AI in Civil Engineering

ORIGINAL ARTICLE

Open Access

Machine learning surrogates for predicting and interpreting flood fragility parameters of masonry walls

Stefanos Georgiadis[1*], Marco Nale[1], Marco Faggioli[1], Željana Nikolić[2] and Elena Benvenuti[1*]

Abstract

Flood-induced out-of-plane collapse of masonry walls is a recurrent cause of severe building damage and loss during flood events, particularly in low-rise residential construction. Assessing this failure mode at scale remains challenging: simplified fragility models are fast but often lack physical transparency, while mechanics-based analyses are interpretable but computationally expensive. This study develops a mechanics-informed surrogate modelling framework that bridges this gap. Machine-learning surrogates, including linear sparse regressions, kernel-based methods, and ensemble models, are trained on a constrained synthetic dataset generated from a mechanics-based wall collapse formulation. This enables fragility parameters to be predicted with accuracy comparable to the reference simulator while retaining a direct link to the governing structural mechanisms within the adopted modelling domain. A central contribution of the work is the explicit integration of explainable artificial intelligence, allowing the role of geometry, openings, vertical connectivity, and boundary restraint to be quantified and interpreted in physically meaningful terms. The results reveal stable and intuitive collapse drivers across boundary configurations, with the upper-floor contribution, thickness-related descriptors, and boundary restraint forming the dominant sensitivity group, while masonry density and opening-related variables have secondary or model-dependent influence. The consistency of these trends across linear and nonlinear surrogate models demonstrates that the framework captures not only fragility estimates but also mechanically consistent sensitivity patterns of wall overturning. By combining mechanics-based simulations with surrogate modelling, the proposed approach enables efficient and physically interpretable assessment of masonry wall fragility.

Keywords Mechanics-informed surrogate modelling, Machine learning, Flood fragility analysis, Masonry walls

1 Introduction

Flooding is among the most frequent and damaging natural hazards worldwide, with impacts exacerbated by climate change, urbanisation, and the ageing of the building stock (Wing et al., 2020; Wang et al., 2024; Yarveysi et al., 2025). While advances in hydrodynamic modelling now allow increasingly detailed representation of flood hazards, the structural vulnerability of buildings remains a dominant source of uncertainty in flood-risk assessment (Huizinga et al., 2017; Nasiri et al., 2019; Russo et al., 2020). In particular, unreinforced masonry (URM) buildings are widely exposed and exhibit pronounced vulnerability to out-of-plane wall failure under flood loading, motivating the development of reliable and physically meaningful fragility models. Large-scale flood loss assessments are mainly based on empirical depth–damage relationships derived from post-event observations or insurance records (Huizinga et al., 2017;

*Correspondence:
Stefanos Georgiadis
stefanos.georgiadis@unife.it
Elena Benvenuti
elena.benvenuti@unife.it
[1] Department of Engineering, University of Ferrara, Ferrara, Italy
[2] Faculty of Civil Engineering, Architecture and Geodesy, University of Split, Split, Croatia

© The Author(s) 2026. **Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit http://creativecommons.org/licenses/by/4.0/.

FEMA factsheet, 2022). These models offer simplicity and scalability but remain specific to each region and are also weakly sensitive to structural configuration, boundary conditions, and failure mechanisms. Indicator-based approaches attempt to extend empirical models by incorporating multiple vulnerability drivers, such as construction quality or building age (Nasiri et al., 2019; Nur and Shrestha, 2017). However, indicator weighting is often subjective and rarely validated against physical collapse or structural response. As a result, both empirical and indicator-based models provide limited insight into the mechanisms governing flood-induced structural failure. Mechanics-based models represent the balance between destabilising actions caused by flooding and stabilising structural strength, allowing for direct interpretation of collapse mechanisms. Experimental investigations have demonstrated the critical role of hydrostatic pressure distribution, boundary constraints, and wall geometry in the instability of masonry walls under flood loading (Jansen et al., 2020). Building on these observations, simplified analytical frameworks have been proposed to estimate collapse conditions and derive fragility functions for masonry walls subjected to flood actions (Milanesi et al., 2018; Korswagen et al., 2025). Probabilistic extensions of these formulations have further incorporated material and geometric uncertainty, improving realism while preserving physical transparency (Jalayer et al., 2016; Nale et al., 2023; Bodda et al., 2021). Despite their interpretability and mechanical basis, repeated simulations of mechanics-based fragility models remain computationally demanding, particularly within Monte Carlo simulations or regional-scale vulnerability assessments (Dong and Frangopol, 2017). This computational load limits their direct applicability in large-scale flood-risk analyses and motivates the use of surrogate modelling approaches. Surrogate models are simplified approximations of more complex, higher-order models. They are used to map input data to outputs when the actual relationship between the two is unknown or computationally expensive to evaluate (Williams and Cremaschi, 2019). Machine-learning (ML) surrogate models provide an alternative solution for emulating mechanics-based simulators, enabling faster evaluation of nonlinear structural behavior after training. Surrogate modelling has been widely applied in hazard and fragility analysis using approaches such as kernel methods, ensemble learning, and sparse regression (Alizadeh et al., 2020; Garzón et al., 2022; Tsokanas et al., 2022; Kim and Boukouvala, 2020; Hwang and Martins, 2018). In the field of flood risk, data-driven and hybrid approaches have been proposed to estimate fragility parameters and determine the uncertainty of parameters that measure vulnerability or risk (Nofal and van de Lindt, 2020; Nofal et al., 2020; Refadah, 2025; Yarveysi et al., 2025; Shrestha and Chini, 2025). However, many ML-based fragility models face physical improbability problems. Without constraints from known laws, models may produce results that violate conservation principles or other scientific truths. They also struggle with interpretability. The lack of mechanistic grounding makes it difficult to explain or trust model behavior, especially in domains subject to high-risk decisions (Holzinger et al., 2026), such as the vulnerability of buildings to flooding. Recent advances in explainable artificial intelligence (XAI) have highlighted the importance of combining predictive accuracy with transparent and physically interpretable model behavior (Wang et al., 2024; Besarra et al., 2025; Bhattarai et al., 2025). Within flood-induced masonry fragility analysis, the systematic integration of XAI concepts with mechanics-based surrogate modelling remains limited. Recent studies have further advanced mechanics-based and surrogate-assisted vulnerability assessment for masonry structures under flood-related actions. Capparelli et al. (2023) proposed a rapid numerical-based vulnerability assessment method for masonry buildings subjected to flood actions, while (Asad et al., 2024) investigated out-of-plane masonry-wall failure under hydrostatic and hydrodynamic flash-flood loading. (Mendoza Cabanzo et al., 2022) used surrogate modelling techniques to derive flood fragility curves for masonry arch bridges exposed to local scour, highlighting the growing role of computationally efficient surrogate approaches in masonry flood-risk analysis. Compared with these studies, the present work focuses on predicting and interpreting probabilistic flood fragility parameters of masonry walls using ML surrogates trained on mechanics-based synthetic data, with XAI tools used to verify the consistency of the learned trends with structural mechanics.

1.1 Knowledge gap and contributions

The review above highlights a gap between physically transparent but computationally intensive mechanics-based fragility models and efficient but often opaque ML surrogates. While analytical formulations provide direct insight into collapse mechanisms (Milanesi et al., 2018; Jansen et al., 2020; Jalayer et al., 2016), data-driven approaches increasingly offer scalability at the expense of interpretability (Nofal and van de Lindt, 2020; Nofal et al., 2020; Rocchi et al., 2022; Tsokanas et al., 2022). This motivates the development of modelling approaches that retain the accuracy of mechanics-based simulations while allowing efficient prediction and physically meaningful interpretation of flood-induced wall collapse (Karniadakis et al., 2021; Willard et al., 2022).

To address this gap, this study develops a mechanics-informed surrogate modelling framework for predicting the fragility parameters—median μ and dispersion β—of masonry walls exposed to flood loading. In this study, the term *mechanics-informed* denotes that the ML surrogates are trained and interpreted using reference outputs generated by a mechanics-based limit-analysis simulator. It does not imply a physics-informed neural-network formulation in which residuals of governing differential equations are directly embedded in the loss function or network constraints (Raissi et al., 2019; Cuomo et al., 2022). Recent structural-mechanics examples explicitly encode geometric, constitutive, and equilibrium relations through loss terms derived from the governing differential equations (He et al., 2024). In contrast, the mechanical information in the present framework enters through the simulator-generated training database, mechanism-based input features, and the subsequent interpretation of surrogate responses using sensitivity and partial-dependence analyses. The resulting synthetic dataset is used to train surrogate models, including Support Vector Regression, Random Forests, and sparse linear regressions. A central objective of the study is the application of XAI techniques, including sensitivity analysis and partial dependence diagnostics, to ensure that surrogate predictions remain consistent with structural mechanics.

The main contributions of this work are threefold. First, a mechanics-informed surrogate modelling framework is developed to predict the flood fragility parameters of masonry walls across multiple boundary conditions, combining simulator-level accuracy with computational efficiency. Second, linear and nonlinear surrogate models are systematically compared to assess predictive performance and physical consistency under different geometric and mechanical configurations. Third, XAI techniques are used to ensure that surrogate predictions remain interpretable and aligned with the mechanics of out-of-plane wall instability. Building on these results, compact guideline representations are derived to support engineering interpretation and preliminary vulnerability assessment.

1.2 Paper organisation

The remainder of the paper is organised as follows. Section 2 presents the methodology; Sect. 3 reports the main results and interpretability analyses; Sect. 4 discusses implications and limitations; and Sect. 5 summarises the key findings and outlines future research directions.

2 Methodology

Figure 1 provides an overview of the proposed mechanics-informed surrogate modelling framework. A mechanics-based collapse model is first used to generate fragility parameters for masonry walls under flood loading, which are used as reference simulator outputs, across a broad range of geometric, material, opening, and boundary

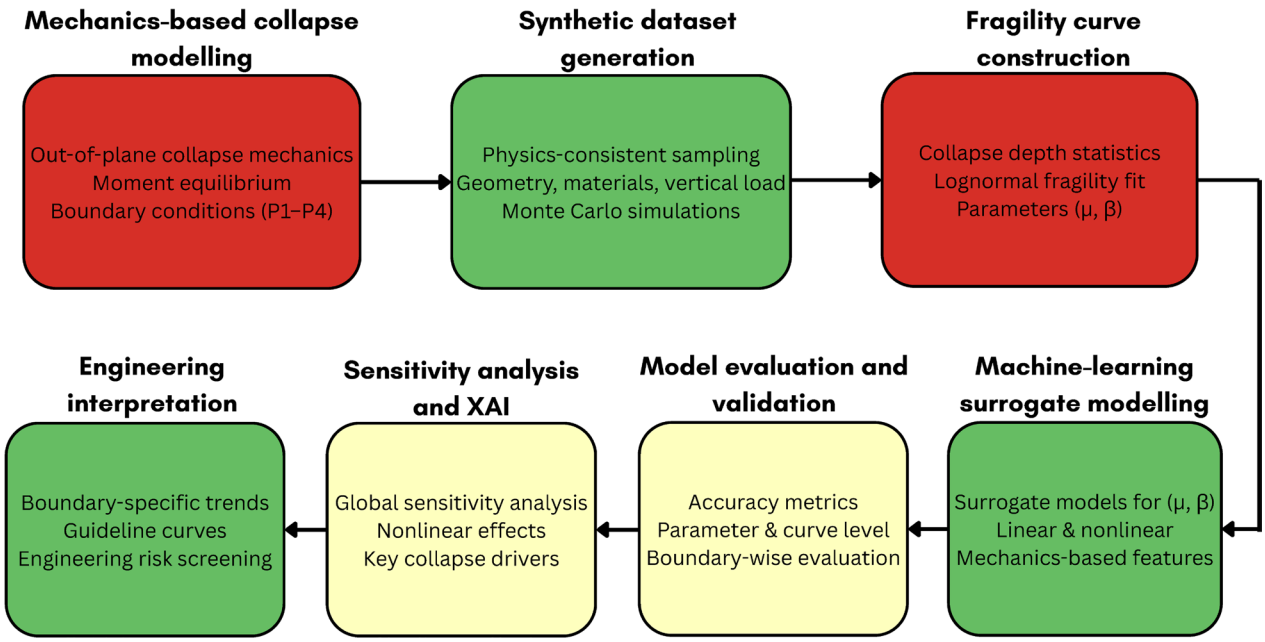


Fig. 1 Overview of the proposed mechanics-informed surrogate modelling framework. A mechanics-based collapse model is used to generate synthetic fragility data, which are approximated by ML surrogates. Model evaluation and sensitivity analysis are subsequently employed to interpret the learned sensitivity trends and derive compact guideline representations for preliminary engineering risk assessment

conditions. Next, these simulator outputs are estimated by ML surrogates, which are evaluated through cross-validation and curve-level diagnostics. Finally, complementary interpretability tools are applied to verify physical consistency and to derive compact, boundary-specific surrogate relations suitable for engineering applications.

2.1 Mechanics-based simulator

A simplified limit-analysis model originally proposed by (Milanesi et al., 2018) is adopted to simulate the out-of-plane collapse of masonry walls subjected to flood loading. The wall is idealised as an assembly of rigid blocks separated by planar cracks, and collapse is identified through the *principle of virtual work*, following the classical limit-analysis framework for masonry structures (Heyman, 1999). Within this framework, the masonry is assumed to exhibit unlimited compressive strength and stiffness, no tensile resistance, and no relative slip between adjacent blocks, and is treated as homogeneous at the panel scale. Scour effects and foundation erosion are neglected, and the wall is assumed to remain fully supported at the base throughout the loading process. The model offers a balance between analytical tractability and physical realism and has been widely used in flood fragility analyses (Jalayer et al., 2016; Nofal and van de Lindt, 2020; Jansen et al., 2020), enabling systematic evaluations without computationally expensive finite-element simulations (Alizadeh et al., 2020; Tsokanas et al., 2022; Besarra et al., 2025; Shrestha and Chini, 2025).

