduced information and provides the basis for sequential updating.

At each information stage, the repair-completion process is represented by a stage-specific discrete-time hazard trajectory:

$$h_i^{(s)} = \left[h_{i1}^{(s)}, h_{i2}^{(s)}, \dots, h_{iK}^{(s)}\right]$$

where K denotes the number of discrete intervals on the repair-time axis, and $k = 1, \dots, K$ indexes an individual interval, and $h_{iK}^{(s)}$ denotes the conditional repair-completion hazard of building i during interval k under information stage s.The hazard trajectory at the current stage is obtained by combining the preceding-stage hazard trajectory, cumulative information, and newly available information:

$$h_i^{(s)} = M_s\left(h_i^{(s-1)}, \mathcal{J}_i^{(s-1)}, \Delta\mathcal{J}_i^{(s)}\right), \quad s = 1,2,3 \tag{8}$$

The updating procedure is summarized as follows:

(1) Retain the preceding-stage hazard trajectory and cumulative information.

(2) Incorporate the information newly available at the current stage.

(3) Apply the corresponding stage-specific updating model.

(4) Obtain the updated repair-completion hazard trajectory.

(5) Propagate the updated hazard trajectory and cumulative information to the subsequent stage.

Accordingly, the four information stages are not treated as four independent prediction models**.** Instead, they form a sequential probabilistic updating chain:

$$h_i^{(0)} \rightarrow h_i^{(1)} \rightarrow h_i^{(2)} \rightarrow h_i^{(3)}.$$

The updated hazard trajectory is subsequently transformed into survival probability and cumulative repair-completion probability, which are used to derive engineering reporting ranges, point summaries, uncertainty measures, and building-stock recovery trajectories.

The term dynamic probabilistic updating refers to stage-wise probabilistic revision triggered by changes in information availability, as illustrated in Fig. 2. It does not represent continuous-time parameter updating or online model retraining. Instead, the proposed framework establishes an initial probabilistic assessment under limited information and progressively refines the repair-completion prediction as additional information becomes available across successive information stages.

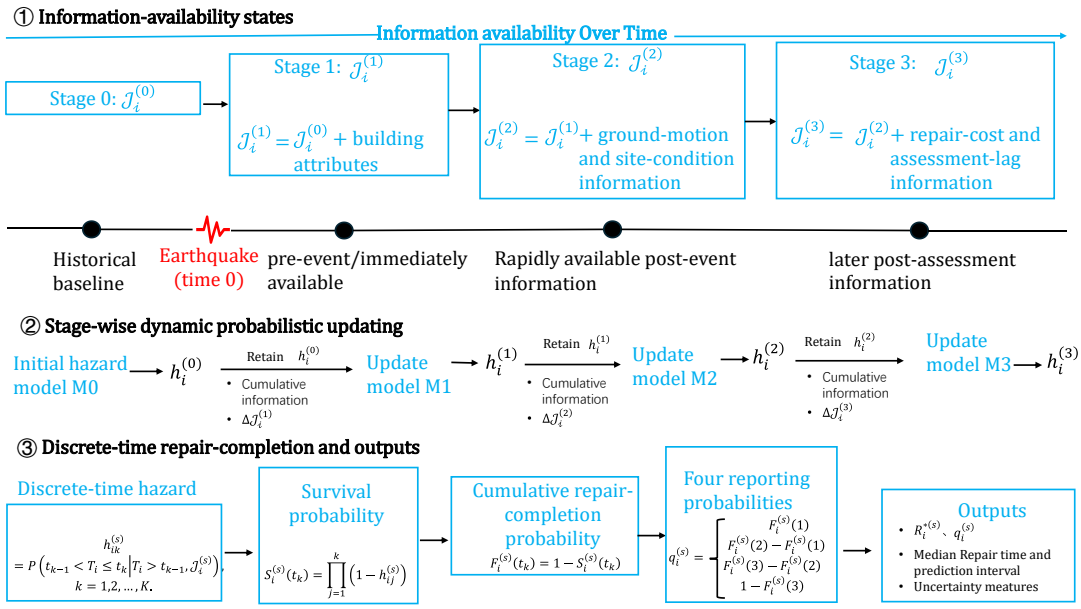


Fig. 2 Information-availability-driven dynamic probabilistic updating framework for post-earthquake building repair time

3.2 Discrete-time repair-completion formulation

Let T_i denote the observed repair completion time of building i, measured from the earthquake occurrence to the recorded completion of the managed repair process. To represent the temporal progression of repair completion probabilistically, the repair-time axis is discretized into a sequence of ordered evaluation times,

$$0 = t_0 < t_1 < t_2 < \cdots < t_K = \tau$$

where t_K denotes the kth evaluation time and τ is the primary prediction horizon. The consecutive discrete-time intervals are defined as

$$B_k = (t_{k-1}, t_k], \quad k = 1,2, \ldots, K \tag{9}$$

Here, B_k represent the kth ordered interval along the repair-time axis, whereas s denotes the information stage at which the prediction is updated. Thus, k describes temporal progression and s describes information availability.

For any condition, define the indicator $I(A)$ as

$$I(A) = \begin{cases} 1, if\ condition\ A\ is\ true, \\ 0, otherwise. \end{cases}$$

For building i, the interval-specific repair-completion indicator is

$$y_{ik} = I(t_{k-1} < T_i \leq t_k) \tag{10}$$

The corresponding at-risk indicator is

$$r_{ik} = I(T_i > t_{k-1}) \tag{11}$$

Thus, building i contributes to interval B_k only if it has not yet completed repair at the beginning of that interval. For an observed repair time, all intervals preceding the repair-completion interval have $y_{ik} = 0$, whereas the interval containing the observed repair completion has $y_{ik} = 1$. Intervals after repair completion are excluded because the building is no longer at risk of the repair-completion event.

For example, if $T_i = 1.37$ years and the interval width is 0.25 years, the repair completion event is assigned to interval $(1.25,1.50]$, resulting in the event sequence

$$0,\ 0,\ \ldots,\ 0,\ 1.$$

Subsequent intervals are excluded because the building is no longer at risk.

At information stage s, the discrete-time repair-completion hazard for building i

during interval B_k is defined as

$$h_{ik}^{(s)} = P\left(t_{k-1} < T_i \leq t_k \mid T_i > t_{k-1}, \mathcal{J}_i^{(s)} \right) \tag{12}$$

Equivalently

$$h_{ik}^{(s)} = P\left(y_{ik} = 1 \mid r_{ik} = 1, \mathcal{J}_i^{(s)} \right)$$

The quantity $h_{ik}^{(s)}$ therefore represents the conditional probability that building i, given that it remains unrepaired at time t_{k-1}, completes repair during the subsequent interval B_k under the information available at stage s (Singer and Willett 1993). The sequence of interval-specific hazards constitutes the stage-specific repair-completion hazard trajectory introduced in Sect. 3.1. The hazard $h_{ik}^{(s)}$ is conditional on the building remaining unrepaired at t_{k-1}.

Based on the interval-specific hazards, the probability that building remains unrepaired through time t_k is represented by the survival function

$$S_i^{(s)}(t_k) = P\left(T_i > t_k \mid \mathcal{J}_i^{(s)} \right)$$

which is obtained as

$$S_i^{(s)}(t_k) = \prod_{j=1}^{k} \left(1 - h_{ij}^{(s)}\right) \tag{13}$$

With

$$S_i^{(s)}(t_0) = 1$$

The corresponding cumulative repair-completion probability is defined as

$$F_i^{(s)}(t_k) = P\left(T_i \leq t_k \mid \mathcal{J}_i^{(s)} \right) = 1 - S_i^{(s)}(t_k)$$

Therefore,

$$F_i^{(s)}(t_k) = 1 - \prod_{j=1}^{k} \left(1 - h_{ij}^{(s)}\right) \tag{14}$$

The function $F_i^{(s)}(t_k)$ represents the predicted probability that building i has completed repair by evaluation time t_k under the information available at stage s. Because

$$0 \le h_{ik}^{(s)} \le 1$$

the resulting cumulative repair-completion probabilities satisfy

$$0 \le F_i^{(s)}(t_1) \le F_i^{(s)}(t_2) \le \cdots \le F_i^{(s)}(t_K) \le 1$$

Thus, the predicted cumulative repair-completion trajectory is monotonically non-decreasing by construction, consistent with the physical interpretation that once a building has completed repair, it remains repaired at all subsequent evaluation times.

The unconditional probability that building i completes repair specifically within interval B_k can subsequently be derived as

$$\pi_{ik}^{(s)} = P\left(t_{k-1} < T_i \le t_k \mid \mathcal{J}_i^{(s)} \right) = S_i^{(s)}(t_{k-1}) h_{ik}^{(s)} = F_i^{(s)}(t_k) - F_i^{(s)}(t_{k-1}) \quad (15)$$

Therefore, the probability masses associated with the individual repair-time intervals are derived from the estimated repair-completion process rather than directly predicted as mutually exclusive multiclass probabilities**.**

If a building has not completed repair by the end of the primary prediction horizon τ, its remaining probability is represented by

$$P\left(T_i > \tau \mid \mathcal{J}_i^{(s)} \right) = S_i^{(s)}(\tau)$$

Accordingly, repair completion beyond the primary prediction horizon is represented through the remaining survival probability rather than by assigning the building to an artificial terminal repair-time class.

The formulation above produces a complete stage-specific cumulative repair-completion distribution,

$$F_i^{(s)} = \left[F_i^{(s)}(t_1), F_i^{(s)}(t_2), \ldots, F_i^{(s)}(t_K) \right]$$

This cumulative distribution is the primary probabilistic representation of repair time in this study. The four broader reporting ranges are derived from $F_i^{(s)}$ only for engineering interpretation and are not used as training categories.

3.3 Stage-wise LightGBM implementation

The discrete-time formulation described in Sect. 3.2 transforms the repair-time prediction problem into estimation of interval-specific conditional repair-completion

probabilities. LightGBM is employed as the common base learner for estimating interval-specific repair-completion hazards in the initial model and the three subsequent stage-specific updating models (Ke et al. 2017). Importantly, LightGBM is not used as a multiclass classifier to assign buildings to predefined repair-time categories. Instead, it is used as a binary probabilistic learner on building–time-interval observations to estimate interval-specific conditional repair-completion hazards.

For model construction, each building is expanded into a sequence of observations corresponding to the discrete intervals during which it remains at risk of repair completion. Let

$$\mathcal{D}^{(s)} = \left\{\left(x_{ik}^{(s)}, y_{ik}\right) : r_{ik} = 1\right\}$$

denote the modeling dataset at information stage s, where y_{ik} is the repair-completion indicator defined in Sect. 3.2 and $r_{ik} = 1$ indicates that building i remains unrepaired at the beginning of interval B_k. Thus, one building may contribute multiple training observations, but only up to the interval in which repair completion occurs.

