TU Delft

Delft University of Technology

Document Version
Final published version

Licence
CC BY

Citation (APA)
Kazemi, P., Turrin, M., Andriotis, C., Entezami, A., Mariani, S., & Ghisi, A. (2026). Surrogate-assisted multi-objective seismic design optimisation of diagrid tall buildings. *Computer-Aided Civil and Infrastructure Engineering*, *50*, Article 100181. https://doi.org/10.1016/j.cacaie.2026.100181

Important note
To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright
In case the licence states “Dutch Copyright Act (Article 25fa)”, this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership.
Unless copyright is transferred by contract or statute, it remains with the copyright holder.

Sharing and reuse
Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy
Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

This work is downloaded from Delft University of Technology.

Computer-Aided Civil and Infrastructure Engineering 50 (2026) 100181

Contents lists available at ScienceDirect

Computer-Aided Civil and Infrastructure Engineering

journal homepage: www.sciencedirect.com/journal/computer-aided-civil-and-infrastructure-engineering

Research Article

Surrogate-assisted multi-objective seismic design optimisation of diagrid tall buildings

Pooyan Kazemi[a], Michela Turrin[b], Charalampos Andriotis[b], Alireza Entezami[a,*], Stefano Mariani[a], Aldo Ghisi[a]

[a] Department of Civil and Environmental Engineering, Politecnico di Milano, Piazza Leonardo da Vinci 32, 20133 Milano, Italy
[b] Architectural Engineering and Technology Department, Delft University of Technology, Julianalaan 134, 2628 BL, Delft, the Netherlands

ABSTRACT

Early-stage design optimisation of tall buildings has been studied for several decades through performance-based design, parametric modelling, evolutionary algorithms and, more recently, machine-learning surrogates. A remaining challenge, however, is the computationally efficient integration of seismic time-history simulation, architectural form generation, multi-response prediction and multi-objective design exploration within a single reproducible workflow. This issue is particularly relevant for diagrid tall buildings, where geometric transformations such as tapering, twisting and curvilinearity strongly affect structural response, material use and environmental indicators. This study proposes a surrogate-assisted framework for early-stage multi-objective seismic design optimisation of diagrid tall buildings. A database of 1000 parametric building configurations was generated and analysed under seven bidirectional ground-motion records, resulting in 7000 OpenSeesPy time-history simulations. The original 22 input variables were transformed into 45 preprocessed input columns through scaling, categorical encoding and conditional masking of geometry-dependent variables. A multi-input, multi-output feed-forward neural network was then trained to predict 91 responses, including 69 structural responses, 14 environmental and cost-related responses and 8 geometry-dependent quantities. The trained surrogate was coupled with NSGA-II to screen candidate designs using selected objectives related to acceleration, displacement, inter-storey drift, von Mises stress and total mass. Approximately 2000 surrogate-based optimisation evaluations were produced over 10 generations and reduced to 408 unique candidate geometries. The proposed framework does not replace final nonlinear finite-element verification; rather, it provides a computationally efficient decision-support tool for rapidly reducing a high-dimensional diagrid design space before detailed engineering assessment.

1. Introduction and background

Designing tall buildings in regions subjected to moderate seismic actions is a complex task at the intersection of architecture and structural engineering. Early stage decisions regarding the form and structural system of a building deeply affect not only its seismic resilience but also its economic viability and environmental footprint. Consequently, there is a growing interest in integrating high-fidelity simulations into the early design process, thereby enabling the performance-based design of buildings (Negendahl, 2015). Recent tools (Gauch et al., 2022) linked form generation with structural and environmental metrics, highlighting a trade-off between carbon emissions and construction costs. Previous research, see e.g., (Granadeiro et al., 2013), demonstrated the benefits of embedding structural and environmental analyses in conceptual design, regarding architectural form generation. These developments underscored the need for early-stage optimisation in seismic regions to balance structural, environmental and economic objectives in tall building design.

To address this high-dimensional decision problem, the architecture, engineering and construction community increasingly adopts parametric modelling and artificial intelligence (AI)-driven methods.