A wall of height Z, thickness t, and effective width $\ell_e = \ell(1-o)$, where o denotes the opening ratio, is exposed to floodwater of density ρ_f. A schematic representation of the wall geometry, coordinate system, and flood loading configuration is provided in Fig. 3. Defining $k = \rho_f g$ with $g = 9.81\ \mathrm{m\,s^{-2}}$, the hydrostatic pressure distribution associated with a local flood depth $h \le Z$ is

$$p(z) = k(h - z), \qquad 0 \le z \le h. \tag{1}$$

Following (Milanesi et al., 2018), the adopted hydrostatic representation is intended to capture the peak flood action associated with the maximum run-up at the wall and is assumed to govern the initiation of overturning mechanisms.

Four idealised restraint configurations are considered, denoted P1–P4, representing increasing levels of edge support and defining the corresponding kinematic collapse mechanisms. A schematic overview of the geometry and boundary conditions is provided in Figs. 2 and 3. Mechanisms P3 and P4 include two lateral wedges defined by an inclination angle α, while P4 additionally includes a horizontal hinge at height z_h separating the wall into lower and upper central blocks. The wedge angle α is selected from tabulated values as a function of the dimensionless geometry and load ratios

$$X_1 = \frac{\ell_e}{Z}, \qquad X_3 = \frac{n}{\rho_m g Z t}, \tag{2}$$

where ρ_m is the masonry density and n is the vertical line load per unit out-of-plane length. In the present implementation, intermediate values of X_1 and, where required, X_3 are obtained by linear interpolation.

Virtual-work collapse condition. Collapse is ruled by a balance between destabilising work due to hydrostatic pressure and stabilising work due to the wall self-weight and vertical loads. In general form,

$$L^-(h) = L^+, \tag{3}$$

where $L^-(h)$ is the external work of hydrostatic pressure up to depth h and L^+ denotes the stabilising work for the assumed kinematic mechanism. Since internal dissipation along fracture lines is neglected, collapse begins when

$$L^-(h) \ge L^+. \tag{4}$$

Both $L^-(h)$ and L^+ depend on the selected mechanism (P1–P4) through the associated rigid-body rotations and virtual displacement fields.

Stabilising actions. The stabilising actions comprise (i) the self-weight of the wall and (ii) vertical loads associated with upper floors. The wall self-weight is

$$W = \rho_m g\, t\, \ell_e Z, \tag{5}$$

and the vertical line load is given by n (units N/m). In the synthetic dataset generation, the total vertical load is parameterised as

$$n = n_{\mathrm{up}}\, \rho_m g\, t\, Z, \tag{6}$$

where n_{up} is an integer representing the number of upper floors contributing to vertical load. With this definition, the dimensionless load ratio in Eq. (2) reduces to $X_3 = n_{\mathrm{up}}$.

For a generic kinematic mechanism, the stabilising work is obtained as the sum of the work contributions of the individual rigid components,

$$L^+ = \sum_i W_i\, \xi_i + \sum_j N_j\, \xi_{N_j}, \tag{7}$$

where W_i and N_j denote the weights of the individual blocks and the vertical loads acting on them, respectively, and ξ_i and ξ_{N_j} are the corresponding virtual vertical displacements dictated by the assumed mechanism. For instance, for P1 (one-edge restraint), the mechanism

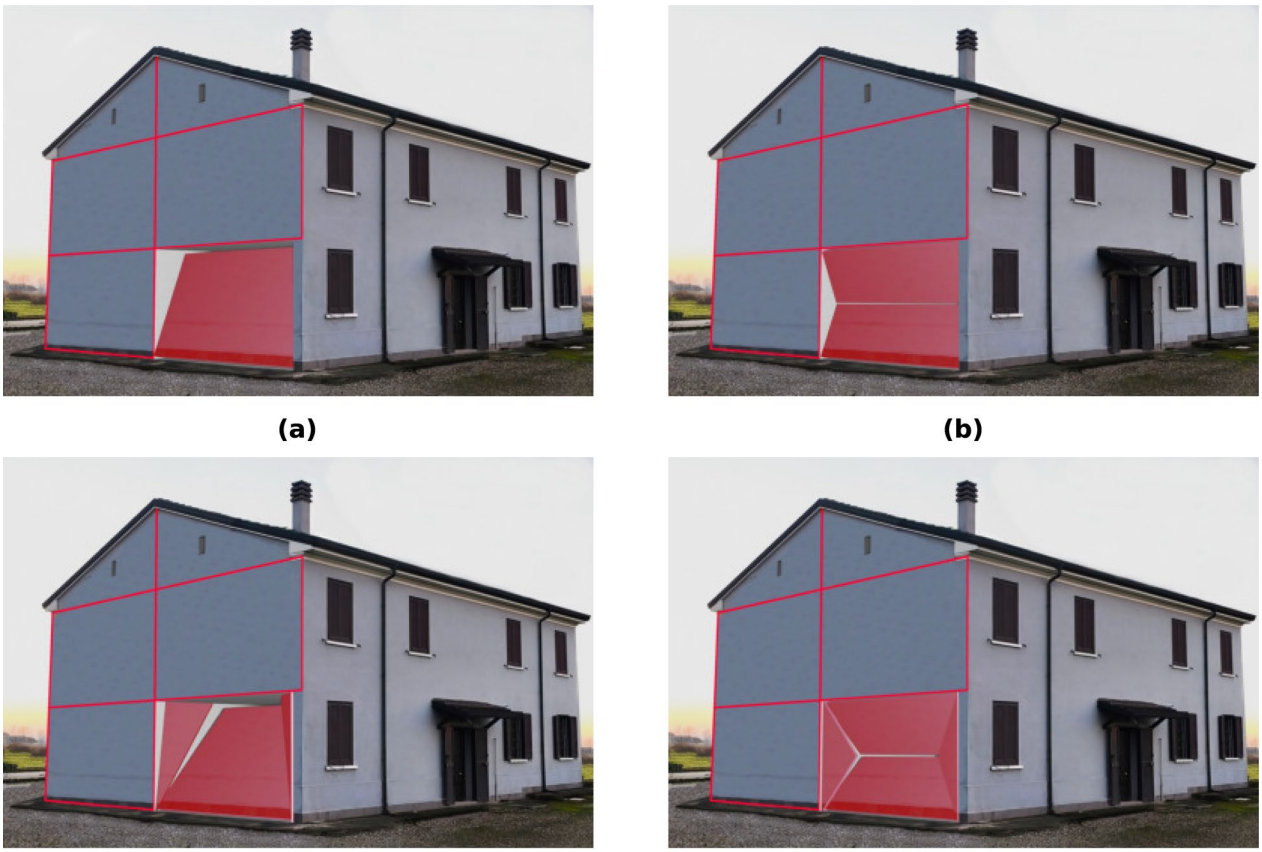


Fig. 2 Schematic representation of the masonry wall and the four idealised boundary restraint conditions considered in the limit-analysis model. Panels (**a**–**d**) correspond to boundary conditions P1–P4, respectively. P1 (one-edge restraint), P2 (two-edge restraint), P3 (three-edge restraint), and P4 (four-edge restraint)

reduces to a single rigid rotation about the base, leading to the compact expression

$$L_{P1}^{+} = (W + n\ell_e)\,\frac{t}{2}, \tag{8}$$

where $n\ell_e$ is the resultant vertical load acting over the effective wall length. For P2, P3, and P4, the calculation of the stabilizing work is more complex due to the geometry of these cases, applying Eq. (7) to the corresponding multi-block mechanisms, following the kinematic descriptions in (Milanesi et al., 2018).

Destabilising flood action. The destabilising work is computed by integrating the hydrostatic pressure distribution multiplied by the mechanism-dependent width-integrated virtual out-of-plane displacement field $\delta(z)$ along the wetted wall height,

$$L^{-}(h) = \int_0^h p(z)\,\delta(z)\,\mathrm{d}z, \tag{9}$$

where $\delta(z) = \int_0^{l_e} u(s,z)ds$ is the width-integrated virtual displacement (m^2), with s the horizontal coordinate and u(s,z) the local virtual out-of-plane displacement for the selected mechanism. For P1, this reduces to a closed-form expression equivalent to the work of a triangular hydrostatic load acting through its resultant lever arm. For P3, the integration is performed on the central region and the two lateral wedges; if the wedge intersection occurs within the wall height, the integration is evaluated in parts on two vertical zones separated at $z_i = \ell_e/(2\tan\alpha)$ (truncated to h and Z), consistent with (Milanesi et al., 2018). For P4, the destabilising work includes contributions from the lower and upper central blocks and the two lateral wedges. Kinematic compatibility between wedge rotations is enforced through

$$\tan\beta_{\mathrm{k}} = \frac{z_h}{Z - z_h}\tan\alpha, \tag{10}$$

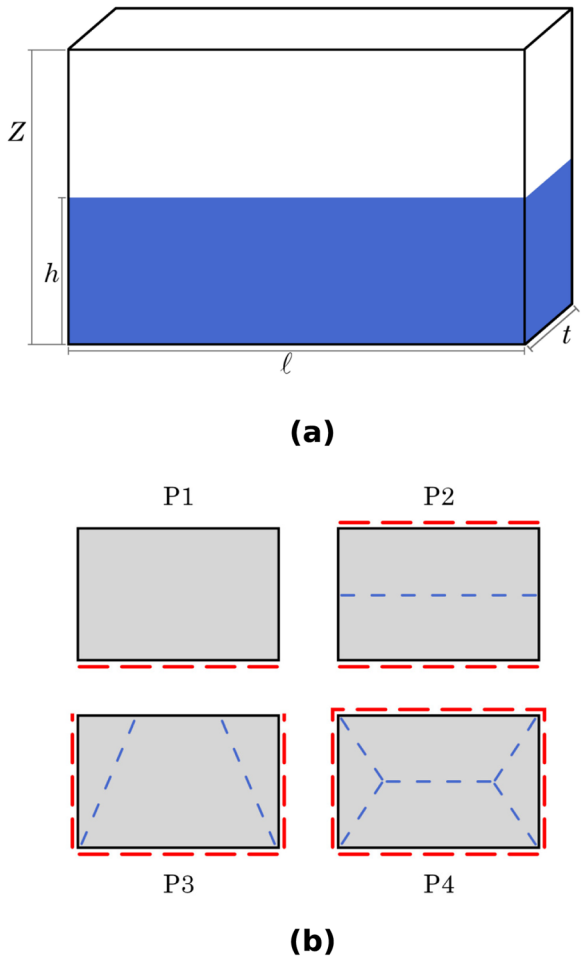


Fig. 3 Schematic representation of a masonry wall subjected to flood loading. (**a**) Three-dimensional view of the wall geometry and loading, showing the wall height Z, thickness t, length ℓ, and flood depth h. (**b**) Idealised boundary conditions (P1, P2, P3, P4) considered in the limit-analysis model: red dashed lines show restrained edges and blue dashed lines display the corresponding kinematic mechanisms

where β_k indicates the wedge inclination in the upper region.

For four-edge restraint (P4), the hinge location z_h is not known *a priori*. Here z_h denotes the hinge elevation, while h_c denotes the collapse flood depth. The admissible hinge positions are explored and the mechanism is identified as the configuration that minimises the safety margin (equivalently the ratio L^+/L^-) for a given flood depth.

For each wall configuration, the critical collapse depth is defined as the smallest flood depth satisfying Eq. (4) for the respective mechanism. The treatment of material and geometric variability and the estimation of fragility parameters from repeated simulations are described in Sect. 2.2.

Table 1 Input variables and constrained sampling ranges used to generate the synthetic dataset

Variable	Symbol	Unit	Sampling Range/ Levels
Wall thickness	t	m	0.20–0.50
Floor height	Z	m	2.4–4.0
Number of upper floors	n_{up}	–	0–2 (integer)
Masonry density	ρ_m	kg m^{-3}	1200–2200
Opening ratio	o	–	0.10–0.40
Wall length	ℓ	m	3.0–12.0
Boundary condition	BC	categorical	P1, P2, P3, P4

2.2 Synthetic dataset and input ranges

For the purposes of the current methodology, it was essential to create a synthetic dataset using the limit-analysis simulator described in Sect. 2.1. Each entry in the dataset represents a masonry wall defined by geometric, material, opening, and boundary parameters. The dataset was generated using independent random sampling of nominal wall configurations, followed by rejection of candidates that violated prescribed model-domain or typological constraints. Continuous variables were sampled within the bounds reported in Table 1, n_{up} was sampled as an integer variable over $\{0, 1, 2\}$, and the boundary condition was sampled from the four mechanisms P1–P4. For P1 configurations, $n_{up} = 0$ was imposed to satisfy the assumptions of the corresponding one-edge overturning mechanism.

This type of simulation-based sampling strategy is commonly adopted in surrogate-assisted reliability and fragility analyses, where a synthetic database generated from a reference mechanical or numerical model is used to train a computationally efficient emulator (Lieu et al., 2022; Banik and Hong, 2023). In the present study, the sampling strategy defines a broad and typologically plausible simulator-domain dataset for surrogate training, model comparison, and sensitivity interpretation. For each sampled wall configuration, the flood collapse depth h_c was calculated from the collapse condition in Eq. (4), and the resulting collapse-depth distribution was used to estimate the corresponding fragility parameters. This synthetic strategy enables systematic exploration of the mechanics-based model while remaining suitable for cases where observational flood-damage data are limited (Huizinga et al., 2017; Nofal and van de Lindt, 2020).

The independent variables were selected to reflect the main factors affecting out-of-plane collapse under flood loading: (i) wall thickness t, (ii) floor height Z, (iii) number of upper floors n_{up}, (iv) masonry density ρ_m, (v) opening ratio o, (vi) wall length ℓ, and (vii) boundary condition (P1–P4). These parameters are consistent with the assumptions of the analytical model of (Milanesi et al., 2018) and with previous flood vulnerability studies on masonry structures (Jalayer et al., 2016; Nofal et al., 2020). They also correspond to architectural and typological descriptors commonly used in masonry vulnerability and exposure-assessment frameworks (Zuccaro et al., 2023; Tosto et al., 2023; Colajanni et al., 2023). The resulting input vector for each realization is therefore

$$\mathbf{x} = \left[t,\, Z,\, n_{\text{up}},\, \rho_m,\, o,\, \ell,\, \text{BC}\right], \tag{11}$$

where BC denotes the categorical boundary condition.