For a building whose observed repair time falls within interval B_m, the corresponding response sequence is

$$y_{i1} = 0, \quad y_{i2} = 0, \quad \dots, \quad y_{i,m-1} = 0, \quad y_{im} = 1$$

No observations are generated for intervals after repair completion because the building is no longer at risk. If repair remains incomplete beyond the primary prediction horizon τ, the building contributes observations with

$$y_{ik} = 0, \quad k = 1, \dots, K,$$

and is treated as administratively right-censored at τ. This representation preserves the temporal structure of the repair process and avoids converting the observed repair time into a single multiclass label.

The input vector for a building–interval observation consists of the interval position on the repair-time axis, the information currently available for that building,

and, for updating stages, the probabilistic state obtained from the preceding stage. Let $z_i^{(s)}$ denote the cumulative building-specific predictors available at stage s. Because the information states are nested,

$$z_i^{(s)} = \left[z_i^{(s-1)}, \Delta z_i^{(s)}\right]$$

where $\Delta z_i^{(s)}$ contains the predictors newly available at stage s.

At the baseline stage, historical repair records are used to estimate the initial repair-completion hazard over the discrete repair-time axis. Because no current-event building-specific information is introduced at this stage, the baseline model represents the historical repair-completion pattern:

$$\widehat{h_{ik}^{(0)}} = M_0(k)$$

where k identifies the position of interval B_k along the repair-time axis. The resulting sequence

$$\widehat{h_i^{(0)}} = \left[\widehat{h_{i1}^{(0)}}, \widehat{h_{i2}^{(0)}}, \ldots, \widehat{h_{iK}^{(0)}}\right]$$

provides the baseline probabilistic state propagated to the first updating stage.

For each subsequent stage $s = 1,2,3$, the corresponding building–interval predictor vector is constructed as

$$x_{ik}^{(s)} = \left[k, \widehat{h_{ik}^{(s-1)}}, z_i^{(s-1)}, \Delta z_i^{(s)}\right]$$

The stage-specific LightGBM model then estimates,

$$\widehat{h_{ik}^{(s)}} = M_s\left(x_{ik}^{(s)}\right) = P\left(y_{ik} = 1 \,\middle|\, r_{ik} = 1, \widehat{h_{ik}^{(s-1)}}, z_i^{(s-1)}, \Delta z_i^{(s)}\right) \quad (16)$$

Operationally, $\widehat{h_{ik}^{(s-1)}}$ represents the component of the preceding hazard trajectory corresponding to interval B_k. Across all k building–interval observations, these interval-specific values collectively propagate the complete preceding hazard trajectory $\widehat{h_i^{(s-1)}}$ into the current updating stage.

This formulation makes the sequential updating mechanism explicit. At stage 1, the historical baseline hazard is combined with newly available building attributes:

$$\widehat{h_{ik}^{(1)}} = M_1\left(k, \widehat{h_{ik}^{(0)}}, z_i^{(1)}\right)$$

At stage 2, the updated stage-1 hazard and the cumulative building information are retained, while newly available ground-motion and site-condition information is incorporated:

$$\widehat{h_{ik}^{(2)}} = M_2\left(k, \widehat{h_{ik}^{(1)}}, z_i^{(2)}\right)$$

At Stage 3, the preceding hazard trajectory and all previously available information are retained, while repair-cost and assessment-lag information are incorporated after the corresponding post-earthquake assessment becomes available. Let u_i =$L_{A,i}$ denote the assessment-information availability time for building i. The Stage 3 update is applied only to intervals satisfying

$$\mathcal{K}_i^{(3)} = \{k: r_{ik} = 1,\ t_{k-1} \geq u_i\}$$

For these intervals,

$$\widehat{h_{ik}^{(3)}} = M_3\left(k, \widehat{h_{ik}^{(2)}}, z_i^{(3)}\right)$$

Accordingly, post-assessment predictors are not used retrospectively for intervals preceding u_i. For those earlier intervals, the Stage 2 hazard estimates are retained, ensuring that Stage 3 only updates subsequent at-risk repair-time intervals using information available at that time. This implementation prevents temporal information leakage and preserves the practical sequence of information acquisition.

At each updating stage, the LightGBM model estimates the current-stage hazard trajectory using the cumulative information available at that stage and the propagated probabilistic state from the preceding stage. Therefore, the stage-specific models do not represent independent predictions with increasingly large feature sets; instead, each model updates the propagated probabilistic state using newly available information.

LightGBM is trained using a binary probabilistic objective. For information stage s, the model parameters are estimated by minimizing the binary negative log-likelihood over the at-risk building–interval observations:

$$\mathcal{L}_s = -\sum_i \sum_{k:r_{ik}=1} \left[y_{ik} \log\left(\widehat{h_{ik}^{(s)}}\right) + (1 - y_{ik}) \log\left(1 - \widehat{h_{ik}^{(s)}}\right) \right] \quad (17)$$

Minimizing this loss encourages the predicted value $\widehat{h_{ik}^{(s)}}$ to approximate the conditional probability of repair completion during interval B_k, given that the building remains unrepaired at the start of the interval. Thus, the model output has a direct probabilistic interpretation as a discrete-time hazard rather than as a multiclass membership probability.

After interval-specific hazards have been predicted for all k intervals, they are reorganized by building to reconstruct the stage-specific hazard trajectory,

$$\widehat{h_i^{(s)}} = \left[\widehat{h_{i1}^{(s)}}, \ldots, \widehat{h_{iK}^{(s)}} \right]$$

The corresponding survival and cumulative repair-completion probabilities are then obtained using the transformations defined in Sect. 3.2:

$$\widehat{h_i^{(s)}} \rightarrow \widehat{S_i^{(s)}}(t_k) \rightarrow \widehat{F_i^{(s)}}(t_k)$$

The resulting cumulative repair-completion distribution is subsequently used for engineering reporting and probabilistic evaluation as described in Sect. 3.4.

To prevent information leakage caused by the building–time-interval expansion, data partitioning is performed at the building level before expansion into interval-level observations**.** All interval observations generated from the same building therefore remain within the same training, validation, or test subset. The validation data are used for model selection and hyperparameter tuning, whereas the hold-out test data remain independent of model development. This is particularly important in the discrete-time formulation because random splitting after interval expansion could otherwise place observations from the same building in both the training and evaluation datasets.

The four LightGBM models $M0$,$M1$, $M2$ and$M3$ are trained offline using their corresponding information states. During application, the term *updating* does not imply retraining a LightGBM model whenever new information about a building becomes available. Instead, when the information state changes, the probabilistic

504 output from the preceding stage and the newly available predictors are supplied to the
505 pre-established model corresponding to the next information stage. The updated
506 hazard trajectory is then propagated forward until additional information becomes
507 available.

508 Among the candidate base-learning algorithms considered in this study,
509 LightGBM is retained as the common learner because the original model comparison
510 showed the lowest mean IBS among logistic regression, random forest, LightGBM,
511 XGBoost, and CatBoost. The same learning algorithm is used across stages so that
512 changes in prediction performance can primarily be attributed to the progressive
513 availability and propagation of information rather than to differences in model
514 families.

515 **3.4 Probabilistic outputs and evaluation**

516 The discrete-time hazard formulation produces a stage-specific cumulative
517 repair-completion distribution for each building. For information stage s, the
518 distribution is represented as

519 $$F_i^{(s)} = \left[F_i^{(s)}(t_1), F_i^{(s)}(t_2), \dots, F_i^{(s)}(t_K)\right],$$

520 where

521 $$F_i^{(s)}(t_k) = P\left(T_i \le t_k \mid \mathcal{J}_i^{(s)}\right).$$

522 This fine-resolution cumulative distribution, rather than the four broader reporting
523 ranges, is the primary probabilistic output of the proposed framework. For
524 engineering interpretation, point summaries, uncertainty measures, and building-stock
525 recovery quantities are subsequently derived from this distribution.

526 The evaluation criteria are selected according to their specific purposes.
527 Probability scoring rules assess agreement between predicted probabilities and
528 observed repair outcomes (Gneiting and Raftery 2007); discrimination metrics
529 evaluate ranking performance at recovery milestones; uncertainty measures quantify
530 the concentration or spread of predicted distributions; and updating metrics
531 characterize probability redistribution between successive information stages.

Therefore, reduced uncertainty or substantial probability redistribution alone is not interpreted as evidence of improved predictive accuracy.

3.4.1 Engineering reporting probabilities

For practical interpretation, the cumulative repair-completion distribution is summarized into four broader ordered repair-time ranges:

$$R_1: T_i \leq 1\ year,$$

$$R_2: 1 < T_i \leq 2\ years,$$

$$R_3: 2 < T_i \leq 3\ years,$$

$$R_4: T_i > 3\ years.$$

The corresponding stage-specific reporting probabilities are derived from the cumulative repair-completion distribution as

$$q_{i1}^{(s)} = F_i^{(s)}(1),$$

$$q_{i2}^{(s)} = F_i^{(s)}(2) - F_i^{(s)}(1),$$

$$q_{i3}^{(s)} = F_i^{(s)}(3) - F_i^{(s)}(2),$$

$$q_{i4}^{(s)} = 1 - F_i^{(s)}(3).$$

Thus,

$$q_i^{(s)} = \left[q_{i1}^{(s)}, q_{i2}^{(s)}, q_{i3}^{(s)}, q_{i4}^{(s)}\right],$$

$$q_{ir}^{(s)} \geq 0, \quad \sum_{r=1}^{4} q_{ir}^{(s)} = 1.$$

These four probabilities are derived reporting quantities rather than direct categorical outputs of LightGBM. The underlying models estimate interval-specific repair-completion hazards, from which the cumulative distribution and reporting-range probabilities are subsequently reconstructed. Therefore, the four reporting probabilities preserve the temporal ordering of repair completion and do not represent mutually exclusive classes predicted by the machine-learning model.