Computational optimisation of structural systems has been an active research topic for several decades, and several important contributions have already demonstrated the value of artificial intelligence and evolutionary search in structural design. Early studies on computational intelligence showed that genetic algorithms, neural networks and mathematical optimisation could be combined for the automated design of large-scale and high-rise structural systems. More recent studies have extended this direction through parametric modelling, surrogate modelling, generative design and multi-objective optimisation. These developments confirm that neither neural networks nor evolutionary algorithms are new by themselves in the context of structural optimisation. The current challenge is instead how to integrate these tools with high-fidelity simulation databases and high-dimensional response prediction in a way that is computationally efficient, reproducible and useful for early-stage engineering decision-making.

* Corresponding author.
Email address: alireza.entezami@polimi.it (A. Entezami).

https://doi.org/10.1016/j.cacaie.2026.100181
Received 21 January 2026; Accepted 5 July 2026
Available online 25 August 2026
1093-9687/© 2026 The Authors. Published by Elsevier Inc. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

The application of computational intelligence to tall-building structural optimisation has a long history: Adeli and Kumar (1995) demonstrated concurrent genetic-algorithm-based optimisation of large steel structures on a massively parallel supercomputer, while Adeli and Park (1996) presented a hybrid AI model integrating neural networks and mathematical optimisation for the fully automated design of super-high-rise steel buildings.

While parametric design supports systematic exploration of geometrical and structural features, exhaustive seismic simulations for each variant are computationally prohibitive. This challenge drove the development of surrogate models and data-driven design frameworks. As an example, Pan et al. (2020) proposed a hybrid approach combining self-organizing maps with multi-layer perceptron neural networks (NNs) to support conceptual design exploration of building geometry and quantitative performance metrics. Machine learning (ML)-based surrogates offer near-instant proxies of simulation outcomes, enabling real-time design feedback. Techniques ranging from simple NNs (Pan et al., 2020) to advanced generative models are employed to predict structural and environmental responses, and guide early design decisions.

Balmer et al. (2024) recently developed a performance-driven design framework using a conditional variational autoencoder that acts as both a forward predictor of structural performance and an inverse generator of design features, enhanced by sensitivity analysis for interpretability. This solution reflects the growing interest in interpretable AI surrogates that not only replicate simulation outputs but also reveal how design parameters affect performance. Related ML regression frameworks have also mapped excitation and monitoring descriptors to structural vibration responses (Entezami et al., 2025; Entezami & Sarmadi, 2025; Entezami et al., 2026). Similar advances include hybrid deep learning (DL) surrogates for seismic performance prediction (Shabbir et al., 2025) and probabilistic frameworks incorporating uncertainty quantification in seismic analysis (Noureldin et al., 2023).

Alongside surrogate modelling, researchers embraced generative design and optimisation algorithms to explore the vast design space for tall buildings. Evolutionary strategies, particularly the non-dominated sorting genetic algorithm II (NSGA-II), proved effective in multi-objective design problems. Bailey and Caldas (2023) integrated NSGA-II with parametric massing studies to generate early-stage building forms optimised in terms of density and solar exposure. NSGA-II and related approaches were also used for seismic isolation systems (Guner et al., 2023) and robust multi-objective optimisation strategies (Nabian & Meidani, 2018).

Generative AI approaches enable exploring a huge number of design alternatives and identifying Pareto-optimal solutions that balance competing criteria. In structural design, ensuring the diversity of solutions while meeting performance targets is also crucial. Danhaive and Mueller (2021) addressed this issue by employing performance-conditioned generative models to explore diverse structural options without compromising engineering goals. Beyond structural and geometric considerations, occupant-centric criteria were integrated; for instance, agent-based simulations of post-earthquake evacuations informed early-stage layout decisions. Related efforts included autoencoder-based feature extraction for seismic learning (Ning & Xie, 2024), transformer-based contrastive denoising autoencoding for vibration-change detection in monitored bridges (Chen et al., 2026), and kernel-enhanced deep anomaly detection for bridge monitoring under freezing-induced variability (Chen et al., 2026). Explainable models such as quantum convolutional neural networks have also been investigated for seismic damage detection (Bhatta & Dang, 2024).