The sampling ranges were selected to represent typical low-rise masonry walls in flood-prone European and Mediterranean contexts while ensuring sufficient variability for surrogate learning. Wall thickness $t = 0.20{-}0.50$ m, floor height $Z = 2.4{-}4.0$ m, and masonry density $\rho_m = 1200{-}2200\ \text{kg m}^{-3}$ cover common low-rise masonry-wall configurations from relatively lightweight masonry units to dense brick or stone masonry. The number of upper floors $n_{\text{up}} = 0{-}2$ represents the vertical load contribution in Eq. (6) while maintaining a conservative low-rise typological envelope. The opening ratio $o = 0.10{-}0.40$ spans lightly perforated to highly perforated wall elevations and is consistent with vulnerability and facade-level masonry studies where openings are treated as important structural descriptors (Del Zoppo et al., 2022; Liberatore et al., 2020; Tosto et al., 2023). The wall length range $\ell = 3{-}12$ m covers short wall panels, intermediate facade bays, and longer exterior wall portions. Thus, the adopted ranges define a mechanics-based modelling envelope rather than an inventory-calibrated frequency distribution for a specific regional building stock.

The constrained sampling rules were applied after the initial random draw. All accepted configurations were required to satisfy $0.05 \le t/Z \le 0.35$, consistent with the dimensionless thickness-ratio range of the adopted limit-analysis model. Additional gross aspect-ratio limits were imposed for restrained mechanisms, with $\ell/Z \le 4.0$ for P3 and $\ell/Z \le 2.75$ for P4. Further typological restrictions were used to avoid unlikely combinations of very thin walls, long spans, large openings, and multiple upper floors. In particular, configurations with $n_{\text{up}} = 2$ required $t \ge 0.24$ m, long walls required increased wall thickness, and highly perforated walls were not combined with the thinnest wall sections.

The constrained sampling procedure was verified through a plausibility audit, which confirmed that no generated nominal wall configuration violated the imposed model-domain or typological constraints. The audit considered both dimensional variables and the governing dimensionless quantities $X_{1,\text{eff}} = \ell(1-o)/Z$, $X_2 = t/Z$, $X_3 = n_{\text{up}}$, and $X_5 = \rho_f/\rho_m$. During sampling, the gross aspect ratio ℓ/Z was used to enforce the P3 and P4 rejection rules, whereas the mechanics-based simulator evaluates the effective aspect ratio $X_{1,\text{eff}}$ because openings reduce the effective wall length.

To account for material and geometric uncertainty around each nominal wall configuration, random perturbations were applied to selected parameters using normally distributed perturbations centred on the sampled nominal values and bounded within physically admissible ranges. This choice reflects the assumption that, for a given nominal wall case, construction tolerances, measurement uncertainty, and local material variability are most likely to remain close to the nominal value, while larger deviations are progressively less likely. Normal-type perturbations are commonly adopted in structural reliability and fragility analyses to propagate uncertainty in geometric and material properties through Monte Carlo-type simulations (JCSS, 2001; Silva et al., 2023; Jalayer et al., 2016).

Specifically, wall thickness, masonry density, and opening ratio were perturbed using standard deviations $\sigma_t = 0.05$ m, $\sigma_{\rho_m} = 50\ \text{kg m}^{-3}$, and $\sigma_o = 0.05$, respectively, while openings reduce the effective width ℓ_e multiplicatively. Values falling outside the admissible modelling range were clipped to prevent nonphysical realizations. For each configuration, the simulator incrementally increases the flood depth until the collapse condition in Eq. (4) is met, yielding a sample of collapse depths $\{h_c^{(i)}\}$.

A total of 10,000 nominal wall configurations were generated, providing a statistically rich dataset spanning the relevant combinations of geometry, material properties, openings, and boundary restraints. For each configuration, the distribution of collapse depths is summarized by the median μ and the logarithmic standard deviation β, which define the two fragility parameters used as learning targets for the surrogate models introduced in Sect. 2.3.

2.3 Model choice and learning setup

This subsection describes the surrogate models used in this study, the definition of the learning targets, the feature representations supplied to the models, and the training and validation pipeline. The goal is to determine the ML surrogate model that achieves the best possible fit of the mechanics-based simulator under all boundary

conditions, while ensuring that the model is capable of adapting to as wide a range of wall cases as possible.

2.3.1 Targets and problem formulation

For each wall configuration, the limit-analysis simulator returns a distribution of collapse flood depths h_c resulting from Monte Carlo perturbations of material and geometric properties. From this distribution, the fragility parameters are defined as

$$\mu = \mathrm{M}(\mathrm{h_c}), \qquad \beta = \sqrt{\ln\left(1 + \frac{\mathrm{Var}(\mathrm{h_c})}{\mathbb{E}(\mathrm{h_c})^2}\right)}, \tag{12}$$

where μ is the median collapse capacity, β the logarithmic dispersion of a lognormal fragility model, $\mathrm{Var}(\mathrm{h_c})$ the variance of the collapse depth distribution, and $\mathbb{E}(h_c)$ its expected value.

Surrogate learning is formulated as two supervised regression problems with shared inputs $\mathbf{x}$ and distinct outputs:

$$\hat{\mu} = f_\mu(\mathbf{x}), \qquad \hat{\beta} = f_\beta(\mathbf{x}). \tag{13}$$

To enforce positivity and improve numerical stability, the natural logarithm of β is used as the learning target, with back-transformation applied at inference. All models are trained using a mean squared error loss on standardized targets.

Predictive performance is also quantified using the coefficient of determination R^2 and the mean absolute error (MAE),

$$\begin{aligned} R^2 &= 1 - \frac{\sum_i (y_i - \hat{y}_i)^2}{\sum_i (y_i - \bar{y})^2}, \\ \mathrm{MAE} &= \frac{1}{N} \sum_i |y_i - \hat{y}_i|, \end{aligned} \tag{14}$$

where $y \in \{\mu, \beta\}$ as appropriate and $\bar{y}$ denotes the sample mean.

2.3.2 Curve-level evaluation of surrogate performance

In addition to parameter-level accuracy, surrogate models are evaluated on their ability to reproduce the full fragility curves implied by (μ, β). In the present study, the intensity measure is the flood water depth h, evaluated at the wall. For each configuration, the fragility function is expressed using the standard lognormal formulation,

$$F(h) = \Phi\left(\frac{\ln(h) - \ln(\mu)}{\beta}\right), \tag{15}$$

where $\Phi(\cdot)$ denotes the standard normal cumulative distribution function.

Predicted curves $F_{\mathrm{pred}}(h)$ are compared directly to the reference curves $F_{\mathrm{true}}(h)$ generated by the simulator. Two global curve-level error metrics are employed: the root-mean-square error,

$$\mathrm{RMSE} = \sqrt{\frac{1}{\mathrm{N}} \sum_{\mathrm{i}=1}^{\mathrm{N}} \left[\mathrm{F_p}(\mathrm{h_i}) - \mathrm{F_t}(\mathrm{h_i})\right]^2}, \tag{16}$$

and the normalised area between curves,

$$\mathrm{ABC} = \frac{1}{\Delta \mathrm{h}} \int \left|\mathrm{F_p}(\mathrm{h}) - \mathrm{F_t}(\mathrm{h})\right| \mathrm{dh}, \tag{17}$$

where $\Delta h = h_{\max} - h_{\min}$ is the evaluated water-depth range, $F_{\mathrm{p}}(h)$ and $F_{\mathrm{t}}(h)$ are the predicted and reference fragility curves, and the integral is evaluated numerically over the considered interval.

To link surrogate accuracy to engineering-relevant performance points, additional errors are computed for the flood depths associated with selected exceedance probabilities q,

$$h_q = \mu \exp\left(\beta\, \Phi^{-1}(q)\right), \tag{18}$$

with relative errors evaluated at $q = 0.10$, 0.50, and 0.90.

2.3.3 Feature engineering and representations

The base predictors correspond to the simulator inputs described in Sect. 2.2: wall thickness t, floor height Z, number of upper floors n_{up}, masonry density ρ_m, opening ratio o, wall length ℓ, and boundary condition (categorical). Continuous features are standardized to zero mean and unit variance, while boundary conditions (P1–P4) are encoded using one-hot indicators.

To integrate engineering knowledge into the learning process, a limited set of composite features is introduced. These include tZ as a slenderness-related measure, the quadratic wall-span descriptor $A = \ell^2$, a surrogate geometric scale measure, reflecting the quadratic scaling of wedge-controlled collapse mechanisms with wall span under increasing boundary restraint, and $o\ell$ as an approximation of the total opening width. For linear and sparse models, polynomial expansions up to second order are generated, including terms such as t^2, Z^2, and tZ, while higher-order monomials are excluded to preserve interpretability and avoid overfitting.

2.3.4 Surrogate models and training procedure

A representative set of regression models including linear, kernel-based, and ensemble families is adopted. Linear baselines include ridge and lasso regression, which

provide transparent parametric relations and enable coefficient regularization. Ridge regression minimizes

$$\mathcal{L}_{\text{ridge}} = \frac{1}{N}\sum_{i=1}^{N}\left(y_i - \hat{y}_i\right)^2 + \lambda\|\mathbf{w}\|_2^2, \tag{19}$$

while lasso regression replaces the ℓ_2 penalty with an ℓ_1 term to promote sparsity.

Support Vector Regression (SVR) with a Radial Basis Function (RBF) kernel is employed to capture smooth nonlinear dependencies. The SVR predictor takes the form

$$\hat{y}(\mathbf{x}) = \sum_{i=1}^{N_{\text{SV}}}\alpha_i \exp\left(-\gamma\|\mathbf{x} - \mathbf{x}_i\|^2\right) + b, \tag{20}$$

with hyperparameters optimized via grid search.

Ensemble tree-based models are represented by Random Forests and Gradient Boosting methods. Random Forest predictions are obtained as

$$\hat{y}_{\text{RF}}(\mathbf{x}) = \frac{1}{N_{\text{trees}}}\sum_{m=1}^{N_{\text{trees}}} T_m(\mathbf{x}), \tag{21}$$

while Gradient Boosting constructs an additive model,

$$F_m(\mathbf{x}) = F_{m-1}(\mathbf{x}) + \nu\, h_m(\mathbf{x}), \tag{22}$$

where $h_m(\mathbf{x})$ denotes the regression tree fitted at iteration m and ν is the learning rate.

All models are trained using stratified 5-fold cross-validation with respect to boundary condition to ensure consistent representation of P1–P4 configurations. Hyperparameters are optimized within the cross-validation loop, and final models are retrained on the full training set before evaluation on a 20% hold-out test set.

2.4 Explainable AI and sensitivity analysis framework

A key objective of this study is the interpretability of the results. To this end, a generalized sensitivity analysis framework is used to interpret the predictions of the surrogates and test their consistency with the physics mechanisms leading to wall overturning under flooding. The objective is to quantify and visualize how geometric, material, and boundary-condition variables influence the fragility parameters and to identify stable, physically meaningful trends across different surrogate families.

(1) Linear and sparse–polynomial sensitivities. For linear and sparse polynomial surrogates (OLS, Ridge, Lasso), the learned model can be expressed in the general form

$$\hat{y} = B + \sum_j w_j x_j + \sum_{j\le k} w_{jk} x_j x_k, \tag{23}$$

where $\hat{y} \in \{\mu, \log\beta\}$ and x_j denote the standardized features. The coefficients w_j and w_{jk} provide direct estimates of first-order and interaction effects in standardized units.

For discrete predictors such as the number of upper floors n_{up}, the surrogate response can be summarized using a low-order polynomial,

$$\hat{y}(n_{\text{up}}) = a_0 + a_1 n_{\text{up}} + a_2 n_{\text{up}}^2,$$

and the stepwise effect of adding one upper floor is evaluated through finite increments,

$$\Delta_k = \hat{y}(k) - \hat{y}(k-1).$$

These increments quantify the change in the predicted fragility parameter between successive upper-floor levels while keeping all other variables unchanged.

(2) Permutation-based global importance. Permutation importance is used as a model-agnostic measure of global sensitivity. For each feature, its values are randomly permuted while all others are held fixed, and the resulting degradation in predictive performance is recorded. Features causing the largest performance loss are interpreted as the most influential. This method enables direct comparison of importance rankings across surrogate families and boundary conditions.

(3) Partial dependence (PD). To characterize nonlinear marginal effects, partial dependence (PD) functions are computed for selected predictors,

$$\begin{aligned}\bar{y}(x_j) &= \mathbb{E}_{\mathbf{x}_{-j}}\left[\hat{y}(x_j, \mathbf{x}_{-j})\right] \\ &= \frac{1}{N}\sum_{k=1}^{N}\hat{y}\left(x_j, \mathbf{x}_{-j}^{(k)}\right),\end{aligned} \tag{24}$$

where $\mathbf{x}_{-j}$ are all features except x_j. PD curves reveal monotonic trends, curvature, and saturation effects that cannot be captured by linear coefficients.

3 Results

This section presents the predictive performance of the ML surrogates for the two fragility parameters—the median collapse depth μ and the lognormal dispersion β—followed by a sensitivity analysis aimed at interpreting the learned relationships. The focus is on quantifying how well the different regression algorithms (linear, kernel-based, and ensemble models) approximate the mechanics-based simulator, and on assessing whether the resulting surrogates behave consistently with the underlying mechanics of wall collapse under flood loading.

Table 2 Cross-validated performance of the main ML models for predicting the fragility median μ. Both models with updated mechanics-based features and feature-based models are evaluated

Model	R^2_{CV}	MAE_{CV} [m]
Ridge Regression (baseline)	0.936	0.125
Lasso Regression (baseline)	0.936	0.124
Gradient Boosting (GBR)	0.994	0.033
Random Forest (RF)	0.995	0.032
Hist. Grad. Boosting (HGB)	0.997	0.024
SVR (RBF kernel, base)	0.998	0.024
SVR (RBF kernel, mech)	0.998	0.023

Unless otherwise specified, all metrics are reported from stratified k-fold cross-validation with respect to boundary condition (P1–P4), and test scores refer to a 20 % hold-out set that was not used during training or hyperparameter tuning. The surrogate models, evaluation metrics, and sensitivity diagnostics discussed in this section are defined in Sects. 2.3–2.4.