3.4.2 Most likely repair-time range

For engineering interpretation, the most likely repair-time reporting range derived from the reporting probabilities at information stage s is defined as

$$r_i^{(*(s))} = \arg \max_{(r\in 1,2,3,4)} q_{(ir)}^{(s)},$$

with the corresponding range

$$R_i^{*(s)} = R_{r_i^{*(s)}}$$

The corresponding maximum reporting probability is

$$q_i^{*(s)} = \max_{r\in\{1,2,3,4\}} q_{ir}^{(s)}.$$

The observed reporting range of building i is denoted by $c_i \in \{1,2,3,4\}$. Overall accuracy is then calculated as

$$Accuracy^{(s)} = \frac{1}{N}\sum_{i=1}^{N} I\left(r_i^{*(s)} = c_i\right).$$

Because the four reporting ranges may have unequal sample sizes, range-specific recall is additionally evaluated as

$$Recall_r^{(s)} = \frac{\sum_{i=1}^{N} I\left(c_i = r, r_i^{*(s)} = r\right)}{\sum_{i=1}^{N} I(c_i = r)}$$

A row-normalized confusion matrix is additionally used to examine whether incorrect most-likely-range predictions occur predominantly between adjacent reporting ranges or across more widely separated repair-time ranges.

3.4.3 Median repair time and prediction interval

Point summaries of repair time are derived from the fine-resolution cumulative repair-completion distribution rather than from the four reporting ranges. For a probability level p, the repair-time quantile is obtained from the cumulative distribution. Because the prediction horizon is finite, the required probability level may not always be reached within$[0, \tau]$. The horizon-truncated quantile is therefore defined as

$$\widetilde{T_{p,i}^{(s)}} = \begin{cases} \inf\left\{t_k : F_i^{(s)(t_k)} \geq p\right\}, & F_i^{(s)(\tau)} \geq p \\ \tau, & F_i^{(s)(\tau)} < p \end{cases}$$

If the cumulative repair-completion probability does not reach p by the prediction horizon τ, the corresponding quantile is truncated at τ for summary purposes. The predicted median repair time is then

$$\widetilde{T_{50,i}^{(s)}} = \widetilde{T_{0.50,i}^{(s)}}$$

The median-time mean absolute error and bias are calculated as

$$\mathrm{MAE}_{50}^{(s)} = \frac{1}{N}\sum_{i=1}^{N}\left|\widetilde{T_{50,i}^{(s)}} - T_i\right|,$$

$$\mathrm{Bias}_{50}^{(s)} = \frac{1}{N}\sum_{i=1}^{N}\left(\widetilde{T_{50,i}^{(s)}} - T_i\right).$$

Similarly, the lower and upper bounds of the central 80% prediction interval are

$$\widetilde{T_{10,i}^{(s)}} = \widetilde{T_{0.10,i}^{(s)}},$$

$$\widetilde{T_{90,i}^{(s)}} = \widetilde{T_{0.90,i}^{(s)}};$$

and its width is

$$W_{80,i}^{(s)} = \widetilde{T_{90,i}^{(s)}} - \widetilde{T_{10,i}^{(s)}}.$$

Thus, both the median repair time and the 80% prediction interval are derived from the fine-resolution cumulative repair-completion distribution. Quantiles that extend beyond the primary prediction horizon are represented by the horizon boundary τ.

3.4.4 Probabilistic accuracy over the repair-time horizon

Probabilistic accuracy is evaluated directly against the observed repair-completion status over time. At evaluation time t_k, the observed cumulative repair-completion state of building i is

$$Z_i(t_k) = I(T_i \le t_k)$$

The Brier score at t_k is

$$BS^{(s)}(t_k) = \frac{1}{N}\sum_{i=1}^{N}\left[F_i^{(s)}(t_k) - Z_i(t_k)\right]^2.$$

A lower Brier score indicates closer agreement between the predicted probability of repair completion by t_k and the corresponding observed binary outcome (Brier 1950).

For equally spaced evaluation times, this reduces to the average Brier score over the k evaluation times:

$$IBS^{(s)} = \frac{1}{K}\sum_{k=1}^{K} BS^{(s)}(t_k) \tag{18}$$

A lower IBS therefore indicates better agreement between the complete predicted cumulative repair-completion trajectory and the observed repair process over the considered time horizon (Graf et al. 1999).

This interpretation is fundamentally different from measuring whether a distribution is merely more concentrated. A concentrated distribution can still have a poor IBS if it concentrates probability at incorrect repair times.

3.4.5 Ranked probability score for the four reporting ranges

Because the four engineering reporting ranges are ordered, the ranked probability score (RPS) is used as a complementary measure of the accuracy of the derived four-range probability distribution (Epstein 1969)**.**

Define the cumulative predicted probability through reporting range as

$$Q_{im}^{(s)} = \sum_{r=1}^{m} q_{ir}^{(s)}, \quad m = 1,2,3,$$

and the corresponding observed cumulative indicator as

$$O_{im} = I(c_i \le m).$$

The building-specific RPS is

$$RPS_i^{(s)} = \frac{1}{3}\sum_{m=1}^{3}\left(Q_{im}^{(s)} - O_{im}\right)^2 \tag{19}$$

and the overall RPS is

$$RPS^{(s)} = \frac{1}{N}\sum_{i=1}^{N} RPS_i^{(s)}.$$

A lower RPS indicates closer agreement between the predicted and observed ordered reporting-range outcomes. Because the score is based on cumulative probabilities across ordered ranges, probability assigned to a range far from the observed range generally produces a larger penalty than probability assigned to an adjacent range.

The roles of IBS and RPS are therefore distinct:

$$IBS \rightarrow \text{fine-resolution cumulative repair-completion trajectory}$$

whereas

$$RPS \rightarrow \text{derived four-range reporting distribution}$$

3.4.6 Milestone-specific Brier score and AUC

In addition to evaluation over the complete repair-time horizon, prediction performance is examined at the practically interpretable 1-, 2-, and 3-year recovery milestones.

For a milestone $t^* \in \{1,2,3\}$, the predicted score for building is

$$F_i^{(s)}(t^*)$$

and the corresponding observed binary outcome is

$$Z_i(t^*) = I(T_i \leq t^*).$$

The milestone-specific Brier score is therefore

$$BS^{(s)}(t^*) = \frac{1}{N}\sum_{i=1}^{N}\left[F_i^{(s)}(t^*) - Z_i(t^*)\right]^2.$$

Brier scores are primarily compared across successive information stages at the same milestone, because their absolute values can depend on the prevalence of repair completion at different evaluation times.

The area under the receiver operating characteristic curve (AUC) is additionally calculated using $F_i^{(s)}(t^*)$ as the prediction score and $Z_i(t^*)$ as the observed binary label. AUC measures the ability of the model to rank buildings according to whether repair has been completed by the specified milestone. Higher AUC values indicate

stronger discrimination, whereas $AUC = 0.5$ corresponds to discrimination no better than random ranking.

Importantly, AUC is interpreted as a **d**iscrimination measure rather than a probability-calibration measure**.** Accordingly, Brier score and IBS are used to evaluate probabilistic accuracy, whereas AUC provides complementary information about ranking ability.

3.4.7 Predictive uncertainty

Predictive uncertainty is characterized separately from predictive accuracy. Three complementary summaries are considered: normalized entropy of the four reporting probabilities, the width of the 80% prediction interval, and the prevalence of a dominant reporting-range probability.

For building i, normalized predictive entropy is defined as

$$H_i^{(s)} = -\frac{\sum_{r=1}^{4} q_{ir}^{(s)} \ln q_{ir}^{(s)}}{\ln 4},$$

with the convention

$$0 \ln 0 = 0.$$

Thus,

$$0 \le H_i^{(s)} \le 1.$$

Values closer to 0 indicate that probability mass is concentrated in fewer reporting ranges, whereas values closer to 1 indicate a more diffuse distribution across the four ranges.

$$q_i^{*(s)} = \max_r q_{ir}^{(s)} \ge 0.50.$$

A building is defined as having a dominant reporting range when

$$P_{\text{dom}}^{(s)} = \frac{1}{N} \sum_{i=1}^{N} I\left(q_i^{*(s)} \ge 0.50\right).$$

A smaller W_{80}, lower entropy, or larger P_{dom} indicates a more concentrated prediction. However, these quantities characterize predictive concentration only and

are not interpreted as evidence of improved probabilistic accuracy unless accompanied by improved probability scores such as IBS or RPS.

3.4.8 Magnitude and effectiveness of successive probability updates

The effect of newly available information is further examined by comparing predictions between successive information stages. Three transitions are considered:

$$0 \to 1, \quad 1 \to 2, \quad 2 \to 3.$$

The magnitude of probability redistribution is quantified using total variation distance between the derived four-range probability vectors:

$$TV_i^{(s-1\to s)} = \frac{1}{2}\sum_{r=1}^{4}\left|q_{ir}^{(s)} - q_{ir}^{(s-1)}\right| \tag{20}$$

The TV distance satisfies

$$0 \le TV_i^{(s-1\to s)} \le 1.$$

A value of 0 indicates that the four-range probability allocation is unchanged, whereas larger values indicate greater redistribution of probability mass.

For a specified threshold a, the proportion of buildings whose probability distribution changes by more than a is explicitly calculated as

$$P_{TV>a}^{(s-1\to s)} = \frac{1}{N}\sum_{i=1}^{N} I\left(TV_i^{(s-1\to s)} > a\right).$$

For example, a =0.05 and a =0.10 are used to summarize the proportions of buildings experiencing relatively noticeable probability redistribution.

However, a large TV only indicates that the prediction changed substantially; it does not indicate that the change was beneficial.

To evaluate whether an update improves the four-range probabilistic prediction for an individual building, define

$$\Delta RPS_i^{(s)} = RPS_i^{(s-1)} - RPS_i^{(s)}$$

if

$$\Delta RPS_i^{(s)} > 0,$$

the updated distribution has a lower RPS and is therefore closer to the observed ordered repair-time outcome. The building is then classified as benefiting from that update.

The proportion of buildings benefiting from update $s-1 \rightarrow s$ is calculated explicitly as

$$P_{\text{benefit}}^{(s-1\rightarrow s)} = \frac{1}{N}\sum_{i=1}^{N} I\left(\Delta RPS_i^{(s)} > 0\right).$$

Thus, TV and benefited proportion answer two different questions:

$$TV\text{:How much did the prediction change?}$$

and

$$P_{\text{benefit}}\text{:How often did the change improve prediction?}$$

3.4.9 Changes in the most likely reporting range

Finally, changes in the most likely reporting range are examined separately from changes in the complete probability distribution.

For building i, define

$$C_i^{(s-1\rightarrow s)} = I\left(r_i^{*(s)} \neq r_i^{*(s-1)}\right).$$

The proportion of buildings whose most likely reporting range changes is therefore

$$P_{\text{change}}^{(s-1\rightarrow s)} = \frac{1}{N}\sum_{i=1}^{N} C_i^{(s-1\rightarrow s)}.$$

A row-normalized transition matrix is additionally constructed. For preceding range a and updated range b, its element is

$$A_{ab}^{(s-1\rightarrow s)} = \frac{\sum_{i=1}^{N} I\left(r_i^{*(s-1)} = a,\ r_i^{*(s)} = b\right)}{\sum_{i=1}^{N} I\left(r_i^{*(s-1)} = a\right)}$$

Diagonal elements represent buildings whose most likely reporting range remains unchanged, whereas off-diagonal elements represent transitions between ranges. Because the four reporting ranges are temporally ordered, the matrix is used to distinguish local changes between adjacent ranges from larger shifts across multiple ranges.