Recent advances have reflected a paradigm shift in AI-enhanced design, leveraging data-driven models, optimisation algorithms and simulation data to make informed decisions. However, their application to seismic performance optimisation of tall buildings remains limited. Previous studies often addressed individual aspects, such as energy efficiency (Singh et al., 2022), embodied carbon (Gauch et al., 2022), or simplified structural proxies, and often focused on low- or mid-rise structures. In contrast, high-rises present unique challenges due to the complex dynamic behaviour under seismic loads, which requires intensive numerical simulations. The high computational cost has hindered the integration of such fidelity in early design phases. This paper addresses this gap by proposing a DL approach to predict the seismic performance under moderate earthquakes across diverse tall building designs. A multi-input, multi-output feed-forward neural network (MIMO-FNN) has been developed and trained on an extensive dataset of buildings and simulated seismic responses. Unlike conventional single-output models, the MIMO-FNN jointly predicts multiple response metrics (e.g., inter-storey drifts, base shear, accelerations), capturing interdependencies among performance criteria. The model uses a parametric input space comprising architectural and geometric design variables together with record-specific ground-motion descriptors.

The novelty of the present study does not lie in introducing parametric modelling, FE simulation, neural-network surrogates or NSGA-II as isolated techniques, since each of these components has already been widely studied. Rather, the contribution lies in formulating a unified computational workflow specifically tailored to early-stage seismic design exploration of diagrid tall buildings. Diagrid structural systems have been recognised as structurally efficient, architecturally flexible and sustainable solutions for tall buildings (Asadi & Adeli, 2017), and the nonlinear seismic behaviour of mid- to high-rise diagrid configurations has been investigated in detail (Asadi & Adeli, 2018). Recent data-driven studies have also examined architectural–structural correlations, feature selection and data augmentation for early-stage diagrid building design (Kazemi et al., 2024b,2024a). However, existing studies have not embedded such systems within a surrogate-assisted multi-objective design-exploration framework that simultaneously accounts for geometric transformation, seismic response, material demand and sustainability-related indicators.

The proposed framework addresses this gap by combining: (i) automated generation of a broad family of tapered, twisted and curvilinear diagrid geometries; (ii) mechanics-based seismic simulations used to construct a database of 7000 OpenSeesPy bidirectional time-history analyses; (iii) conditional preprocessing of transformation-dependent geometric variables to preserve the physical meaning of the input design space; (iv) a MIMO-FNN surrogate that jointly predicts 91 structural, environmental and geometry-dependent responses; and (v) surrogate-assisted multi-objective optimisation in a discrete high-dimensional design space. This positions the work as a computational workflow-level and application-level contribution, rather than as a claim of standalone algorithmic novelty. The value of the proposed approach is that it enables rapid screening and ranking of candidate diagrid tall-building geometries at the early design stage, while maintaining a direct connection to mechanics-based FE simulation data.

It should be emphasised that the term optimisation in this study refers to early-stage seismic design optimisation. The design variables are the architectural and geometric descriptors of the diagrid tall-building configurations, while the optimisation objectives are selected structural and performance-related responses predicted by the trained surrogate model. The aim is therefore to identify promising candidate geometries for subsequent high-fidelity verification. The trained MIMO-FNN provides real-time predictions of structural performance, which are embedded in an NSGA-II optimisation routine to efficiently explore the design space. This integration allows identification of solutions that balance multiple structural-response and material-demand criteria. Trained on high-fidelity simulation data, the surrogate model displays strong predictive accuracy. In addition, it extracts interpretable insights, such as the influence of design parameters and correlations among performance metrics, to serve as a transparent decision-support tool rather than a black-box. The proposed workflow consists of four main steps: (1) automated generation of a large dataset of architectural and structural models; (2) seismic analysis using finite element (FE) simulations in a parallel computing environment; (3) collection and preparation of the simulation data for surrogate model development; and (4) exploitation

Table 1
Parametric design space and associated variables defining the architectural form-generation framework for the considered tall-building configurations.

Condition	Parameter	Name	Description	Cardinality	Values
	X1	Number of sides of the top plan	Number of polygon edges defining the top plan	10	[3,4,...,12]
	X2	Number of sides of the bottom plan	Number of polygon edges defining the bottom plan	10	[3,4,...,12]
	X3	Top–bottom plan orientation	Relative alignment between the top and bottom plan (vertex to vertex alignment = 0, vertex to edge = 1)	2	[0,1]
	X4	Total number of stories	Total number of floors	6	[40,44,48,...,60]
	X5	Inter-storey height	Distance between two consecutive floors (m)	3	[3.5,4.0,4.5]
	X6	Form-transformation approach along height	Procedure by which the top and bottom are joined through the elevation (tapered/twisted/curvilinear)	3	[0,1,2]
If X6 = 0	X7	Tapering	Relative area scaling of the top plan with respect to the bottom plan	6	[50%,60%,...,90%, 100%]
If X6 = 1	X8	Twisting angle	Angular rotation of the top plan (°)	5	[30,45,60,75,90]
If X6 = 2	X9	Vertical location of curvilinearity for reference story	Parameter to define the position of curvilinearity	3	[0.25,0.5,0.75]
	X10	Reference story curvilinearity scale factor	Parameter to define the intensity of curvilinearity	2	[0.85,1.15]
			Total solutions	**61,200**	