3.1 Model performance and comparison

Table 2 summarises the cross-validated performance of the core surrogate families considered in this work for predicting the fragility median μ. Linear baseline models (Ridge and Lasso) already achieve reasonably high accuracy, with $R^2_{CV} \approx 0.94$ and $MAE_{CV} \approx 0.12$ m. These findings suggest that linear models, despite achieving reliable performance, cannot capture the full performance of the simulator. This confirms that nonlinear relationships dominate the simulator's calculations. Nevertheless, tree-based ensemble methods yield a marked improvement in predictive performance. Gradient Boosting, Random Forests, and Histogram Gradient Boosting all exceed $R^2_{CV} = 0.99$, with mean absolute errors reduced to the range 0.02–0.03 m. The SVR models provide similarly high accuracy, confirming that the median collapse-depth response can be emulated very accurately over the constrained sampling domain.

The SVR_{RBF} model augmented with mechanics-based features was retained as the final surrogate for μ because it combines near-optimal predictive accuracy with the same mechanics-informed feature representation used for interpretation. This model achieves $R^2_{CV} = 0.998$ and $MAE_{CV} = 0.023$ m, indicating an almost exact reproduction of the mechanics-based simulator, which was the desired outcome. These performance metrics quantify the ability of the surrogate models to reproduce the outputs of the adopted mechanics-based simulator. Therefore, the reported accuracy demonstrates simulator-level emulation accuracy, while the external validity of the predictions remains bounded by the assumptions of the reference analytical model. These assumptions and their implications are discussed further in Sect. 4.

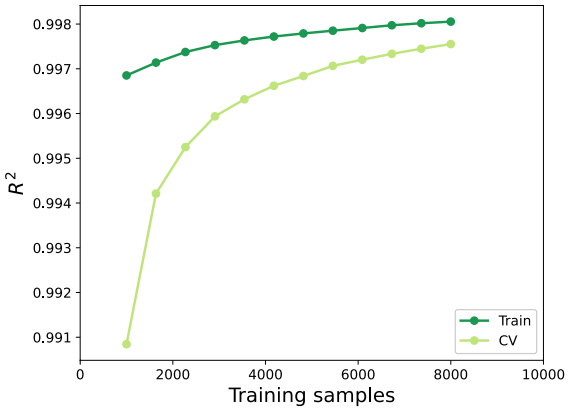


Fig. 4 Learning curve of the SVR model for prediction of the fragility median μ

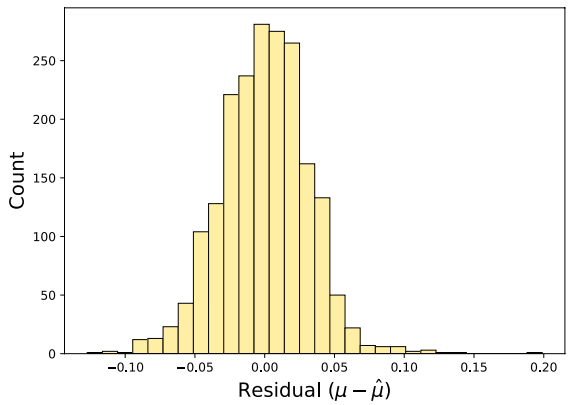


Fig. 5 Residual distribution of the SVR model for μ on the test set

Table 3 Cross-validated performance of the main models for predicting the fragility dispersion β

Model	R^2_{CV}	MAE_{CV}
Ridge Regression (baseline)	0.698	0.013
Lasso Regression (baseline)	0.698	0.013
Gradient Boosting (GBR)	0.920	0.007
Random Forest (RF)	0.945	0.006
Hist. Grad. Boosting (HGB)	0.943	0.007
SVR (RBF kernel, base)	0.702	0.015
SVR (RBF kernel, mech)	0.724	0.014

Figure 4 presents the learning curve of the best-performing SVR model. The cross-validation score increases

steadily and does not saturate, indicating that the surrogate remains data-limited rather than model-limited. The residual distribution in Fig. 5 shows an approximately Gaussian centered pattern, with slightly heavier tails for higher fragility medians, reflecting the heteroskedastic response of the ML surrogate model.

Table 3 reports the corresponding results for the fragility dispersion parameter β. In contrast to the median, prediction of β proves substantially more challenging, with systematically lower coefficients of determination across all models. Linear baselines achieve $R^2_{\mathrm{CV}} \approx 0.70$ and $\mathrm{MAE}_{\mathrm{CV}} \approx 0.013$, highlighting the limited ability of purely linear mappings to capture the complex interactions governing dispersion in the collapse response.

Ensemble methods provide a significant improvement for β. Gradient Boosting, Random Forests, and Histogram Gradient Boosting all reach $R^2_{\mathrm{CV}} \approx 0.92-0.95$, with $\mathrm{MAE}_{\mathrm{CV}}$ values close to $0.006-0.007$. Among these, Random Forests deliver the best overall performance in terms of mean absolute error, achieving $R^2_{\mathrm{CV}} = 0.945$ and $\mathrm{MAE}_{\mathrm{CV}} = 0.006$. The relatively lower performance of SVR for β, with R^2_{CV} values around $0.70-0.72$, implies that the dispersion parameter is influenced by different physical drivers than the fragility median μ.

The learning and residual curves for the selected β -surrogate exhibit similar qualitative behavior to those for μ: cross-validation scores improve with increasing sample size, and residuals remain centred with modest dispersion. However, the lower R^2 compared with the μ surrogate indicates that β carries additional variability that is not fully explained by the current input set. This is expected because β describes the spread of collapse depths generated by the Monte Carlo perturbations around each nominal wall configuration, whereas μ primarily represents the central collapse capacity. The current input features describe the nominal wall geometry, material density, opening ratio, and boundary condition, but they do not explicitly encode uncertainty descriptors such as coefficients of variation, correlation structure, local material heterogeneity, or boundary-restraint uncertainty. Therefore, part of the variability controlling β remains implicit in the simulator-generated response and is less directly observable from the current macroscopic input features.

Table 4 Curve-level and parameter-level evaluation metrics for the surrogate fragility model

Metric	Value	Interpretation
Mean curve RMSE	0.018	Excellent curve agreement
Std. curve RMSE	0.015	Low variability
Mean ABC	0.008	Very small curve mismatch
Std. ABC	0.005	Low variability
Rel. error $h_{0.10}$	0.017	~1.7% quantile error
Rel. error $h_{0.50}$	0.013	~1.3% quantile error
Rel. error $h_{0.90}$	0.015	~1.5% quantile error

The curve-level metrics summarised in Table 4 confirm that the surrogate models accurately reproduce the full fragility behavior of the numerical simulator. The mean curve RMSE is 0.018, with a corresponding mean area-between-curves (ABC) of 0.008, indicating that the predicted fragility curves deviate from the true curves by only about 1–2% across the intensity measure domain. This level of agreement is consistent across the test set, as reflected by the low standard deviations. Parameter-level errors corroborate this performance. Importantly, engineering-relevant performance points are reproduced with similar accuracy, with relative errors of 1.7, 1.3, and 1.5% for the 10th, 50th, and 90th percentile intensity levels, respectively. These results demonstrate that the combined SVR–RF surrogate framework provides a valid approximation of the simulator-generated fragility curves.

3.2 Feature importance and sensitivity analysis

The relative influence of geometric, material, and boundary-related variables on the fragility median μ is examined across surrogate models and boundary configurations (P1–P4). The input features correspond to the variables defined in Eq. (11) and summarised in Table 1. Because the surrogate is trained on the constrained simulator-domain dataset described in Sect. 2.2, the feature-importance results quantify sensitivity trends within this modelling envelope. This interpretation is consistent with model-agnostic interpretability studies showing that permutation feature importance and partial-dependence estimates depend on the data distribution, feature dependencies, and the sampled region of the input space (Molnar et al., 2023). Within this envelope, consistent importance patterns are observed across the model families, indicating stable rankings of the dominant predictors.

The overall feature-importance ranking indicates that the fragility median is primarily controlled by variables associated with vertical load contribution, wall sectional resistance, and boundary restraint. The number of upper floors n_{up} is the leading global predictor, followed by wall thickness t and the boundary condition. Floor height Z and masonry density ρ_m form a secondary group, while the quadratic wall-span descriptor and opening ratio have substantially smaller first-order contributions. This pattern is consistent with the mechanics-based formulation, in which collapse resistance depends on the balance between destabilising hydrostatic actions and stabilising

effects associated with wall section, self-weight, vertical load, and restraint.

The boundary-resolved feature-importance results provide further insight into these mechanisms. For the base-only configuration (P1), where $n_{\mathrm{up}} = 0$ by definition, the response is governed mainly by thickness-related and geometric descriptors, especially the coupled descriptor tZ, wall thickness t, masonry density ρ_m, and wall-scale variables such as ℓ or A. For the restrained configurations (P2–P4), n_{up} becomes the dominant predictor, followed by thickness-related descriptors t and tZ. Masonry density remains stabilising but secondary, while geometric variables such as Z, A, and ℓ have lower or boundary-dependent influence.

Opening-related variables display consistently low first-order importance within the explored parameter space. This reflects the effective-width representation adopted in the analytical model, where openings enter through $\ell_e = \ell(1 - o)$. The composite opening-length term $o\ell$ also remains generally secondary, although it can partly reflect wall-scale geometry in the P1 configuration. A more detailed interpretation is provided in Sect. 4.1.

Taken together, the feature-importance results support the following qualitative grouping of influence for the fragility median:

$$(n_{\mathrm{up}},\ t,\ \mathrm{BC}) \ >\ (\mathrm{tZ},\ \rho_{\mathrm{m}},\ \mathrm{Z}) \ >\ (\mathrm{A},\ \ell,\ \mathrm{o},\ o\ell). \quad (25)$$

This expression provides a compact summary of the dominant sensitivity groups within the adopted simulator-domain dataset. In the boundary-resolved analysis, the grouping changes slightly because the boundary condition is fixed within each subset: P1 is dominated

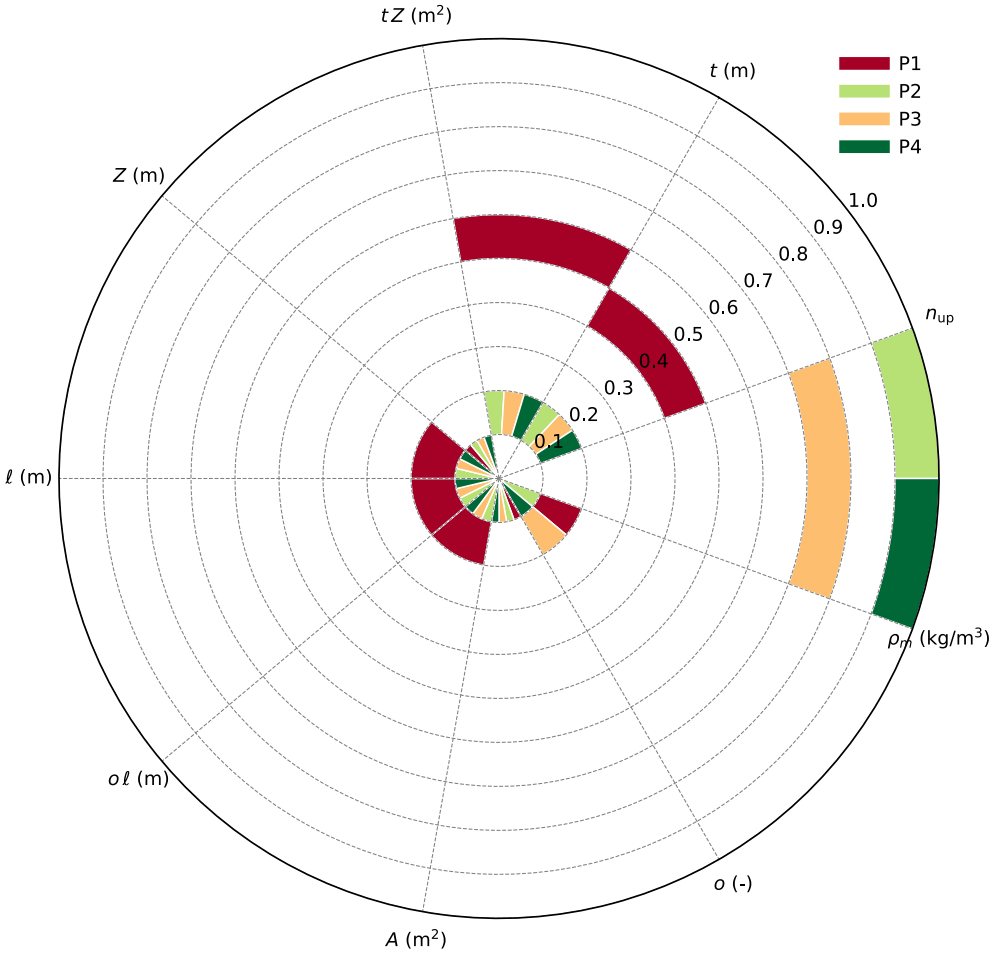


Fig. 6 Permutation-based feature importance for the fragility median μ, resolved by boundary condition. Feature importances are shown in a polar representation, with colours indicating the four boundary conditions P1–P4. Thickness-related descriptors and n_{up} for the restrained cases form the leading sensitivity group; boundary condition is important in the global ranking, while density and geometric descriptors are secondary and opening-related variables are generally low or boundary-dependent

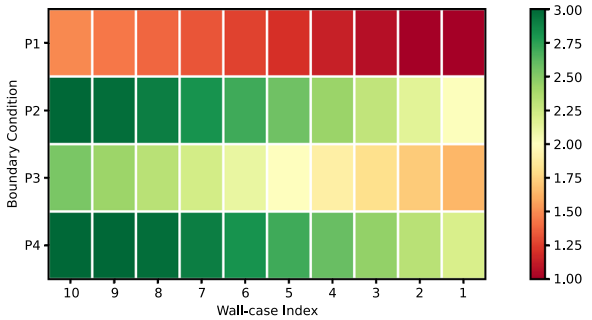


Fig. 7 Risk heatmap of the predicted median collapse depth μ as a function of boundary condition (P1–P4) and wall-case index. Each index corresponds to a predefined discrete triplet of opening ratio o, masonry density ρ_m, and wall thickness t, as listed in Table 5. P1 is evaluated with $n_{up} = 0$, while P2–P4 are evaluated with $n_{up} = 2$. Red colours indicate lower values of μ (higher collapse risk), while green colours correspond to increased collapse resistance. The ordering P1 < P3 < P2 < P4 is preserved across the representative wall cases

Table 5 Representative wall-case indices used in the risk map. Each index is a predefined discrete triplet of opening ratio o, masonry density ρ_m, and wall thickness t

Index	o	ρ_m (kg/m^3)	t (m)
1	0.29	1600	0.24
2	0.27	1620	0.26
3	0.25	1650	0.28
4	0.24	1675	0.30
5	0.23	1700	0.32
6	0.22	1750	0.34
7	0.22	1800	0.36
8	0.21	1850	0.38
9	0.21	1900	0.40
10	0.20	2000	0.42

by tZ, t, and wall-scale descriptors, whereas P2–P4 are dominated by n_{up}, t, and tZ. The findings are visualized in Fig. 6, which presents the permutation-based feature importance resolved by boundary condition.