3.4.10 Building-stock cumulative repair-completion curve

To evaluate the probabilistic prediction at the building-stock scale, the mean predicted cumulative repair-completion probability is calculated as

$$\overline{F^{(s)}}(t_k) = \frac{1}{N}\sum_{i=1}^{N} F_i^{(s)}(t_k).$$

The corresponding empirical cumulative repair-completion proportion is

$$F_{\text{emp}}(t_k) = \frac{1}{N}\sum_{i=1}^{N} I(T_i \le t_k).$$

Comparison of

$$\overline{F^{(s)}}(t_k)$$

with

$$F_{\text{emp}}(t_k)$$

examines whether the aggregation of building-level probabilistic predictions reproduces the observed building-stock recovery trajectory.

4 Experimental Design and Results

The proposed framework was evaluated using building repair records from the 2010 Darfield and 2011 Christchurch earthquakes. The 2010 records were used to establish the historical baseline, whereas the 2011 Christchurch earthquake was treated as the target event for stage-wise probabilistic updating. Building attributes, ground-motion and site-condition information, and repair-cost and assessment-lag information were incorporated sequentially according to their availability. Model performance was evaluated in terms of point prediction, probabilistic accuracy, predictive uncertainty, and updating characteristics.

4.1 Experimental data and settings

The dataset contains 49,474 building repair records, including 15,189 records from the 2010 Darfield earthquake and 34,285 records from the 2011 Christchurch earthquake. Repair times range from 0.515 to 4.684 years for the 2010 records and from 0.208 to 4.216 years for the 2011 records.

The 2010 Darfield earthquake records were used to establish the historical baseline repair-completion model, whereas the 2011 Christchurch earthquake was treated as the target event for stage-wise probabilistic updating. The 2011 dataset was divided at the building level into training, validation, and independent hold-out test sets using proportions of 70%, 15%, and 15%, corresponding to 23,999, 5,143, and 5,143 buildings, respectively, as summarized in Table 1. The validation set was used for model validation and hyperparameter tuning, while the hold-out test set was reserved exclusively for final independent performance evaluation.

A uniform interval width of 0.25 years was adopted over a prediction horizon of 4.25 years, resulting in 17 discrete repair-time intervals:

$$t_k = 0.25k, \quad k = 1, \ldots, 17.$$

These intervals are used to construct the building–time-interval observations required by the discrete-time hazard model rather than as 17 categorical repair-time classes. Records with repair times exceeding 4.25 years are treated as right-censored at the prediction horizon.

To prevent information leakage, the building-level split was performed before expansion into interval-level observations. Preceding-stage hazard predictions used to train subsequent updating models were generated using three-fold building-level out-of-fold prediction, ensuring that the preceding-stage hazard supplied for each training building was obtained from a model that had not been fitted using that building. Validation and test predictions were generated only by models fitted on the training data.

For the Stage 3 model, repair-cost and assessment-lag predictors were introduced only for post-assessment intervals in which the building remained at risk of repair completion; intervals preceding the assessment time retained the Stage 2 hazard estimates.

Table 1 Dataset partitioning for the 2011 target earthquake event

Dataset	Sample size	Proportion	Purpose
Training set	23,999	70%	Model training and parameter estimation
Validation set	5143	15%	Model validation and hyperparameter tuning
Hold-out test set	5,143	15%	Final independent performance evaluation
Total	34,285	100%	/

4.2 Base-learner selection and model settings

To select a common base learner for the discrete-time repair-completion framework, five candidate algorithms—LightGBM, XGBoost, CatBoost, logistic regression, and random forest—were compared using three-fold cross-validation performed at the building level under the most information-rich condition. For each algorithm, the model was trained on the building–time-interval data described in Sect. 3.2 to estimate the interval-specific conditional repair-completion hazard,

$$\widehat{h_{ik}}$$

The predicted hazards were then transformed into survival probabilities,

$$\widehat{S_i}(t_k) = \prod_{j=1}^{k} \left(1 - \widehat{h_{ij}}\right),$$

and cumulative repair-completion probabilities,

$$\widehat{F_i}(t_k) = 1 - \widehat{S_i}(t_k).$$

The integrated Brier score (IBS), calculated from the reconstructed cumulative repair-completion distributions over the full prediction horizon, was used as the primary criterion for base-learner selection. Therefore, the comparison evaluates the probabilistic accuracy of the complete repair-completion trajectory rather than categorical classification performance.

As shown in Table 2, LightGBM achieved the lowest mean IBS of 0.112430 under three-fold building-level cross-validation, followed by XGBoost (0.112685) and CatBoost (0.113784). Logistic regression and random forest yielded higher mean IBS values of 0.116776 and 0.122420, respectively. The reported mean and standard deviation of IBS were calculated across the three cross-validation folds. The relatively small standard deviations indicate stable probabilistic performance across the folds.

Accordingly, LightGBM was selected as the common base learner for the baseline model $M0$ and the three stage-specific updating models $M1$,$M2$, and $M3$. Using the same base learner across all information stages allows changes in prediction performance to be attributed primarily to the progressive incorporation of newly available information and the propagation of the preceding probabilistic state, rather than to differences in model families.

Table 2 Comparison of candidate base learners based on IBS under three-fold building-level cross-validation

Algorithm	Mean IBS	Standard deviation
LightGBM	**0.112430**	0.001431
XGBoost	0.112685	0.001416
CatBoost	0.113784	0.001258
Logistic Regression	0.116776	0.001427
Random Forest	0.122420	0.001005

LightGBM achieved the lowest mean IBS, slightly outperforming XGBoost and CatBoost and clearly outperforming logistic regression and random forest. It was therefore selected as the common base learning algorithm for the initial probabilistic prediction and all subsequent updating stages. For all candidate algorithms, IBS was calculated after transforming the predicted interval-specific hazards into cumulative repair-completion probabilities, rather than from direct repair-time category predictions.

4.3 Point-summary prediction performance

Point-summary performance was evaluated across the four information stages using two quantities derived from the predicted cumulative repair-completion distribution: the most likely repair-time reporting range and the median repair time $\widetilde{T_{50}}$. The former is obtained from the four reporting probabilities derived from $\widehat{F_i^{(s)}}(t)$, whereas the latter is determined from the fine-resolution cumulative repair-completion trajectory. These quantities are therefore used only as summaries of the probabilistic prediction rather than as direct model-training targets.

As shown in Table 3, both measures improve consistently as additional information becomes available. Accuracy of the most likely reporting range increases from 43.78% at Stage 0 to 53.07%, 65.81%, and 85.42% at Stages 1–3, respectively. Meanwhile, the MAE of $\widetilde{T_{50}}$ decreases from 0.879 years to 0.689, 0.534, and 0.385 years. These results indicate that the stage-wise updating process progressively improves both identification of the most probable engineering reporting range and estimation of the central repair time.

The bias of $\widetilde{T_{50}}$ changes from −0.498 years at Stage 0 to 0.056, 0.014, and 0.076 years at the three subsequent stages. The large negative bias at the baseline stage indicates systematic underestimation of repair time when only historical information is available. After building-specific and earthquake-specific information is incorporated, the bias becomes substantially smaller and remains close to zero. Although Stage 2 yields the smallest absolute bias, Stage 3 achieves the lowest MAE and therefore provides the best overall point-summary agreement with the observed repair times.

Because the four reporting ranges contain unequal numbers of buildings, range-specific Recall was additionally examined for the final information stage. As shown in Table 4, Recall is 80.22%, 90.01%, 82.00%, and 82.00% for the $0-1$,$1-2$, $2-3$, and 3-year ranges, respectively. The relatively balanced Recall values indicate that the high final-stage Accuracy is not achieved solely by favoring the most frequent reporting range.

Fig.3 presents the row-normalized confusion matrix for the final-stage most-likely-range prediction. The diagonal values show that correct identification rates are high for all four reporting ranges, especially for the $1-2$ year range (90.0%). Off-diagonal values indicate that most misclassifications occur between adjacent reporting ranges rather than across multiple ranges. For example, among buildings with observed repair times of $0-1$ year, 19.8% are misclassified into the neighboring $1-2$year range, while none are misclassified into the $2-3$ year or 3-year ranges. Similarly, for observed repair times of $2-3$years and years, misclassifications are concentrated mainly in the adjacent ranges. This pattern indicates that, even when the exact most likely reporting range is not correctly identified, the predicted repair-time range usually remains temporally close to the observed outcome.

Taken together, Table 3, Table 4, and Fig. 3 show that the proposed stage-wise updating framework improves both the accuracy and robustness of point-summary repair-time prediction, while preserving the ordinal structure of the reporting ranges.

Table 3 Point-summary prediction performance across successive information stages

stage	Accuracy	$\widetilde{T_{50}}$ MAE/year	$\widetilde{T_{50}}$ Bias/year
0	43.78%	0.878653	-0.498405
1	53.07%	0.689245	0.056145
2	65.81%	0.534270	**0.014490**
3	**85.42%**	**0.384562**	0.076196

Table 4 Range-specific performance of the most likely repair-time reporting range at the final information stage

Observed repair-time range	Number of buildings	Correctly predicted	Recall
0-1	239	192	80.22%
1-2	2251	2026	90.01%
2-3	1141	936	82.00%
＞3	1512	1240	82.00%
Overall	5143	4394	Accuracy = 85.4%

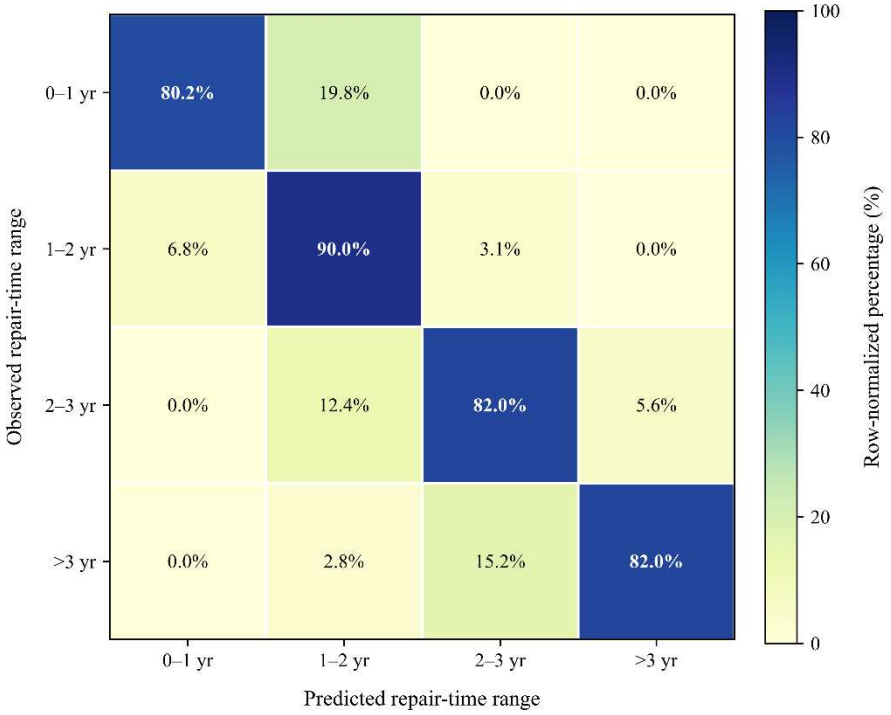


Fig. 3 Row-normalized confusion matrix of the most likely repair-time reporting range at the final information stage.