of the surrogate via an optimisation algorithm to generate new building systems that meet structural, architectural and sustainability objectives. Finally, some concluding remarks and proposals for future work are provided in Section 6. Steps 1–2 above were already described in Kazemi et al. (2025) and are only briefly summarised in Section 2. This work focuses instead on: the surrogate model creation in Section 3; MIMO-FNN training and validation in Section 4; and the use of optimisation as a generative AI tool for expanding the design space beyond the initial dataset in Section 5.

1.1. Computational contribution and paper organisation

The proposed study is organised around three distinct but connected components: the computational method, the simulation database and the optimisation application. The computational method consists of a surrogate-assisted design-exploration workflow in which a MIMO-FNN is trained to approximate a high-dimensional mapping between diagrid building descriptors, ground-motion descriptors and multiple response quantities. The simulation database provides the mechanics-based information required to train this surrogate and is generated through automated parametric modelling and OpenSeesPy time-history analyses. The optimisation application then uses the trained surrogate within NSGA-II to screen new candidate geometries at low marginal computational cost.

Accordingly, Section 2 defines the architectural and structural modelling environment used to generate the diagrid design space and the FE simulation database. Section 3 describes the input and output variables used for surrogate construction. Section 4 presents the surrogate-modelling method, including preprocessing, mathematical formulation, network architecture and training strategy. Section 5 then applies the trained surrogate to multi-objective seismic design optimisation and discusses the computational savings, error propagation and optimised candidate solutions. This organisation separates the proposed method, the supporting database and the optimisation application while preserving the logical dependency between them.

2. Architectural and structural modelling

To develop a robust and diverse parametric dataset, models were generated, incorporating vertical transformations of the diagrid geometry such as tapering, twisting and curvature. These forms, prevalent in contemporary high-rise architecture, are critical in influencing the seismic behaviour.

The design space was defined by the ten architectural parameters (X1–X10) listed in Table 1. The ten architectural parameters defined a discrete design space comprising 61,200 admissible configurations, as summarised in Table 1. This design space encompasses regular and vertically transformed geometries, including tapered, twisted and curvilinear configurations. This approach captured a wide range of geometries, from regular to complex high-rise shapes. Fig. 1 presents 20 representative models, with details provided in Table 2.