To provide a synthesis view of how boundary restraint and wall characteristics jointly affect collapse capacity, Fig. 7 shows a risk heatmap of the median collapse flood depth μ as a function of boundary condition (P1–P4) and wall-case index. Each wall-case index corresponds to a predefined discrete triplet of opening ratio, masonry density, and wall thickness, as summarised in Table 5. The revised indices were selected so that all heatmap cells satisfy the constrained sampling rules. To isolate the effects of boundary condition and representative wall case, the remaining parameters are held constant across all heatmap cells ($Z = 3$ m, $\ell = 5$ m). The number of upper floors is fixed to $n_{up} = 2$ for P2–P4, while $n_{up} = 0$ for P1 by definition. For the restrained configurations with $n_{up} = 2$, the wall thickness satisfies $t \geq 0.24$ m. The heatmap is therefore interpreted as a scenario-based visualization of the combined influence of opening ratio, density, thickness, and boundary condition.

The heatmap reveals a consistent ordering of collapse capacity across boundary conditions, with configurations featuring increased restraint exhibiting higher values of the fragility median μ for all wall cases defined in Table 5. Increasing the wall-case index systematically shifts the collapse capacity toward larger flood depths, while the boundary condition controls the overall capacity level at which this shift occurs. In particular, the constrained heatmap preserves the ordering P1 < P3 < P2 < P4. Boundary conditions with top restraint (P2 and P4) form a higher-capacity group compared with those without top restraint (P1 and P3), highlighting the role of boundary restraint and upper-diaphragm action in the flood vulnerability of walls with comparable representative configurations.

3.3 Partial dependence analysis

Partial dependence results are used to examine the functional relationship between key predictors and the fragility median μ under different boundary conditions. Figures 8 and 9 illustrate representative PD profiles for the one-edge (P1) and four-edge (P4) configurations, respectively. The plotted predictors are n_{up}, wall thickness t, masonry density ρ_m, opening ratio o, wall length ℓ, and floor height Z. These profiles complement the feature-importance analysis by showing how the selected surrogate responds to controlled variations of individual inputs over the sampled modelling envelope.

Across the considered boundary conditions, wall thickness t exhibits a clear monotonic positive effect on μ. As shown in Figs. 8 and 9, increasing t consistently shifts the predicted collapse capacity toward larger flood depths. The response span is moderate for P1 and larger for P4, indicating that the influence of thickness becomes more pronounced under stronger boundary restraint. The same positive trend is observed in Fig. 10, where the response is expressed as the normalized median collapse depth μ/Z.

The influence of the number of upper floors n_{up} is positive and concave for the restrained configurations. Because n_{up} is an integer variable, the plotted response is interpreted through changes between successive floor levels. The PD curves for restrained cases, represented by the P4 response in Fig. 9, indicate that the increase in μ from zero to one upper floor is larger than the additional

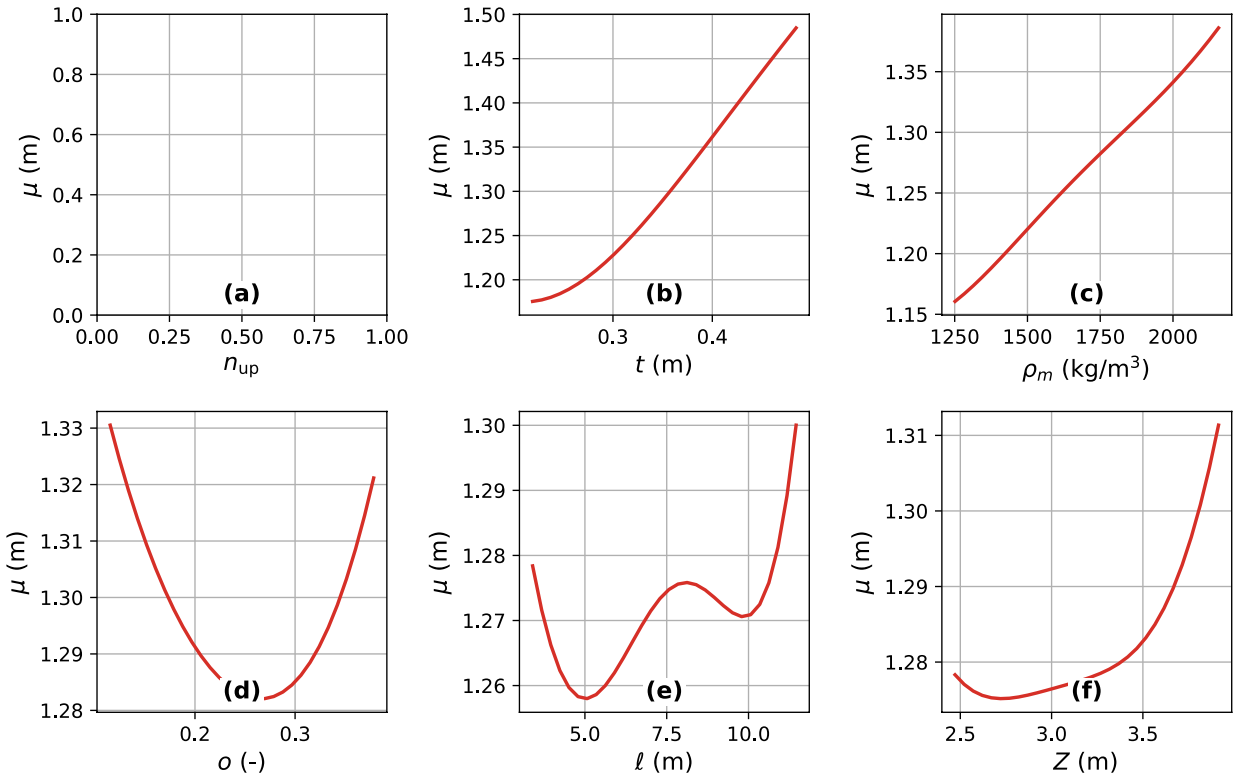


Fig. 8 Partial dependence responses for the fragility median under the P1 (one-edge) boundary condition

increase from one to two upper floors. This pattern indicates diminishing marginal gains from additional vertical load and restraint. For P1, $n_{\text{up}} = 0$ is imposed by definition, so upper-floor effects are not evaluated for this mechanism. The normalized responses in Fig. 10 show the same qualitative behaviour for P2–P4.

The masonry density ρ_m also has a positive influence on the fragility median μ, although its response span remains smaller than those of wall thickness and upper-floor contribution. The opening ratio o and wall length ℓ show lower-magnitude and boundary-dependent effects: under P1 their responses are mildly non-monotonic, whereas under P4 the opening-ratio response is small and positive and the wall-length response remains weak over the plotted range. The normalized responses in Fig. 10 lead to the same overall interpretation, with density remaining positive and opening-related and wall-scale geometric effects remaining secondary.

The floor height Z shows a weak or boundary-dependent response in the absolute-depth PD curves. Under P1, the response is nearly flat, while under P4 it is mildly non-monotonic. Fig. 10 provides the corresponding normalized response μ/Z, which decreases with increasing Z across the plotted boundary conditions. Thus, the absolute-depth and normalized-depth profiles provide complementary information: the former describes changes in predicted collapse depth, while the latter describes the corresponding depth-to-height trend within the adopted modelling domain.

4 Discussion

This section synthesizes the results presented in Sect. 3 and interprets them in light of the mechanics governing flood-induced out-of-plane collapse of masonry walls. The discussion focuses on (i) the mechanistic interpretation of the dominant sensitivity trends, (ii) the prediction of fragility dispersion and the role of uncertainty descriptors, (iii) the influence of boundary restraint and interaction effects, and (iv) the consistency of surrogate interpretations with the adopted limit-analysis formulation. The computational efficiency of the proposed framework, together with its validation scope, applicability boundaries, and future extensions, is also addressed.

4.1 Mechanistic interpretation of dominant sensitivity trends

The sensitivity results presented in Sect. 3 indicate that the fragility median μ is governed by a group of

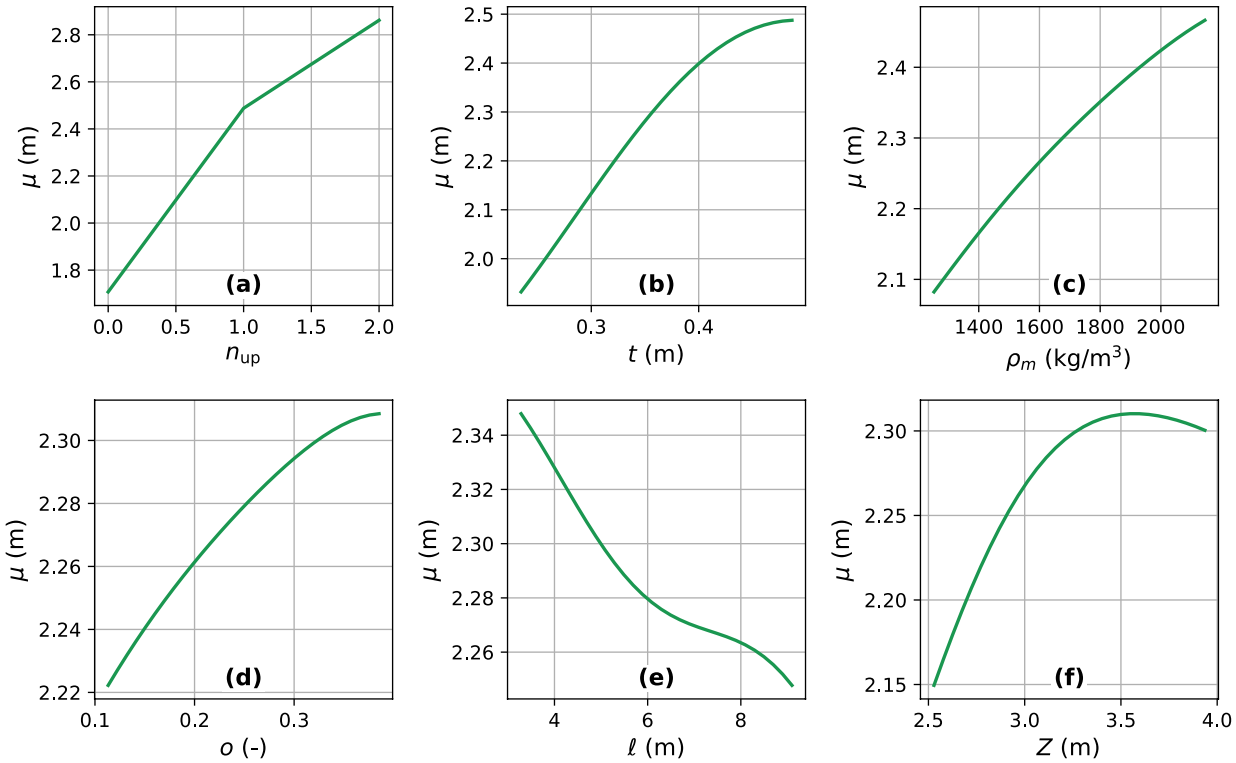


Fig. 9 Partial dependence responses for the fragility median under the P4 (four-edge) boundary condition

mechanically related variables associated with vertical load contribution, wall sectional resistance, and boundary restraint. In particular, n_{up}, wall thickness t, boundary condition, and the coupled thickness–height descriptor tZ emerge as the dominant sensitivity group, while masonry density ρ_m provides a secondary stabilising contribution. This finding is mechanically consistent with limit-analysis principles: increasing thickness enhances both sectional capacity and the stabilising moment resisting flood-induced overturning (Milani et al., 2007).

The dependence of the fragility median on wall thickness can be summarized by a compact low-order approximation,

$$\mu(t) \approx b_0 + b_1\, t + b_2\, t^2 \tag{26}$$

with

$$\frac{\partial \mu}{\partial t} = b_1 + 2b_2 t > 0 \quad \text{for } t \in [0.20, 0.50]\ \text{m}. \tag{27}$$

Eq(27) states the condition that the predicted collapse capacity increase with the thickness over the considered range. This interpretation is consistent with prior studies showing that boundary restraint limits rigid-body rotation and promotes geometry-dependent resisting mechanisms in out-of-plane wall response (Destro Bisol et al., 2024).

Furthermore, in the corresponding analysis with the upper floors as input feature, the relationship between μ and n_{up} shows a clear concave trend. Both feature-importance rankings and PD profiles show that the first attached floor produces the largest increase in μ, while additional floors yield diminishing marginal gains. This behaviour can be approximated by

$$\mu(n_{up}) \approx a_0 + a_1 n_{up} + a_2 n_{up}^2, \qquad a_1 > 0, \qquad a_2 < 0. \tag{28}$$

which captures the saturation of stabilising effects as vertical connectivity increases.