4.4 Probabilistic prediction performance

Probabilistic prediction performance was evaluated using the complete repair-completion distributions reconstructed from the predicted discrete-time hazards. Two complementary scoring rules were considered. The integrated Brier score (IBS) evaluates the agreement between the predicted cumulative repair-completion

trajectory $\widehat{F_i^{(s)}}(t_k)$ and the observed repair-completion status over the full prediction horizon, whereas the ranked probability score (RPS) evaluates the derived probability distribution across the four ordered engineering reporting ranges. Lower values of both metrics indicate better probabilistic prediction performance, although they assess different representations of the predicted repair-time distribution.

As shown in Table 5, probabilistic prediction performance improves consistently as additional information becomes available. The IBS decreases from 0.118576 at Stage 0 to 0.092618, 0.073747, and 0.055423 at Stages 1–3, respectively, corresponding to an overall reduction of approximately 53.3%. The RPS similarly decreases from 0.181270 to 0.136475, 0.106633, and 0.072000, representing an overall reduction of approximately 60.3%.

The simultaneous reductions in IBS and RPS indicate improvement at two different probabilistic levels. The reduction in IBS shows that the predicted cumulative repair-completion probabilities become progressively more consistent with the observed repair process over the full repair-time horizon. In contrast, the reduction in RPS indicates improved probability allocation across the four ordered reporting ranges. Thus, the improvement is not limited to identifying a more likely repair-time range; it also reflects better agreement of the complete temporal probability structure with the observed outcomes.

The largest improvements occur after the incorporation of earthquake-specific and subsequent assessment information. Relative to the preceding stage, adding building attributes reduces IBS from 0.118576 to 0.092618, while the addition of ground-motion and site-condition information further reduces it to 0.073747. The final incorporation of repair-cost and assessment-lag information produces the lowest IBS of 0.055423. A similar monotonic pattern is observed for RPS. These results demonstrate that the stage-wise updating mechanism progressively refines both the fine-resolution cumulative repair-completion distribution and the derived engineering reporting probabilities.

To further examine probabilistic prediction at practically relevant recovery milestones, Brier scores were calculated at 1, 2, and 3 years using the cumulative repair-completion probabilities $\widehat{F_i^{(s)}}(1)$, $\widehat{F_i^{(s)}}(2)$, and $\widehat{F_i^{(s)}}(3)$, respectively. As shown in Fig. 4(a), Brier scores generally decrease across successive information stages at all three milestones, with more pronounced improvements at the 2- and 3-year milestones. This indicates that the incorporation of newly available information improves probability estimates not only over the full repair-time horizon but also at specific short-, medium-, and longer-term recovery milestones.

Fig. 4(b) presents the corresponding AUC values at 1, 2, and 3 years. In contrast to the Brier score, AUC is used only to evaluate discrimination between buildings that have and have not completed repair by a specified milestone. The AUC values generally increase as information is progressively incorporated, with the baseline model showing the weakest discrimination and the subsequent updating stages providing progressively stronger separation.

Taken together, the reductions in IBS, RPS, and milestone-specific Brier scores, together with the increases in AUC, demonstrate that the proposed stage-wise updating framework progressively improves both probabilistic accuracy and discrimination as information becomes available. Importantly, these results evaluate agreement with observed repair outcomes and are distinct from the predictive-concentration measures examined separately in Sect. 4.5.

Table 5 Probabilistic prediction performance across successive information stages

Information stage	IBS	RPS
0	0.118576	0.181270
1	0.092618	0.136475
2	0.073747	0.106633
3	**0.055423**	**0.072000**

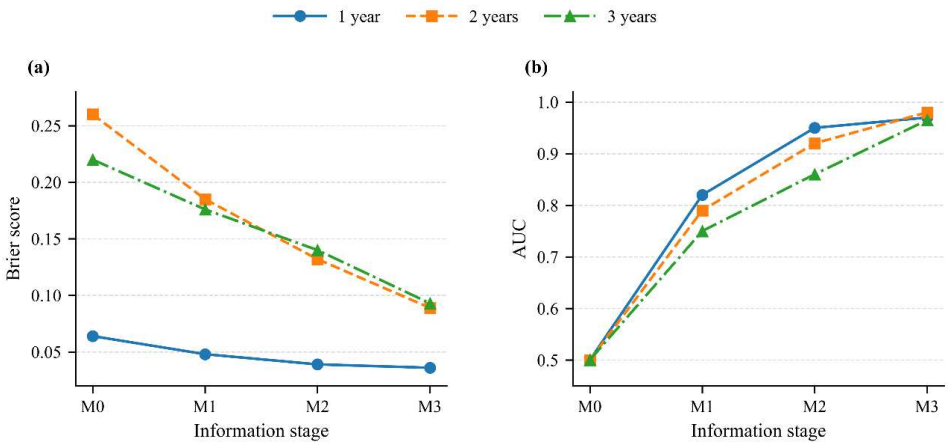


Fig. 4 Probabilistic prediction performance at the 1-, 2-, and 3-year repair-completion milestones across successive information stages: (a) Brier score and (b) AUC

4.5 Prediction uncertainty across successive information stages

Predictive uncertainty was examined separately from probabilistic accuracy using three complementary measures defined in Sect. 3.4: normalized predictive entropy, the width of the 80% prediction interval (W_{80}), and the proportion of buildings with a dominant repair-time reporting probability of at least 0.50. These measures characterize the concentration and temporal spread of the predicted repair-time distribution and are not interpreted, by themselves, as measures of probabilistic accuracy.

As shown in Fig. 5(a), the mean normalized predictive entropy decreases progressively from 0.943 at Stage 0 to 0.843, 0.765, and 0.707 at Stages 1–3, respectively. The initial median entropy is 0.955 and follows the same decreasing trend across the subsequent stages. Because entropy is calculated from the four derived reporting probabilities, this reduction indicates that probability mass becomes progressively more concentrated among the engineering reporting ranges as additional information becomes available.

Fig. 5(b) shows a similar reduction in the temporal spread of the predicted repair-completion distribution. The mean W_{80}, which is derived from the 10th and 90th percentiles of the fine-resolution cumulative repair-completion distribution, decreases from approximately 2.75 years at Stage 0 to 2.25, 2.00, and 1.81 years at Stages 1–3.

The corresponding median values decrease from approximately 2.73 to 2.24, 2.01, and 1.75 years. The progressive narrowing of W_{80} indicates that the range of plausible repair times becomes smaller as information accumulates.

A complementary pattern is observed in Fig. 5(c). The proportion of buildings for which the maximum reporting-range probability is at least 0.50 increases from 0% at Stage 0 to 26.2%, 37.2%, and 53.8% at Stages 1–3, respectively. Thus, a progressively larger fraction of buildings develops a clearly dominant repair-time reporting range as building-specific, earthquake-specific, and assessment-related information is incorporated.

Taken together, the decreasing entropy and W_{80}, together with the increasing proportion of buildings with a dominant reporting probability, show that the predicted repair-time distributions become progressively more concentrated and temporally narrower across successive information stages. These results characterize a reduction in predictive uncertainty rather than probabilistic accuracy itself. The latter is evaluated independently using IBS, RPS, and Brier scores in Sect. 4.4. The simultaneous improvement in probabilistic accuracy and reduction in predictive uncertainty therefore provides complementary evidence that the stage-wise updating framework progressively refines the predicted repair-completion process as additional information becomes available.

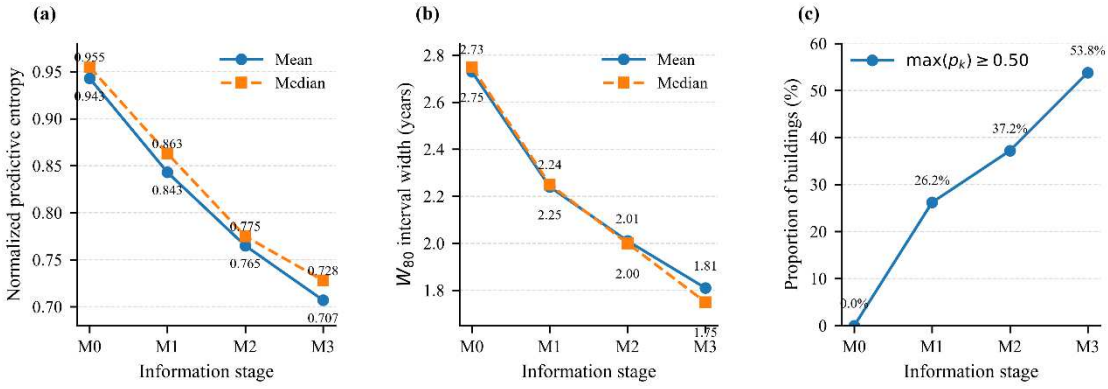


Fig. 5 Predictive uncertainty across successive information stages: (a) mean and median normalized predictive entropy; (b) mean and median width of the 80% prediction interval; and (c) proportion of buildings with a maximum repair-time reporting-range probability of at least 0.50

978 **4.6 Probabilistic updating characteristics between successive information stages**

979 To examine how newly available information modifies existing repair-time
980 predictions, probabilistic changes were evaluated across the three successive updates,
981 $0-1$, $1-2$, and $2-3$. The analysis considers changes in both the most likely
982 repair-time reporting range and the complete four-range probability vector derived
983 from the cumulative repair-completion distribution.