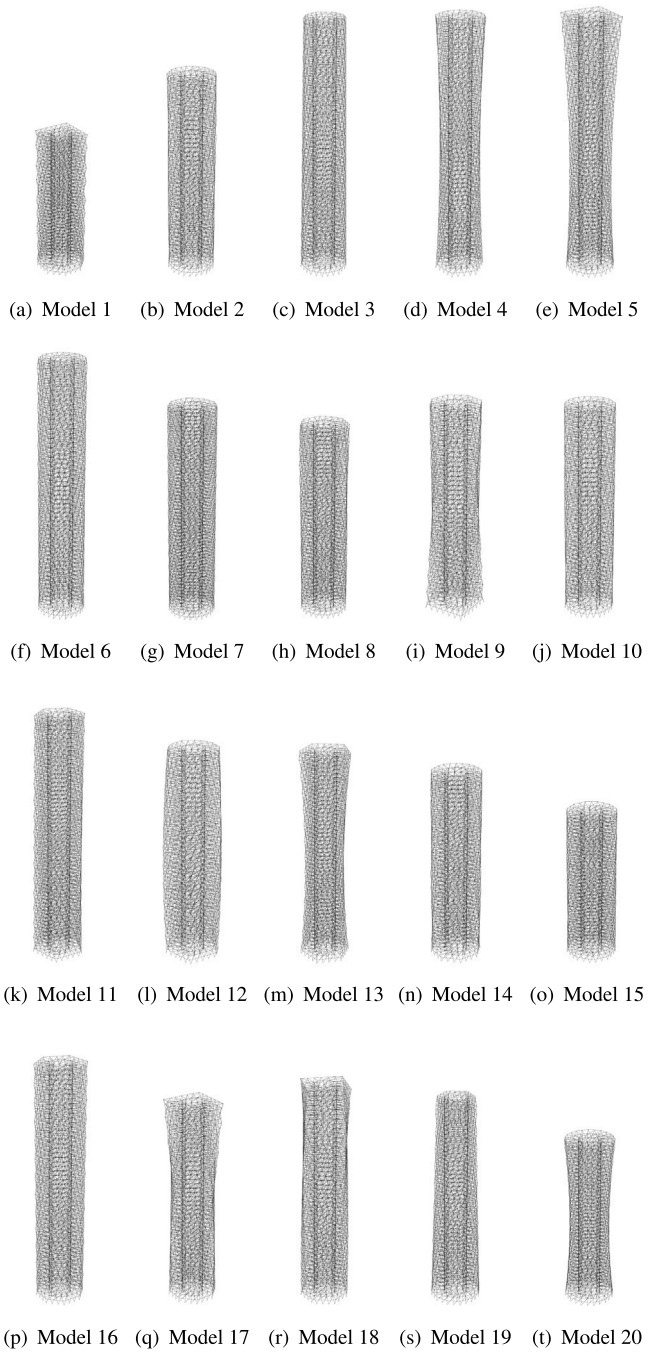
The vertical transformation method (X6) comprises three options: 0–tapering, 1–twisting and 2–curvilinear. The variables from X7 to X10 are conditionally defined based on the selected X6 type. For tapering, X7 specifies intensity, with 50% indicating a highly tapered form and 100% a prismatic one. For twisting, X8 defines the rotation angle measured between the top plan and the base, ranging from 30∘ to 90∘ rotation. For curved models, parameter X9 denotes the vertical elevation of the reference story (first, second, or third height quartile), while X10 specifies its scale, that is 0.85 for concave and 1.15 for convex.

Since considered tall buildings vary in terms of the number of stories, the core area, accommodating vertical transportation and mechanical systems, must scale with height. A dodecagonal core geometry was adopted across all models to ensure comparability of structural responses. Core dimensions were scaled using a linear relationship implemented in Grasshopper™, informed by field data: the core-to-floor area ratio ranged from 19.3% for the 140 m tall buildings to 23.8% for the 270 m tall buildings.

To reduce the computational burden, 1000 models were selected from the 61,200 configurations listed in Table 1 using Latin hypercube sampling (via the pyDOE2 Python library). The selected configurations were then processed using the Colibri plug-in, with X1 to X10 as input design parameters. Finite-element models and Python scripts storing the seismic analysis outputs and post-processing parameters were automatically generated so that, for each model, an image and a corresponding CSV record were produced.

OpenSees (McKenna et al., 2010) was used for FE dynamic simulations in the time domain, with an assumed linearly elastic response of the structures.

This modelling choice should be interpreted as a deliberate simplification consistent with the intended use of the framework at the early design stage and under moderate seismic demand. Since the training database comprises 7000 time-history simulations across 1000 parametric building configurations, adopting nonlinear analyses would have been computationally prohibitive at this scale; the linear-elastic assumption therefore represents a necessary trade-off between modelling fidelity and the feasibility of populating a sufficiently large and diverse training set for the surrogate. The resulting surrogate is therefore aimed at comparative screening and ranking of alternative building configurations, not at replacing nonlinear response-history analysis for final design, code verification or detailed performance assessment.



(a) Model 1 (b) Model 2 (c) Model 3 (d) Model 4 (e) Model 5

(f) Model 6 (g) Model 7 (h) Model 8 (i) Model 9 (j) Model 10

(k) Model 11 (l) Model 12 (m) Model 13 (n) Model 14 (o) Model 15

(p) Model 16 (q) Model 17 (r) Model 18 (s) Model 19 (t) Model 20

Fig. 1. Visualisation of 20 sample tall building models.

Nonlinear phenomena such as yielding, local buckling, connection nonlinearity, stiffness degradation and strength deterioration may modify the absolute response levels and, in some cases, alter part of the ranking among candidate solutions. For this reason, the optimised designs identified in the present study should be understood as promising preliminary candidates to be examined further through more refined nonlinear FE analyses. The present framework is therefore positioned as a design-space reduction tool: it helps identify a smaller set of candidate geometries for subsequent nonlinear verification, rather than providing final performance-assessment results by itself.