Since n_{up} is an integer variable, the marginal contribution of additional upper floors is more appropriately described through finite increments rather than through a continuous derivative. The incremental contribution can therefore be written as

$$\Delta_k = \mu(k) - \mu(k-1), \qquad k = 1, 2, \tag{29}$$

where the PD profiles indicate $\Delta_1 > \Delta_2 > 0$. This behaviour reflects the progressive restraint of wall rotation,

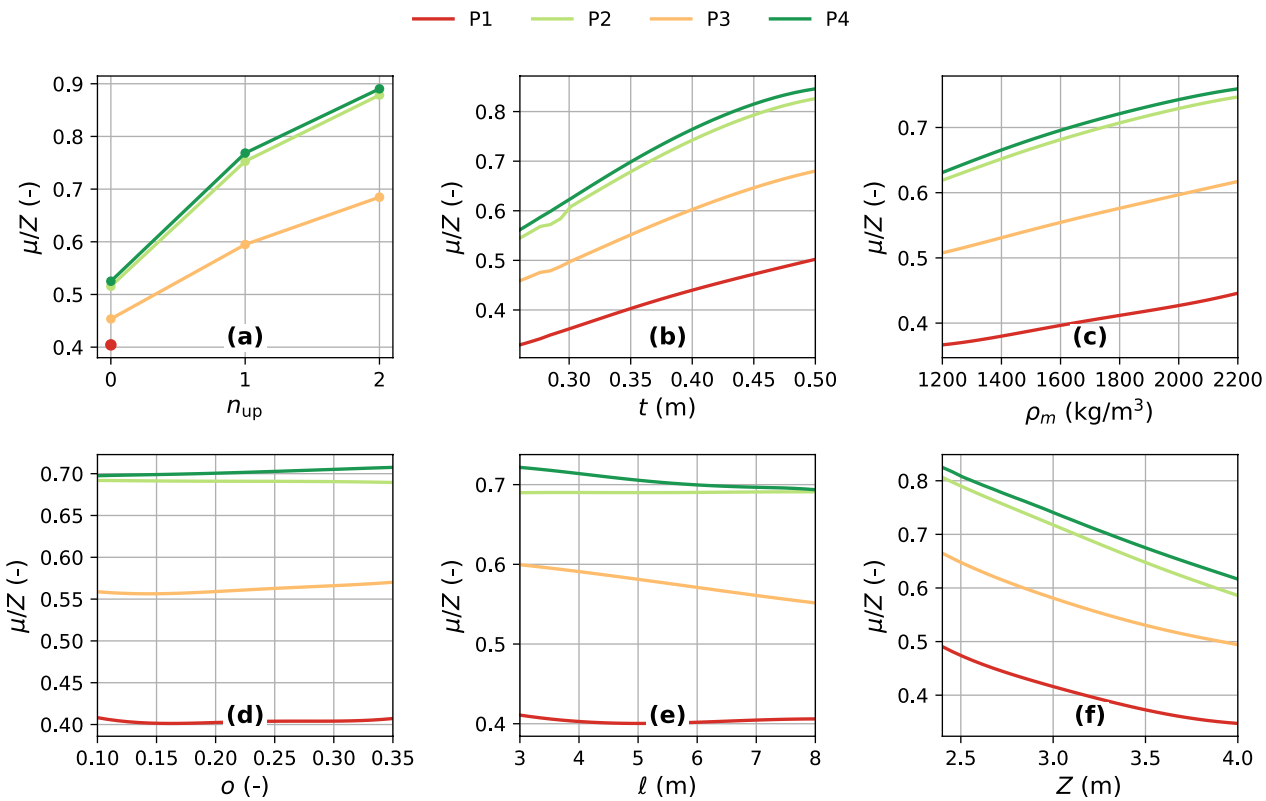


Fig. 10 Partial-dependence-style responses of the normalized fragility median μ/Z for boundary conditions P1–P4. The panels show the effects of n_{up}, wall thickness t, masonry density ρ_m, opening ratio o, wall length ℓ, and floor height Z

for which additional vertical load provides a positive but diminishing stabilising contribution.

Geometric parameters related to wall extent, such as plan dimensions and floor height, influence collapse capacity through their effect on overturning lever arms and the geometry of admissible failure mechanisms. The plotted wall-length variable ℓ shows a boundary-dependent response and remains less influential than wall thickness, upper-floor contribution, and masonry density in the selected partial-dependence profiles. Changes in ℓ affect collapse capacity through the mechanism geometry and, for laterally restrained configurations, through the effective contribution of cross-wall restraint. Similarly, the influence of floor height Z is weaker and less systematic than that of thickness, upper-floor contribution, or density. Under the absolute-depth formulation used here, changes in Z affect both the mechanism geometry and the conversion from dimensionless capacity to absolute flood depth. Therefore, the normalized response μ/Z, shown in Fig. 10 provides a complementary depth-to-height interpretation of the same surrogate predictions. Larger walls may be more prone to geometry-dependent mechanisms when their edges are restrained, because the failure shape is then controlled more by admissible mechanism geometry than by simple rocking.

The opening ratio, by comparison, plays a secondary and low-magnitude role, primarily modulating effective resisting width and local lever arms rather than altering the collapse mechanism. These interpretations are consistent with prior limit-analysis and numerical studies showing that boundary restraint promotes geometry-dependent wedge mechanisms in masonry walls (Asad et al., 2024; Noor-E-Khuda and Thambiratnam, 2024). This interpretation is specific to the adopted effective-width representation of openings. In more detailed masonry models, openings may also affect local stress redistribution, pier and spandrel behaviour, lintel action, and crack patterns. These mechanisms provide a clear direction for future extensions using richer geometric descriptors.

Overall, the marginal relations for $\mu(t)$ and $\mu(n_{up})$ provide compact, physically interpretable summaries of the dominant trends observed in the surrogate predictions.

4.2 Prediction of dispersion and uncertainty descriptors

The comparatively lower predictive performance observed for the dispersion parameter β, relative to the

fragility median μ, reflects the different physical meaning of the two targets. The median μ represents the central collapse capacity and is mainly controlled by nominal wall-scale descriptors such as t, Z, ρ_m, n_{up}, o, ℓ, and boundary condition. These quantities are explicitly included in the surrogate input vector, which explains the very high accuracy achieved for μ.

By contrast, β describes the spread of collapse depths generated by Monte Carlo perturbations around each nominal wall configuration. It is therefore more closely related to uncertainty propagation than to nominal collapse capacity alone. In the present framework, the surrogate inputs describe the nominal wall configuration, while descriptors of the uncertainty model itself, such as coefficients of variation, parameter standard deviations, correlation structure, spatial material heterogeneity, or boundary-restraint uncertainty, are not varied as independent inputs. Quantities such as σ_t, σ_{ρ_m}, and σ_o are used internally to perturb the simulator inputs, but they remain fixed across the dataset. The surrogate therefore learns the response under a prescribed uncertainty model, which explains why β is more difficult to predict than μ, even though the selected ensemble model still provides good accuracy.

The lower coefficient of determination for β may also reflect local nonlinear amplification of uncertainty. Near mechanism transitions, strongly restrained configurations, or regions where small perturbations modify the admissible collapse mechanism, similar nominal wall configurations can produce different levels of dispersion in collapse depth. Future extensions could improve dispersion prediction by augmenting the input space with uncertainty descriptors such as CoV_t, CoV_{ρ_m}, CoV_o, or corresponding standard deviations (JCSS, 2001; Silva et al., 2023; Nale et al., 2023); explicit opening descriptors such as opening width, height, position, pier width, and spandrel height (Liberatore et al., 2020); and material or interface descriptors such as tensile strength, mortar cohesion, friction coefficient, or degradation indicators (JCSS, 2001; Silva et al., 2023). Including such variables would help distinguish nominal capacity trends from the uncertainty mechanisms that control fragility dispersion.

4.3 Role of boundary restraint and interaction effects

Boundary conditions play a central role in modulating both the magnitude and structure of fragility sensitivity. For the base-only configuration (P1), wall response is governed mainly by sectional properties, and upper floors have no influence by definition because $n_{\text{up}} = 0$ is imposed for this mechanism. As restraint is introduced through either top or lateral connectivity (P2 and P3), vertical loads, wall thickness, and geometric descriptors act jointly, and interaction effects become non-negligible. Under full restraint (P4), the response becomes more sensitive to combined geometric and boundary effects. Parameters such as wall span and opening ratio may therefore contribute through mechanism geometry, although their influence remains secondary compared with vertical load contribution and thickness-related resistance. This behaviour is consistent with statically redundant configurations, in which local geometric discontinuities can alter load paths and collapse mechanisms.

The ordering of collapse capacity observed in the heatmap in Fig. 7 can be interpreted in terms of wall connectivity. Configurations without top restraint (P1 and P3) form the lower-capacity group, whereas configurations with top restraint (P2 and P4) form the higher-capacity group. This highlights the stabilising role of upper restraint and vertical load contribution in the adopted model, while also confirming that the absolute level of collapse capacity remains strongly dependent on the assigned boundary condition.

Guideline curves for engineering interpretation. To complement the boundary-specific interaction effects discussed in Sects. 3.2 and 4.3, Fig. 11 presents a set of guideline curves of the fragility median μ intended to support engineering interpretation and preliminary screening. These curves do not introduce additional results but represent controlled evaluations of the surrogate models under fixed, physically consistent conditions, translating the learned relationships into a compact and design-oriented form.

Figure 11 shows the predicted fragility median μ as a function of masonry density ρ_m, parameterized by wall thickness t and resolved by boundary condition (P1–P4). To isolate the combined effects of density, thickness, and boundary restraint, the remaining inputs are fixed to representative values ($o = 0.25$, $\ell = 5$ m, $Z = 3$ m, with $n_{\text{up}} = 0$ for P1 and $n_{\text{up}} = 2$ for P2–P4), consistent with the constrained low-rise masonry-wall domain and with the heatmap analysis presented earlier in Sect. 3.2. For consistency with the constrained sampling rules, the lowest displayed thickness is boundary-dependent: $t = 0.20$ m is used for P1, where $n_{\text{up}} = 0$, whereas $t = 0.24$ m is used as the lowest strictly valid thickness for P2–P4, where $n_{\text{up}} = 2$.

These guideline curves are intended for preliminary screening and engineering interpretation under the assumptions of the adopted hydrostatic, non-sliding, no-scour limit-analysis model. Within this scope, they provide a compact visualization of the dominant mechanisms governing flood-induced out-of-plane wall instability and support transparent comparison of wall configurations. Detailed assessment remains necessary

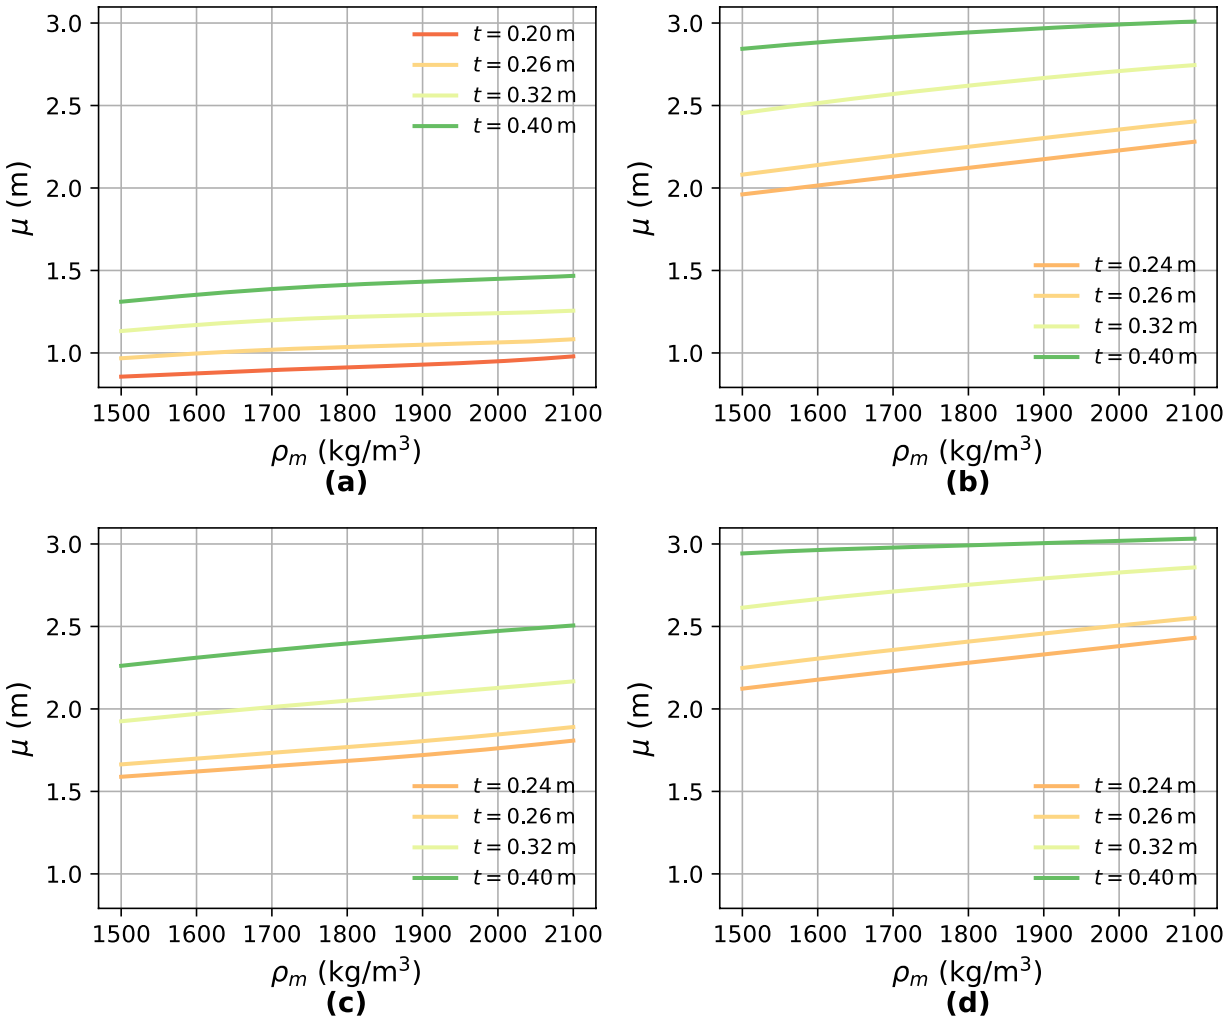


Fig. 11 Guideline curves of the predicted fragility median μ as a function of masonry density ρ_m, parameterized by wall thickness t, for the four boundary conditions considered. Panels (a)–(d) correspond to boundary conditions P1–P4, respectively. The remaining parameters are held constant ($o = 0.25$, $\ell = 5$ m, $Z = 3$ m), with $n_{\text{up}} = 0$ for P1 and $n_{\text{up}} = 2$ for P2–P4. The curves are intended for preliminary interpretation only and are valid within the idealized hydrostatic, non-sliding, no-scour assumptions of the adopted analytical model

for dynamic flood environments, debris-impact scenarios, or cases where sliding, local damage, or foundation degradation may control the response. They help answer practical questions such as: *"Given an existing wall with known thickness, density, and boundary condition, what collapse depth can be expected relative to another wall?"*

For example, under full restraint (P4), a wall with thickness $t = 0.32$ m and density 1800 kg/m^3 exhibits a median collapse depth approximately 0.5 m higher than a thinner wall with $t = 0.24$ m at the same density. Conversely two walls with the same thickness $t = 0.26$ m but densities of 1600 and 2000 kg/m^3, respectively, differ in collapse capacity by roughly 0.20 m. This comparison illustrates that, for otherwise similar existing walls, differences in thickness typically lead to larger changes in collapse capacity than differences in density, while the absolute values remain strongly dependent on the boundary configuration that shifts collapse to greater flood depths.