984 Fig. 6 presents the row-normalized transition matrices of the most likely repair-
985 time reporting range. All three matrices are dominated by their main diagonals,
986 indicating that the most likely range remains unchanged for the majority of buildings
987 after each update. The diagonal proportions generally increase as information
988 accumulates, from approximately 63–74% for the $0-1$ update to 72–80% for $1-$
989 2and 79–86% for $2-3$. Off-diagonal transitions are concentrated mainly near the
990 main diagonal, showing that changes occur predominantly between adjacent reporting
991 ranges rather than across multiple ranges. This pattern suggests that successive
992 updates generally refine rather than radically alter the ordinal repair-time prediction.

993 Because the most likely range may remain unchanged even when its underlying
994 probabilities change substantially, total variation distance (TV) was further used to
995 quantify redistribution of the complete four-range probability vector. As shown in
996 Fig. 7(a), the median TV values for the $0-1$, $1-2$, and $2-3$ updates are 0.030,
997 0.138, and 0.092, respectively. The largest redistribution therefore occurs during the
998 $1-2$ update, indicating that the addition of ground-motion and site-condition
999 information produces the strongest modification of building-level repair-time
1000 probabilities.

1001 Fig. 7(b) further shows that 28.0% and 1.5% of buildings have $TV>0.05$ and
1002 $TV>0.10$, respectively, for the $0-1$update. These proportions increase to 93.5%
1003 and 67.1% for $1-2$, and are 79.4% and 45.0% for $2-3$. Thus, building attributes
1004 produce relatively limited probability redistribution, whereas earthquake-specific
1005 ground-motion and site-condition information alters the predicted distribution for

most buildings. Repair-cost and assessment-lag information also produces substantial adjustments, although their overall magnitude is smaller than that of the preceding update.

The magnitude of an update, however, does not necessarily indicate whether the prediction becomes more accurate. Therefore, the benefited proportion defined in Sect. 3.4 was used to assess update effectiveness. The proportions of buildings benefiting from the three successive updates are 54.5%, 69.9%, and 76.8%, respectively. Notably, although the median TV decreases from 0.138 for the $1-2$ update to 0.092 for the update$2-3$, the benefited proportion increases to its highest value in the final update. This indicates that repair-cost and assessment-lag information produces smaller probability redistributions than ground-motion and site-condition information, but these adjustments more frequently move the predicted distribution closer to the observed repair-time outcome.

Overall, the three information updates exhibit distinct effects. Building attributes provide relatively modest adjustments to the baseline prediction; ground-motion and site-condition information produces the largest redistribution of repair-time probabilities; and repair-cost and assessment-lag information produces smaller but more frequently beneficial refinements. These results demonstrate that the magnitude and effectiveness of probabilistic updating are complementary but distinct properties of the stage-wise updating process.

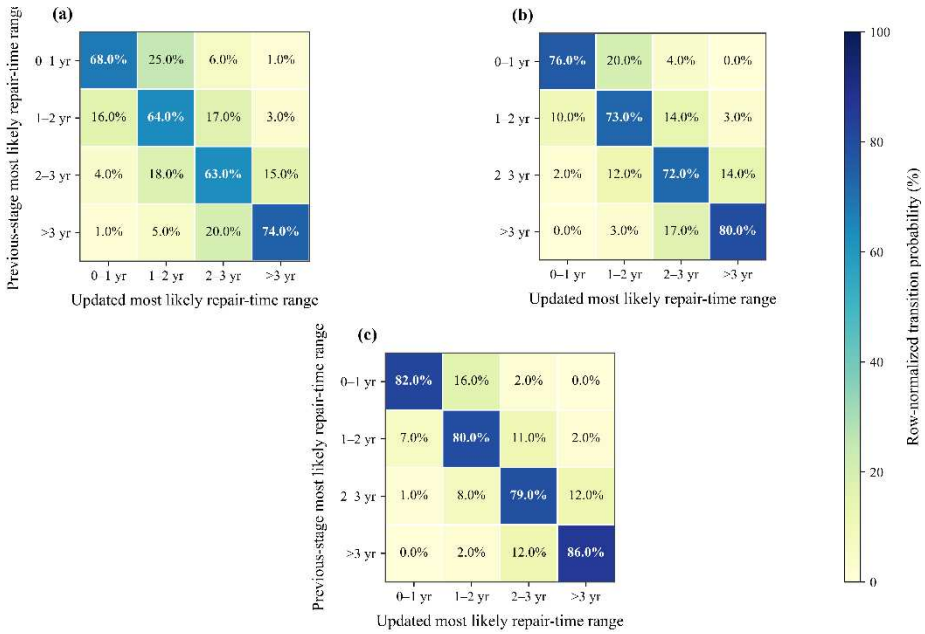


Fig. 6 Transition matrices of the most likely repair-time reporting range between successive information stages: (a) Stage 0 → Stage 1; (b) Stage 1 → Stage 2; and (c) Stage 2 → Stage 3

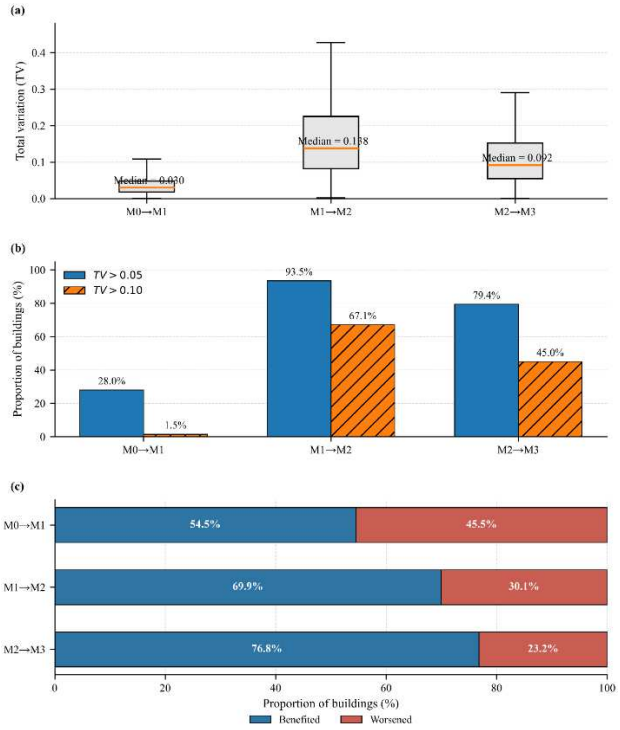


Fig. 7 Probabilistic updating characteristics between successive information stages: (a) total variation distance (TV) between the four-range probability distributions; (b) proportions of buildings exceeding selected TV thresholds; and (c) proportions of buildings benefiting from each probabilistic update

4.7 Building-stock repair-completion prediction

To further evaluate whether the building-level probabilistic predictions reproduce the overall recovery process, the predicted cumulative repair-completion probabilities were aggregated at the building-stock scale. For information stage s, the portfolio-average cumulative repair-completion probability is defined as

$$\overline{F^{(s)}}(t_k) = \frac{1}{N}\sum_{i=1}^{N} F_i^{(s)}(t_k),$$

where N denotes the number of buildings in the hold-out test set. The corresponding empirical cumulative repair-completion proportion is

$$F_{\mathrm{obs}}(t_k) = \frac{1}{N}\sum_{i=1}^{N} I(T_i \le t_k).$$

Fig.8 compares the predicted building-stock cumulative repair-completion curves at the four information stages with the observed repair-completion curve. At Stage 0, the prediction based only on historical repair information substantially underestimates the observed recovery trajectory over most of the repair-time horizon. As additional information becomes available, the predicted curves progressively move toward the observed curve.

The incorporation of building attributes at Stage 1 produces a moderate improvement, while the addition of ground-motion and site-condition information at Stage 2 results in a more pronounced upward adjustment of the predicted repair-completion trajectory. After repair-cost and assessment-lag information is incorporated at Stage 3, the predicted curve shows the closest agreement with the observed building-stock recovery process over most of the considered time horizon.

Nevertheless, the final-stage cumulative repair-completion probability remains below the observed curve during part of the medium- and long-term recovery period. This indicates that, although the stage-wise updating process substantially improves building-stock-level recovery prediction, the overall repair-completion speed is still underestimated to some extent.

Overall, the progressively improved agreement between the predicted and observed cumulative curves demonstrates that the proposed updating framework improves not only building-level repair-time predictions but also the representation of the aggregate repair-completion process at the building-stock scale.

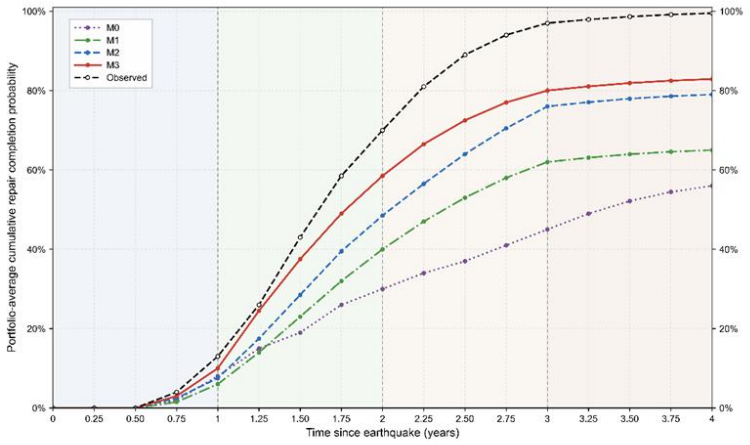


Fig.8 Predicted and observed cumulative repair-completion curves at the building-stock scale across successive information stages

4.8 Feature importance across successive hazard-updating stages

To examine how progressively available information contributes to stage-wise hazard updating, SHAP analysis was applied to the three LightGBM updating models (Lundberg and Lee 2017). Mean absolute SHAP values were aggregated over at-risk building–interval observations to characterize predictor importance. Fig. 9(a) shows all displayed Stage 1 predictors with non-zero importance, whereas Fig. 9(b) and (c) present the leading predictors at Stages 2 and 3 for clarity. The repair-time interval index and propagated preceding-stage hazard are excluded from the displayed rankings, although they remain model inputs. In Fig. 9, red labels denote engineered or building-specific derived variables constructed or matched during feature engineering, whereas black labels denote original variables directly available from the source datasets.

At Stage 1, only two building-related predictors exhibit non-zero mean absolute SHAP values. As shown in Fig. 9(a), the building area ratio $R_{A,i}$ has the largest mean absolute SHAP value of 0.0363, followed by the number of storeys with a value of

0.0262. The remaining Stage 1 predictors have mean absolute SHAP values of zero and are therefore not displayed. This result indicates that the Stage 1 hazard update is primarily driven by these two building-geometry descriptors among the available building attributes.