4.4 Computational efficiency and scalability

Surrogate modelling is widely used to approximate expensive simulation outputs at substantially lower computational cost, enabling large-scale evaluation and uncertainty analysis (Shirzadi et al., 2025). Similar motivations have been reported in structural reliability and fragility applications, where repeated Monte Carlo or nonlinear structural simulations can become

computationally demanding and surrogate models are used to reduce the cost of reliability evaluation (Skandalos et al., 2022). Beyond predictive accuracy, a key advantage of the proposed framework is its computational efficiency relative to direct mechanics-based simulation. The reference approach requires Monte Carlo evaluation of collapse mechanisms through incremental flood-depth analysis and stability checks, and for some boundary conditions also requires internal optimization such as hinge-position search.

Once trained, the surrogate predicts the fragility parameters (μ, β) through algebraic feature transformations and regression model evaluations. If t_{sim} denotes the average time required for one mechanics-based evaluation and t_{ml} the corresponding surrogate inference time, the single-instance speedup is

$$S_1 = \frac{t_{\text{sim}}}{t_{\text{ml}}}, \qquad t_{\text{ml}} \ll t_{\text{sim}}. \tag{30}$$

For N independent wall configurations, the total runtime scales approximately as

$$T_{\text{sim}}(N) \approx N\, t_{\text{sim}}, \qquad T_{\text{ml}}(N) \approx N\, t_{\text{ml}}, \tag{31}$$

so that the absolute time savings increase with the number of evaluated wall instances.

In the present study, the computational advantage was quantified by comparing boundary-weighted mechanics-based runtimes with measured post-training surrogate inference times. The total simulator time was estimated as

$$T_{\text{sim}}(N) = \sum_{i=1}^{4} N_{P_i} t_{P_i}, \tag{32}$$

where N_{P_i} is the number of walls with boundary condition P_i, and t_{P_i} is the corresponding mean mechanics-based evaluation time per wall. The measured mean simulator times were 0.0086 s for P1, 2.6250 s for P2, 0.0267 s for P3, and 3.3795 s for P4. All timings were measured on a Windows 10 workstation with 8 logical CPUs, using Python 3.11.9 and scikit-learn 1.2.0 in CPU-only mode. Similar computational gains have been reported in seismic fragility applications; for example, (He et al., 2024) reported that a GRU-based surrogate completed 6,100 structural response evaluations in a few seconds, compared with approximately 5 h using the reference Bouc–Wen model.

Using the boundary-weighted estimate, the mechanics-based runtimes for $N \in \{10^2, 10^3, 10^4\}$ were approximately 135.20 s, 22.53 min, and 3.76 h, respectively. The corresponding surrogate runtimes were 0.0398 s, 0.1688 s, and 1.3543 s, giving speedups of approximately 3.39×10^3, 8.01×10^3, and 9.98×10^3. For the largest case, the trained surrogate models predicted both fragility parameters for 10,000 wall configurations in 1.35 s, compared with an estimated 3.76 h for the corresponding mechanics-based evaluations.

These results demonstrate a substantial reduction in computational cost after training. This scalability is particularly relevant for regional risk assessments, parametric studies, uncertainty propagation, and rapid scenario screening, where thousands of wall configurations may need to be evaluated repeatedly.

4.5 Validation scope, applicability boundaries, and future extensions

The preceding sections show that the proposed surrogates can reproduce the reference mechanics-based simulator with high accuracy and provide physically interpretable sensitivity trends. The practical use of these predictions depends on the scope and assumptions of the simulator used to generate the training data. This subsection therefore clarifies the validation level of the surrogate framework, the potential bias associated with the adopted limit-analysis assumptions, the intended engineering use of the wall-level fragility estimates, and the future extensions required for broader flood-risk applications.

Validation scope, simulator assumptions, and bias

The high predictive accuracy reported in Sect. 3.1 demonstrates that the proposed surrogates accurately emulate the adopted mechanics-based simulator. The physical validation level of the framework is therefore governed by the validation scope of the reference limit-analysis model. Quantifying physical model bias would require comparison with independent reference values from experiments, post-event observations, or higher-fidelity numerical simulations that explicitly include mechanisms outside the simplified formulation.

The reference formulation of (Milanesi et al., 2018) assumes rigid-block mechanisms, unlimited compressive strength and stiffness, no tensile resistance, no relative slip between adjacent blocks, homogeneous wall behaviour, and full foundation support without scour. That study partly examined the influence of these idealizations through nonlinear finite-element comparisons, showing good agreement for the benchmark mechanisms considered. It also noted that finite compressive strength would shift the pivoting hinges inward and increase vulnerability, although this effect was limited for the tested configurations; variations in compressive strength and stiffness had marginal influence, while tensile strength produced larger but still limited variations. Other assumptions, such as no sliding and no scour, were acknowledged as part of the simplified model scope rather than fully quantified.

Consequently, the direction of potential bias is mechanism-dependent. The surrogate predictions may become unconservative when the real response is governed by mechanisms outside the adopted simulator, such as mortar-joint sliding, interface cracking, local crushing, foundation degradation, or support loss before the idealized overturning mechanism develops. This interpretation is consistent with interface-based masonry models, which show that cohesive–frictional joint behaviour is important for reproducing cracking and sliding mechanisms (Serpieri et al., 2017), and with flood-scour studies showing that foundation exposure and stiffness degradation can increase structural vulnerability (Dong et al., 2022). Thus, ignoring sliding or scour may lead to overestimation of collapse capacity when those mechanisms govern the actual response. This limitation does not affect the surrogate's ability to emulate the adopted simulator, but it defines the physical applicability domain of the present framework.

Engineering applicability and vulnerability-modelling scope

Within this applicability domain, the proposed framework should be viewed as a mechanics-informed complement to empirical depth–damage and indicator-based vulnerability models. Empirical functions remain useful for large-scale loss estimation, but they aggregate different structural behaviours into simplified relationships and cannot explicitly represent boundary restraint, vertical connectivity, wall thickness, or collapse mechanism. Indicator-based models include multiple vulnerability attributes, but their weighting may remain qualitative and may not capture the nonlinear interaction between geometry, material properties, and restraint conditions. By contrast, the present surrogate framework retains a direct link to a mechanics-based collapse model while providing the computational efficiency required for large-scale use.

The proposed framework is primarily intended for rapid screening, regional-scale vulnerability assessment, and preliminary engineering evaluation of masonry wall fragility under flood loading. The outputs of the framework are wall-level fragility parameters for individual exposed masonry wall units under specified boundary conditions. For building-level applications, these wall-level predictions can be embedded within a broader system-level assessment that accounts for diaphragm action, wall-to-floor and wall-to-wall connections, cross-wall interaction, load redistribution, structural redundancy, and possible progressive collapse. This distinction is important because ignoring system-level interaction may either overestimate risk, by neglecting redundancy and alternative load paths, or underestimate risk, by neglecting progressive collapse triggered by the failure of a critical wall.

The applicability of the framework is also bounded by the hydrostatic loading assumption of the reference simulator. Therefore, the present formulation is most appropriate for quasi-static or hydrostatic-equivalent flood actions, where the peak demand can be represented by a pressure distribution associated with flood depth or maximum run-up. However, the present analysis does not consider more complex scenarios such as those including high-velocity floods, debris-impact scenarios, or rapidly varying hydraulic conditions, where hydrodynamic demand, impulsive loading, and transient damage mechanisms may control the response.

Future extensions using richer simulators and additional input features

Future extensions of the framework should therefore focus on enriching the mechanics-based data-generation stage rather than only modifying the ML algorithms. The same surrogate methodology could be applied to training datasets generated from higher-fidelity numerical models or enriched analytical formulations that explicitly include joint sliding, cohesive–frictional interface behaviour, finite compressive strength, local crushing, foundation degradation, and scour. Similarly, the hydraulic input space could be expanded to include flow velocity, Froude number, hydrodynamic pressure coefficients, debris-impact demand, flood duration, and rate of water-level rise, allowing the surrogate to move beyond hydrostatic or hydrostatic-equivalent loading assumptions. At the building scale, future work should couple wall-level fragility predictions with system-level models that account for diaphragm action, cross-wall interaction, load redistribution, redundancy, and progressive collapse. For inventory-specific applications, the sampling strategy should also be calibrated to regional exposure data, so that input ranges and correlations among wall geometry, openings, material properties, and construction typology reflect the target building stock. This is important because model-agnostic feature-importance and partial-dependence measures depend on the data distribution and sampled input domain (Molnar et al., 2023). Finally, inventory-informed uncertainty models could be introduced by varying not only nominal wall properties and sampling rules but also uncertainty descriptors, such as coefficients of variation, material-strength variability, boundary-restraint uncertainty, and foundation-condition indicators. These extensions would preserve the computational advantages of the proposed surrogate approach while broadening its applicability to more realistic flood-damage scenarios.

5 Conclusion

This study proposed and evaluated a mechanics-informed surrogate modelling framework for predicting and interpreting the flood fragility parameters of masonry walls subjected to out-of-plane loading. By combining a mechanics-based collapse simulator with complementary ML surrogates, the framework achieves near-simulator accuracy for the fragility median μ while maintaining physical transparency across the constrained low-rise masonry-wall domain considered in this study. The approach bridges the gap between detailed analytical modelling and scalable flood-risk assessment, enabling efficient fragility estimation while retaining consistency with the adopted mechanics-based simulator.

The results demonstrate that surrogate models can reliably reproduce full fragility curves when the median capacity is accurately captured. Although the dispersion parameter β exhibits greater variability and comparatively lower predictability than μ, this behaviour reflects its dependence on uncertainty propagation around each nominal wall configuration, which is not fully described by the current macroscopic input features. Nevertheless, curve-level evaluation confirms that focusing sensitivity and interpretability analysis on μ preserves probabilistic collapse behaviour relevant for engineering decision-making. Unified sensitivity diagnostics revealed stable and low-dimensional response structures, with upper-floor contribution, wall thickness, boundary restraint, and the coupled thickness–height descriptor emerging as the dominant sensitivity group, while masonry density remained secondary and opening-related variables showed low first-order influence within the adopted model formulation. The feature-importance analysis, partial-dependence profiles, risk heatmap, and guideline curves provide complementary levels of interpretation, linking the surrogate predictions to the governing mechanics of wall overturning and restraint. The sign-level consistency between sparse interpretable surrogates and nonlinear models further supports the robustness and physical interpretability of the proposed framework.

Future research can build on the present framework by enriching both the simulator and the input space while preserving the same surrogate-modelling workflow. Additional loading and failure mechanisms, such as hydrodynamic forces, impact actions, sliding and contact behaviour, local crushing, foundation degradation, scour, and soil–structure interaction effects, could be incorporated through higher-fidelity simulators or extended analytical formulations. Coupling wall-level surrogates with simplified building-scale models will enable assessment of system-level collapse and damage propagation. Further work should also explore enriched input spaces including uncertainty descriptors, explicit opening geometry, boundary-restraint stiffness, and foundation-condition indicators, as well as inventory-informed sampling strategies calibrated to regional exposure data, so that input ranges and correlations better reflect target building stocks. Hybrid calibration strategies that integrate post-event damage observations with mechanics-based synthetic data could further improve transferability across regions and construction practices.

Overall, the proposed surrogate modelling approach combines computational efficiency, simulator-level accuracy, and physical interpretability. It provides a practical route for rapid flood fragility assessment of masonry walls, supports transparent engineering interpretation of the dominant collapse drivers, and offers a transferable methodology for analogous problems in structural and hazard risk modelling.

Acknowledgements
The authors gratefully acknowledge the financial support of the Italy-Croatia Interreg Project STRENGTH (ID: ITHR0200318). The authors also thank the anonymous reviewers for their constructive comments and suggestions, which helped improve the clarity and quality of the manuscript.

Author contributions
SG contributed to the methodology development, analysis, and manuscript preparation. MN contributed to the methodology development and manuscript preparation. MF contributed to the manuscript preparation. ZN contributed to the methodology development and manuscript preparation. EB contributed to the methodology development, manuscript preparation, and funding. All authors read and approved the final manuscript.

Funding
This research has received financial support from the Italy-Croatia Interreg Project STRENGTH (ID: ITHR0200318).

Data availability
The data supporting the findings of this study are available from the corresponding author upon reasonable request.

Materials availability
Not applicable.

Code availability
The code used to generate the results of this study is available from the corresponding author upon reasonable request.

Declarations

Ethics approval and consent to participate
All authors have consented to the publication of this work. No ethical issues are associated with this study.

Consent for publication
All contributors have reviewed and approved the final version of the manuscript for publication.

Competing interests
The authors declare no competing interests.