After earthquake-specific ground-motion and site-condition information becomes available, the feature-importance structure changes substantially. Fig. 9(b) presents the leading predictors ranked by mean absolute SHAP value. The rupture distance $R_{\mathrm{rup},i}$ is the most influential displayed predictor, with a mean absolute SHAP value of 0.1406, followed by PGA (0.0643), $Z_{1.0,i}$ (0.0614), $R_{\mathrm{jb},i}$ (0.0495), and $V_{S30,i}$ (0.0357). Liquefaction vertical displacement also shows a notable contribution (0.0267), whereas $R_{A,i}$ and the number of storeys have smaller displayed values of 0.0141 and 0.0088, respectively. These results indicate that earthquake-specific seismic-demand, source-distance, and site-condition information becomes more influential after its incorporation into the updating process.

At Stage 3, Fig. 9(c) shows the leading predictors with the largest mean absolute SHAP values. Assessment lag $L_{A,i}$ becomes the dominant displayed predictor, with a mean absolute SHAP value of 0.2401, followed by $R_{\mathrm{rup},i}$(0.0895) and the loss ratio LR_i (0.0816). Earthquake-specific variables remain important, including $Z_{1.0,i}$ (0.0599), PGA (0.0510), $R_{\mathrm{jb},i}$(0.0396), liquefaction vertical displacement (0.0273), and $V_{S30,i}$(0.0243). In contrast, $R_{A,i}$ and the number of storeys have relatively smaller mean absolute SHAP values of 0.0129 and 0.0084, respectively.

Overall, the SHAP results show a clear shift in dominant information across stages: building geometry is most relevant in the first update, earthquake-specific seismic and site information dominates the second, and assessment- and loss-related information becomes most influential in the final stage. Because mean absolute SHAP values measure contribution magnitude, they should not be interpreted as indicating the direction or causal effect of individual predictors.

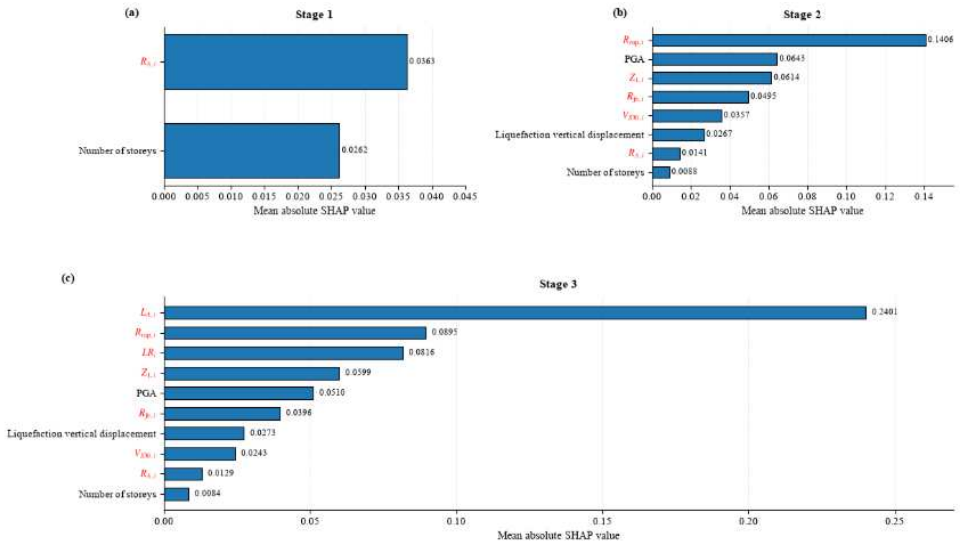


Fig. 9 Feature importance across successive hazard-updating stages based on mean absolute SHAP values: (a) all predictors with non-zero importance at Stage 1; (b) leading predictors at Stage 2; and (c) leading predictors at Stage 3. Red labels indicate engineered or building-specific derived variables, whereas black labels indicate original variables.

5 Conclusions and outlook

5.1 Main findings

This study addresses post-earthquake building repair-time prediction under evolving information conditions by developing an information-availability-driven stage-wise probabilistic updating framework**.** Rather than directly assigning buildings to repair-time categories, the proposed method represents repair completion as a discrete-time event process. LightGBM is used to estimate interval-specific conditional repair-completion hazards, which are subsequently transformed into survival probabilities and cumulative repair-completion distributions. Historical repair information provides the baseline probabilistic state, after which building attributes, ground-motion and site-condition information, and repair-cost and assessment-lag information are progressively incorporated as they become available. At each stage, the preceding hazard trajectory and cumulative information are retained and combined with newly available information to update the repair-completion prediction.

The results show consistent improvements in both point-summary and probabilistic prediction performance. Accuracy of the most likely repair-time

reporting range increases from 43.78% at Stage 0 to 85.42% at Stage 3, while the MAE of the median repair time $\widetilde{T_{50}}$ decreases from 0.879 to 0.385 years. Probabilistic accuracy improves similarly: IBS decreases from 0.118576 to 0.055423 and RPS from 0.181270 to 0.072000, corresponding to overall reductions of approximately 53.3% and 60.3%, respectively. These results demonstrate that progressively available information improves both the complete cumulative repair-completion trajectory and the derived engineering reporting probabilities.

Predictive uncertainty also decreases as information accumulates. The mean normalized predictive entropy decreases from 0.943 at Stage 0 to 0.707 at Stage 3, while the mean width of the 80% prediction interval decreases from approximately 2.75 to 1.81 years. Meanwhile, the proportion of buildings with a dominant reporting-range probability of at least 0.50 increases from 0% to 53.8%. Importantly, these uncertainty measures are interpreted separately from probabilistic accuracy; their simultaneous improvement indicates that the updated distributions become both more accurate and more concentrated.

The three updating stages exert different effects on the predicted distributions. Incorporating building attributes produces relatively limited changes, with a median TV of 0.030. The addition of ground-motion and site-condition information produces the largest redistribution, increasing the median TV to 0.138, whereas the final incorporation of repair-cost and assessment-lag information yields a smaller median TV of 0.092 but the highest benefited proportion of 76.8%. These results show that the magnitude of an update and its effectiveness are distinct: earthquake-specific information produces larger changes, whereas later assessment information more frequently improves individual probabilistic predictions.

At the building-stock scale, the predicted cumulative repair-completion curves progressively approach the empirical recovery trajectory as information is incorporated, with the final-stage prediction showing the closest agreement over most of the repair-time horizon. Some underestimation remains during the medium- and long-term recovery period, indicating that the aggregate recovery speed is not yet

fully captured. Feature-importance analysis further shows that the dominant explanatory information evolves with the updating stage: building characteristics are more influential during the early update, earthquake-specific demand and site variables become prominent after the event, and assessment-lag and loss-related variables become particularly important in the final stage.

Overall, the principal contribution of this study is not a single final repair-time model developed under complete-information conditions, but a stage-wise mechanism for propagating and revising repair-completion probabilities under evolving information availability**.** The framework establishes an initial probabilistic assessment under limited information and progressively refines the repair-completion hazard trajectory as additional information becomes available, providing probabilistic support for post-earthquake repair-progress assessment and recovery resource allocation.

5.2 Limitations

Several limitations should be considered when interpreting the results. First, the study is based primarily on insurance claims and managed-repair records from the Canterbury earthquake sequence. These data may not fully represent uninsured buildings, unreported minor-damage cases, demolished buildings, or repair projects outside the managed-repair system. Consequently, the conclusions are most directly applicable to building populations with similar insurance and repair-management contexts.

Second, repair time is defined as the calendar duration from earthquake occurrence to recorded managed-repair completion. This period may include damage investigation, insurance assessment, engineering design, permitting, waiting periods, repair construction, and completion registration. It should therefore be interpreted as an overall recovery-process duration rather than pure construction time.

Third, the four information stages represent changes in information availability rather than fixed post-earthquake calendar times. The time required to obtain ground-motion products, detailed assessments, repair-cost estimates, and other information

may vary among buildings and earthquake events. Therefore, the proposed framework represents information-availability-dependent updating rather than prediction at uniform post-earthquake time points.

Fourth, although the discrete-time formulation preserves the temporal ordering of repair completion, temporal resolution remains dependent on the selected interval width and prediction horizon. In the present implementation, repair time is discretized at 0.25-year intervals over a 4.25-year horizon. Predictions within the same interval cannot be distinguished without further temporal refinement, and observations extending beyond the prediction horizon are represented through the remaining survival probability rather than an explicitly modeled longer-term trajectory.

In addition, building-level seismic-demand and site parameters are partly obtained through strong-motion station matching and spatial interpolation, and their accuracy is affected by station density, interpolation assumptions, and local spatial variability. Several potentially important recovery-process variables—including contractor availability, labor and material supply, financing conditions, permitting progress, owner decisions, and aftershock effects—are also unavailable in the current dataset.

Finally, the framework has mainly been evaluated within a single earthquake sequence. Its transferability across regions, earthquake types, building stocks, and insurance-management systems requires independent external validation. The SHAP results should therefore be interpreted as predictive associations within the available data rather than as causal effects of individual variables on repair duration.

5.3 Future work

Future research should focus on improving data coverage, temporal modeling, probabilistic calibration, and practical recovery applications. First, repair datasets from additional earthquake events, regions, and insurance-management systems should be incorporated to conduct cross-event and spatially independent validation. Additional recovery-process information, including contractor and labor availability, material supply, financing arrangements, permitting and approval status, owner

decisions, and aftershock effects, would further improve representation of the mechanisms governing repair completion. These extensions are also consistent with the limitations already identified in the current dataset.

Second, the discrete-time hazard framework can be further developed by examining finer or adaptive time discretization, alternative hazard-learning algorithms, and continuous-time survival formulations. Particular attention should be given to probability calibration, uncertainty associated with model parameters, and prediction under incomplete repair observations. Ensemble learning, bootstrap-based uncertainty quantification, and calibration methods may provide more robust building-level and building-stock recovery distributions.

Third, the information-updating mechanism could be extended beyond the four information states considered here. Instead of assuming a fixed set of information stages, future implementations could accommodate additional intermediate states as inspection, engineering assessment, financing, permitting, contractor mobilization, and repair-progress information becomes available. This would allow the same retain–update–propagate mechanism to operate under more detailed recovery-information sequences.

Finally, the predicted cumulative repair-completion probabilities can be more closely integrated with recovery decision making. Building- and portfolio-level probabilities of completion within specified time horizons could support identification of buildings with high long-term repair risk, prioritization of inspections and assessments, evaluation of the value of collecting additional information, and allocation of repair and insurance-processing resources. Such developments would extend the framework from probabilistic repair-time prediction toward operational post-earthquake recovery planning and regional resilience assessment.

CRediT authorship contribution statement

Haihui Zhao: Investigation, Data curation. **Jianming Wang**: Supervision,Methodology. **Zifa Wang**: Project administration, Funding acquisition, Conceptualization. **Dengke Zhao**: Visualization,Validation. **Siran Yang**: Visualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Funding Declaration

This work was supported by the Basic Research Program (Special Fund for Fundamental Scientific Research) of the Institute of Engineering Mechanics, China Earthquake Administration (Grant No. 2023A01), and by the National Natural Science Foundation of China (General Program) (Grant No.52378544).

Data availability

The building repair data used in this study are not publicly available due to confidentiality and data-use restrictions. Processed data may be available from the corresponding author upon reasonable request and subject to applicable permissions. Ground-motion and site-condition data were obtained from publicly available GeoNet and NZSMDB resources.

References

Basaglia A, Aprile A, Spacone E, Pelà L (2020) Assessing community resilience, housing recovery and impact of mitigation strategies at the urban scale: a case study after the 2012 Northern Italy Earthquake. Bull Earthq Eng 18:6039–6074. https://doi.org/10.1007/s10518-020-00919-8

Bektaş N, Kegyes-Brassai O (2025) Developing a machine learning-based rapid visual screening method for seismic assessment of existing buildings on a case study data from the 2015 Gorkha, Nepal earthquake. Bull Earthq Eng 23:4981–5019. https://doi.org/10.1007/s10518-024-01924-x

Bradley BA (2013) A New Zealand-specific pseudospectral acceleration ground-motion prediction equation for active shallow crustal earthquakes based on foreign models. Bull Seismol Soc Am 103:1801–1822. https://doi.org/10.1785/0120120021

Bradley BA (2014) Site-specific and spatially-distributed ground-motion intensity estimation in the 2010–2011 Canterbury earthquakes. Soil Dyn Earthq Eng 61–62:83–91. https://doi.org/10.1016/j.soildyn.2014.01.025

Bradley BA, Hughes M (2013) Spatially-distributed ground motion intensity maps: application for site-specific liquefaction evaluations in Christchurch. In: Proceedings of the 2013 New Zealand Society for Earthquake Engineering conference, Wellington, New Zealand, 26–28 April 2013, paper 49, 8 pp

Brier GW (1950) Verification of forecasts expressed in terms of probability. Mon Weather Rev 78:1–3. https://doi.org/10.1175/1520-0493(1950)078<0001:VOFEIT>2.0.CO;2

Bruneau M, Chang SE, Eguchi RT, Lee GC, O'Rourke TD, Reinhorn AM, Shinozuka M, Tierney K, Wallace WA, von Winterfeldt D (2003) A framework to quantitatively assess and enhance the seismic resilience of communities. Earthq Spectra 19:733–752. https://doi.org/10.1193/1.1623497

Burton HV, Deierlein G, Lallemant D, Lin T (2016) Framework for incorporating probabilistic building performance in the assessment of community seismic

1302 resilience. J Struct Eng 142:C4015007.
1303 https://doi.org/10.1061/(ASCE)ST.1943-541X.0001321
1304 Comerio MC (2006) Estimating downtime in loss modeling. Earthq Spectra 22:349–
1305 365. https://doi.org/10.1193/1.2191017
1306 Cook DT, Liel AB, Haselton CB, Koliou M (2022) A framework for operationalizing
1307 the assessment of post-earthquake functional recovery of buildings. Earthq
1308 Spectra 38:1972–2007. https://doi.org/10.1177/87552930221081538
1309 De Martino G, Di Ludovico M, Prota A, Moroni C, Manfredi G, Dolce M (2017)
1310 Estimation of repair costs for RC and masonry residential buildings based on
1311 damage data collected by post-earthquake visual inspection. Bull Earthq Eng
1312 15:1681–1706. https://doi.org/10.1007/s10518-016-0039-9
1313 Di Ludovico M, Prota A, Moroni C, Manfredi G, Dolce M (2017a) Reconstruction
1314 process of damaged residential buildings outside historical centres after the
1315 L'Aquila earthquake: part I—"light damage" reconstruction. Bull Earthq Eng
1316 15:667–692. https://doi.org/10.1007/s10518-016-9877-8
1317 Di Ludovico M, Prota A, Moroni C, Manfredi G, Dolce M (2017b) Reconstruction
1318 process of damaged residential buildings outside historical centres after the
1319 L'Aquila earthquake: part II—"heavy damage" reconstruction. Bull Earthq
1320 Eng 15:693–729. https://doi.org/10.1007/s10518-016-9979-3
1321 Di Michele F, Stagnini E, Pera D, Rubino B, Aloisio R, Askan A, Marcati P (2023)
1322 Comparison of machine learning tools for damage classification: the case of
1323 L'Aquila 2009 earthquake. Nat Hazards 116:3521–3546.
1324 https://doi.org/10.1007/s11069-023-05822-4
1325 Doğan A, Başeğmez M, Aydın CC (2025) Assessment of the seismic vulnerability in
1326 an urban area with the integration of machine learning methods and GIS. Nat
1327 Hazards 121:9613–9652. https://doi.org/10.1007/s11069-025-07185-4
1328 Epstein ES (1969) A scoring system for probability forecasts of ranked categories. J
1329 Appl Meteorol 8:985–987. https://doi.org/10.1175/1520-
1330 0450(1969)008<0985:ASSFPF>2.0.CO;2

FEMA (2018) Seismic performance assessment of buildings, volume 1—methodology, 2nd edn. FEMA P-58-1. Federal Emergency Management Agency, Washington, DC. https://www.fema.gov/sites/default/files/documents/fema_p-58-1-se_volume1_methodology.pdf. Accessed 15 August 2026

Gehl P, Fayjaloun R, Sun L, Tubaldi E, Negulescu C, Özer E, D'Ayala D (2022) Rapid earthquake loss updating of spatially distributed systems via sampling-based Bayesian inference. Bull Earthq Eng 20:3995–4023. https://doi.org/10.1007/s10518-022-01349-4

Ghimire S, Guéguen P (2024) Host-to-target region testing of machine learning models for seismic damage prediction in buildings. Nat Hazards 120:4563–4579. https://doi.org/10.1007/s11069-023-06394-z

Giordano PF, Iacovino C, Quqa S, Limongelli MP (2022) The value of seismic structural health monitoring for post-earthquake building evacuation. Bull Earthq Eng 20:4367–4393. https://doi.org/10.1007/s10518-022-01375-2

Gneiting T, Raftery AE (2007) Strictly proper scoring rules, prediction, and estimation. J Am Stat Assoc 102:359–378. https://doi.org/10.1198/016214506000001437

Graf E, Schmoor C, Sauerbrei W, Schumacher M (1999) Assessment and comparison of prognostic classification schemes for survival data. Stat Med 18:2529–2545. https://doi.org/10.1002/(SICI)1097-0258(19990915/30)18:17/18<2529::AID-SIM274>3.0.CO;2-5

Haindl M, Smith IFC, Beyer K (2024) Using image-based inspection data to improve response predictions of earthquake-damaged unreinforced masonry buildings. Bull Earthq Eng 22:7117–7148. https://doi.org/10.1007/s10518-024-02023-7

Ke G, Meng Q, Finley T, Wang T, Chen W, Ma W, Ye Q, Liu TY (2017) LightGBM: a highly efficient gradient boosting decision tree. In: Guyon I, von Luxburg U, Bengio S, Wallach H, Fergus R, Vishwanathan S, Garnett R (eds) Advances in

neural information processing systems 30. Curran Associates, Red Hook, pp 3146–3154

Khakurel S, Dhakal RP, Yeow TZ, Saha S (2021) Residential building repair cost and claim settlement time from the Canterbury Earthquake Sequence. In: Proceedings of the 2021 New Zealand Society for Earthquake Engineering annual conference, Christchurch, New Zealand, 14–16 April 2021, paper 157, 10 pp

Li L, Chang-Richards A, Boston M, Elwood K (2025) A hybrid data-driven approach for modeling post-earthquake housing recovery. Earthq Spectra 41:3045–3073. https://doi.org/10.1177/87552930251345841

Lundberg SM, Lee SI (2017) A unified approach to interpreting model predictions. In: Guyon I, von Luxburg U, Bengio S, Wallach H, Fergus R, Vishwanathan S, Garnett R (eds) Advances in neural information processing systems 30. Curran Associates, Red Hook, pp 4765–4774

Miles SB, Chang SE (2006) Modeling community recovery from earthquakes. Earthq Spectra 22:439–458. https://doi.org/10.1193/1.2192847

Molina Hutt C, Vahanvaty T, Kourehpaz P (2022) An analytical framework to assess earthquake-induced downtime and model recovery of buildings. Earthq Spectra 38:1283–1320. https://doi.org/10.1177/87552930211060856

Negulescu C, Benaïchouche A, Lemoine A, Le Roy S, Pedreros R (2020) Adjustability of exposed elements by updating their capacity for resistance after a damaging event: application to an earthquake–tsunami cascade scenario. Nat Hazards 104:753–793. https://doi.org/10.1007/s11069-020-04189-0

Sattar S, Ryan K, Arendt L, Bonowitz D, Comerio M, Davis C, Deierlein G, Johnson KJ (2021) Recommended options for improving the built environment for post-earthquake reoccupancy and functional recovery time. FEMA P-2090/NIST SP 1254. Federal Emergency Management Agency, Washington,

1387 DC, and National Institute of Standards and Technology, Gaithersburg, MD.
1388 https://doi.org/10.6028/NIST.SP.1254
1389 Singer JD, Willett JB (1993) It's about time: using discrete-time survival analysis to
1390 study duration and the timing of events. J Educ Stat 18:155–195.
1391 https://doi.org/10.3102/10769986018002155
1392 Tomar A, Burton HV, Mosleh A (2022) Dynamic updating of post-earthquake
1393 damage and functional restoration forecasts of water distribution systems
1394 using Bayesian inferencing. Earthq Spectra 38:109–127.
1395 https://doi.org/10.1177/87552930211038016
1396 Xiong C, Huang J, Lu X (2020) Framework for city-scale building seismic resilience
1397 simulation and repair scheduling with labor constraints driven by time–history
1398 analysis. Comput Aided Civ Infrastruct Eng 35:322–341.
1399 https://doi.org/10.1111/mice.12496

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- TableA1.docx