Received: 9 February 2026 Revised: 16 June 2026 Accepted: 13 October 2026
Published online: 09 November 2026

References

Alizadeh, R., Allen, J. K., & Mistree, F. (2020). Managing computational complexity using surrogate models: A critical review. *Research in Engineering Design, 31*, 275–298. https://doi.org/10.1007/s00163-020-00336-7

Asad, M., Thamboo, J., Zahra, T., & Thambiratnam, D. P. (2024). Out of-plane failure analysis of masonry walls subjected to hydrostatic and hydrodynamic loads induced by flash floods. *Engineering Structures, 298*, Article 117041. https://doi.org/10.1016/j.engstruct.2023.117041

Banik, S.S., Hong, H.P.: Wind-induced fragility of a monopole structure via artificial neural network based surrogate analysis. Engineering Structures 278, 115515 (2023) https://doi.org/10.1016/j.engstruct.2022.115515

Besarra, I., Opdyke, A., Aquino, D. H., Santiago, J., Mendoza, J. E., & Lagmay, A. M. F. A. (2025). Flood fragility and vulnerability functions for residential buildings in the province of Leyte, Philippines. *Journal of Flood Risk Management*. https://doi.org/10.1111/jfr3.13043

Bhattarai, Y., Chaudhary, V., Walker, C., Talchabhadel, R., & Sharma, S. (2025). Ensemble learning for enhancing critical infrastructure resilience to urban flooding. *Scientific Reports, 15*, 36901. https://doi.org/10.1038/s41598-025-20970-2

Bodda, S. S., Gupta, A., & Sewell, R. T. (2021). Application of risk-informed validation framework to a flooding scenario. *ASCE-ASME Journal of Risk and Uncertainty in Engineering Systems, Part A: Civil Engineering*. https://doi.org/10.1061/ajrua6.0001172

Capparelli, G., Liguori, F. S., Madeo, A., & Versace, P. (2023). A rapid numerical-based vulnerability assessment method for masonry buildings subject to flood. *International Journal of Disaster Risk Reduction, 97*, Article 104001. https://doi.org/10.1016/j.ijdrr.2023.104001

Colajanni, P., Cucchiara, C., D'Anna, J., Pennisi, S., & Pagnotta, S. (2023). Vulnerability and seismic exposure of residential building stock in the historic center of alcamo. *Applied Sciences, 13*(12), 7092. https://doi.org/10.3390/app13127092

Cuomo, S., Di Cola, V. S., Giampaolo, F., Rozza, G., Raissi, M., & Piccialli, F. (2022). Scientific machine learning through physics-informed neural networks: Where we are and what's next. *Journal of Scientific Computing, 92*, 88. https://doi.org/10.1007/s10915-022-01939-z

Destro Bisol, G., Prajapati, S., Sorrentino, L., & AlShawa, O. (2024). Vertical spanning wall elastically restrained at the top: Validation and parametric dynamic analysis. *Bulletin of Earthquake Engineering, 22*, 3951–3978. https://doi.org/10.1007/s10518-024-01909-w

Dong, Y., & Frangopol, D. M. (2017). Probabilistic life-cycle cost-benefit analysis of portfolios of buildings under flood hazard. *Engineering Structures, 142*, 290–299. https://doi.org/10.1016/j.engstruct.2017.03.063

Dong, Z., Frangopol, D. M., & Li, Y. (2022). Failure risk assessment method of masonry structures under earthquakes and flood scouring. *Mechanics Based Design of Structures and Machines, 50*(10), 3602–3625. https://doi.org/10.1080/15376494.2021.1884322

FEMA factsheet: Flood assessment structure tool (FAST). Technical report, Federal Emergency Management Agency (FEMA) (2022). Factsheet describing the FAST tool for building-level flood risk estimation. https://www.fema.gov/

Garzón, A., Kapelan, Z., Langeveld, J., & Taormina, R. (2022). Machine learning-based surrogate modeling for urban water networks: Review and future research directions. *Water Resources Research*. https://doi.org/10.1029/2021WR031808

He, W., Li, J., Kong, X., & Deng, L. (2024). Multi-level physics informed deep learning for solving partial differential equations in computational structural mechanics. *Communications Engineering, 3*, 3, Article 151. https://doi.org/10.1038/s44172-024-00303-3

He, S., Liao, Y., Sun, P. P., & Zhang, R. (2024). Deep learning enabled seismic fragility evaluation of structures subjected to mainshock-aftershock earthquakes. *Urban Lifeline, 2*, Article 2. https://doi.org/10.1007/s44285-024-00013-4

Heyman, J. (1999). The stone skeleton. *Essays in the History of Structural Engineering* (pp. 89–112). Springer. https://doi.org/10.1007/978-1-4757-2006-8_7

Holzinger, A., Longo, L., Cangelosi, A., & Del Ser, J. (2026). Research frontiers in machine learning & knowledge extraction. *Machine Learning and Knowledge Extraction, 8*(1), 6. https://doi.org/10.3390/make8010006

Huizinga, J., De Moel, H., Szewczyk, W.: Global flood depth-damage functions: Methodology and the database with guidelines. Technical report, Joint Research Centre (JRC), European Commission (2017). https://doi.org/10.2760/16510 . https://ec.europa.eu/jrc

Hwang, J. T., & Martins, J. R. R. A. (2018). A fast-prediction surrogate model for large datasets. *Aerospace Science and Technology, 75*, 74–87. https://doi.org/10.1016/j.ast.2017.12.030

Jalayer, F., Carozza, S., De Risi, R., Manfredi, G., & Mbuya, E. (2016). Performance-based flood safety-checking for non-engineered masonry structures. *Engineering Structures, 106*, 109–123. https://doi.org/10.1016/j.engstruct.2015.10.042

Jansen, L., Korswagen, P. A., Bricker, J. D., Pasterkamp, S., de Bruijn, K. M., & Jonkman, S. N. (2020). Experimental determination of pressure coefficients for flood loading of walls of Dutch terraced houses. *Engineering Structures, 216*, Article 110647. https://doi.org/10.1016/j.engstruct.2020.110647

Joint Committee on Structural Safety: JCSS Probabilistic Model Code. https://www.jcss-lc.org/jcss-probabilistic-model-code/. Parts I–III, including sections on dimensions, masonry, and model uncertainties (2001)

Karniadakis, G. E., Kevrekidis, I. G., Lu, L., Perdikaris, P., Wang, S., & Yang, L. (2021). Physics-informed machine learning. *Nature Reviews Physics, 3*, 422–440. https://doi.org/10.1038/s42254-021-00314-5

Kim, S. H., & Boukouvala, F. (2020). Machine learning-based surrogate modeling for data-driven optimization: A comparison of subset selection for regression techniques. *Optimization Letters, 14*, 989–1010. https://doi.org/10.1007/s11590-019-01428-7

Korswagen, P. A., Bricker, J. D., Jonkman, S. N., & Kolen, B. (2025). Experimental and numerical analysis of a masonry wall under hydrostatic and debris impact loads. *Frontiers in Built Environment, 11*, Article 1644699. https://doi.org/10.3389/fbuil.2025.1644699

Liberatore, L., AlShawa, O., Marson, C., Pasca, M., & Sorrentino, L. (2020). Out-of-plane capacity equations for masonry infill walls accounting for openings and boundary conditions. *Engineering Structures, 207*, Article 110198. https://doi.org/10.1016/j.engstruct.2020.110198

Lieu, Q. X., Nguyen, K. T., Dang, K. D., Lee, S., Kang, J., & Lee, J. (2022). An adaptive surrogate model to structural reliability analysis using deep neural network. *Expert Systems with Applications, 189*, Article 116104. https://doi.org/10.1016/j.eswa.2021.116104

Mendoza Cabanzo, C., Santamaría, M., Sousa, H. S., & Matos, J. C. (2022). In-plane fragility and parametric analyses of masonry arch bridges exposed to flood hazard using surrogate modeling techniques. *Applied Sciences, 12*(4), 1886. https://doi.org/10.3390/app12041886

Milanesi, L., Pilotti, M., Belleri, A., Marini, A., & Fuchs, S. (2018). Vulnerability to flash floods: A simplified structural model for masonry buildings. *Water Resources Research, 54*(10), 7177–7197. https://doi.org/10.1029/2018WR022577

Milani, G., Zuccarello, F. A., Olivito, R. S., & Tralli, A. (2007). Heterogeneous upper-bound finite element limit analysis of masonry walls out-of-plane loaded. *Computational Mechanics, 40*, 911–931. https://doi.org/10.1007/s00466-006-0151-9

Molnar, C., König, G., Bischl, B., & Casalicchio, G. (2023). Model-agnostic feature importance and effects with dependent features: A conditional subgroup approach. *Data Mining and Knowledge Discovery, 37*, 2903–2941. https://doi.org/10.1007/s10618-022-00901-9

Nale, M., Benvenuti, E., Chiozzi, A., Minghini, F., & Tralli, A. (2023). Effect of uncertainties on seismic fragility for out-of-plane collapse of unreinforced masonry walls. *Journal of Building Engineering*. https://doi.org/10.1016/j.jobe.2023.106936

Nasiri, H., Yusof, M. J. M., Ali, T. A. M., & Hussein, M. K. B. (2019). District flood vulnerability index: Urban decision-making tool. *International Journal of Environmental Science and Technology, 16*(5), 2249–2258. https://doi.org/10.1007/s13762-018-1797-5

Nofal, O. M., & van de Lindt, J. W. (2020). Minimal building flood fragility and loss function portfolio for resilience analysis at the community level. *Water*. https://doi.org/10.3390/w12082277

Nofal, O. M., van de Lindt, J. W., & Do, T. Q. (2020). Multi-variate and single-variable flood fragility and loss approaches for buildings. *Reliability Engineering & System Safety*. https://doi.org/10.1016/j.ress.2020.106971

Noor-E-Khuda, S., & Thambiratnam, D. P. (2024). Unreinforced masonry walls under two-way out-of-plane bending: Comparison of FE model predictions with analytical models. *Structures, 61*, Article 106126. https://doi.org/10.1016/j.istruc.2024.106126

Nur, I., & Shrestha, K. K. (2017). An integrative perspective on community vulnerability to flooding in cities of developing countries. *Procedia Engineering, 198*, 958–967. https://doi.org/10.1016/j.proeng.2017.07.141

Raissi, M., Perdikaris, P., & Karniadakis, G. E. (2019). Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics, 378*, 686–707. https://doi.org/10.1016/j.jcp.2018.10.045

Refadah, S. S. (2025). Development in flood forecasting: A comprehensive review of complex and machine learning models. *Physics and Chemistry of the Earth, Parts A/B/C, 139*, Article 103975. https://doi.org/10.1016/j.pce.2025.103975

Rocchi, A., Chiozzi, A., Nale, M., Nikolic, Z., Riguzzi, F., Mantovan, L., Gilli, A., & Benvenuti, E. (2022). A machine learning framework for multi-hazard risk assessment at the regional scale in earthquake and flood-prone areas. *Applied Sciences*. https://doi.org/10.3390/app12020583

Russo, B., Velasco, M., Locatelli, L., Sunyer, D., Yubero, D., Monjo, R., Martínez-Gomariz, E., Forero-Ortiz, E., Sánchez-Muñoz, D., Evans, B., & Gonzalez Gómez, A. (2020). Assessment of urban flood resilience in Barcelona for current and future scenarios. The RESCCUE project. *Sustainability*. https://doi.org/10.3390/su12145638

Serpieri, R., Albarella, M., Sacco, E.: A 3d microstructured cohesive–frictional interface model for the analysis of brick-masonry structures. International Journal of Solids and Structures 110–111, 165–178 (2017) https://doi.org/10.1016/j.ijsolstr.2017.01.036

Shirzadi, N., Lau, D., & Stylianou, M. (2025). Surrogate modeling for building design: Energy and cost prediction compared to simulation-based methods. *Buildings, 15*, Article 2361. https://doi.org/10.3390/buildings15132361

Shrestha, A., & Chini, C. M. (2025). Risk-based stormwater fragility curves for urbanized coastal flooding. *Journal of Hydrology*. https://doi.org/10.1016/j.jhydrol.2025.133436

Silva, L. C. M., Milani, G., & Lourenço, P. B. (2023). Probabilistic-based discrete model for the seismic fragility assessment of masonry structures. *Structures, 52*, 506–523. https://doi.org/10.1016/j.istruc.2023.04.015

Skandalos, K., Chakraborty, S., & Tesfamariam, S. (2022). Seismic reliability analysis using a multi-fidelity surrogate model: Example of base-isolated buildings. *Structural Safety, 97*, Article 102222. https://doi.org/10.1016/j.strusafe.2022.102222

Tosto, C., Leggieri, V., Ruggieri, S., & Uva, G. (2023). Investigation of architectural typological parameters influencing seismic vulnerability of masonry buildings in historical centres: The case of Puglia. *Procedia Structural Integrity, 44*, 2036–2043. https://doi.org/10.1016/j.prostr.2023.01.260

Tsokanas, N., Pastorino, R., & Stojadinović, B. (2022). A comparison of surrogate modeling techniques for global sensitivity analysis in hybrid simulation. *Machine Learning and Knowledge Extraction, 4*(1), 1–21. https://doi.org/10.3390/make4010001

Wang, Y., Zhang, Q., Lin, K., Liu, Z., Liang, Y., Liu, Y., & Li, C. (2024). A novel framework for urban flood risk assessment: Multiple perspectives and causal analysis. *Water Research*. https://doi.org/10.1016/j.watres.2024.121591

Willard, J., Jia, X., Xu, S., Steinbach, M., & Kumar, V. (2022). Integrating physics-based modeling with machine learning: A survey. *ACM Computing Surveys, 55*(4), 1–34. https://doi.org/10.1145/3497769

Williams, B.A., Cremaschi, S.: Surrogate model selection for design space approximation and surrogatebased optimization. In: Proceedings of the 9th International Conference on Foundations of Computer-Aided Process Design, pp. 353–358. Elsevier, Amsterdam, Netherlands (2019). https://doi.org/10.1016/B978-0-12-818597-1.50056-4

Wing, O. E. J., Pinter, N., Bates, P. D., & Kousky, C. (2020). New insights into US flood vulnerability revealed from flood insurance big data. *Nature Communications*. https://doi.org/10.1038/s41467-020-15264-2

Yarveysi, F., Jafarzadegan, K., Tripathy, S. S., Moftakhari, H., & Moradkhani, H. (2025). A data-driven framework for an efficient block-level coastal flood risk assessment. *International Journal of Disaster Risk Reduction*. https://doi.org/10.1016/j.ijdrr.2025.105478

Del Zoppo, M., Belliazzi, S., Lignola, G. P., Di Ludovico, M., & Prota, A. (2022). Analytical structural vulnerability and classification of italian residential building stock to tsunami hazard. *Bulletin of Geophysics and Oceanography, 63*(4), 555–574. https://doi.org/10.4430/bgo00398

Zuccaro, G., Dolce, M., Perelli, F. L., De Gregorio, D., & Speranza, E. (2023). Cartis: A method for the typological-structural characterization of Italian ordinary buildings in urban areas. *Frontiers in Built Environment, 9*, Article 1129176. https://doi.org/10.3389/fbuil.2023.1129176

Publisher's Note

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