The plug-in Alpaca 4D was adapted to generate Python scripts that can be run through OpenSeesPy. The Delft Blue supercomputing cluster was leveraged to carry out 7000 simulations, given that each of the 1000 models was analysed under seven earthquake excitations, detailed in a previous work (Kazemi et al., 2025) and included in the NGA-West2 repository using the PEER online ground-motion selection tool (Pacific Earthquake Engineering Research Center, 2024). Bidirectional excitations were applied along two orthogonal horizontal directions. The use of seven ground-motion records should be interpreted in relation to the scope of the present work. The objective of this study is not to perform a complete probabilistic seismic reliability assessment or a collapse-level fragility evaluation, for which a substantially larger suite of records would be required. Rather, the aim is to construct and demonstrate a surrogate-assisted workflow for early-stage design-space exploration under moderate seismic demand. In this context, seven records provide a computationally feasible compromise between seismic variability and the need to analyse a large number of parametric building configurations. Since 1000 building models were considered, the adopted record set already led to 7000 bidirectional time-history simulations, which formed the training database for the MIMO-FNN surrogate.

From a statistical standpoint, the seven records are therefore used to represent a finite design-oriented sample of seismic excitation rather than the full stochastic variability of earthquake ground motions. This choice is consistent with common response-history practice in which seven records are often used to estimate mean engineering demand parameters in design applications, while larger ground-motion suites are typically adopted when record-to-record dispersion, collapse capacity or formal uncertainty quantification are the main objectives. To avoid overstatement, the results reported in this paper are consequently interpreted as valid within the considered design space and selected ground-motion suite. Broader generalisation to other seismic scenarios would require additional records selected from databases such as PEER NGA-West2 and constitutes a natural extension of the present work.

More specifically, the parametric tall-building definitions generated in Grasshopper™ were translated into OpenSeesPy models through the Alpaca 4D interface, ensuring that the geometry-generation stage and the FE-analysis stage remained fully connected within the automated workflow. For each sampled building configuration, the corresponding FE model was generated automatically from the architectural descriptors X1–X10, including plan shape, number of storeys, inter-storey height and vertical transformation type. Linear-elastic bidirectional time-history analyses were then performed under the selected ground-motion records, with horizontal excitations applied along two orthogonal directions.

The response extraction and post-processing stage was also automated. For each analysed model and loading case, the computed responses were collected both at the storey level and at selected reference nodes, consistently with the diaphragm assumptions adopted in the FE models. The extracted quantities included lateral displacements, velocities, accelerations, inter-storey drifts, torsional indicators, internal forces in key structural members and base reactions. These response quantities were subsequently consolidated into the simulation database used for surrogate training and optimisation.

Regarding FE-model validation, the present article focuses on the surrogate-modelling and surrogate-assisted optimisation stages rather than on introducing a new benchmark-validation campaign for the OpenSees models. The broader modelling workflow used to generate the FE database was documented in the companion modelling study, whereas the present paper uses the resulting simulation database to develop and test the MIMO-FNN surrogate. Accordingly, the optimised solutions identified here should be interpreted as preliminary candidates for design-space reduction and should be confirmed through dedicated high-fidelity FE analyses before practical adoption.

Additional indicators related to economic cost and environmental impact (i.e., carbon footprint) were obtained from the building properties; see Section 3.1.

Table 2
Specification of selected 20 sample tall building models.

Model	X1	X2	X3	X4	X5	X6	X7	X8	X9	X10	Total area (m^2)	Total façade area (m^2)	Total weight (tons)
927	4	8	1	40	3.5	0	0.80	–	–	–	52,926	19,085	13,100
631	9	12	1	44	4.5	1	–	75	–	–	67,431	27,479	18,900
683	9	9	0	56	4.5	0	0.70	–	–	–	71,252	32,447	21,900
515	10	6	1	56	4.5	2	–	–	0.50	0.85	83,472	35,286	23,200
366	4	10	1	60	4.0	2	–	–	0.25	0.85	88,664	34,088	24,900
656	8	11	1	56	4.5	0	0.90	–	–	–	81,064	34,142	25,800
657	9	11	1	60	3.5	0	1	–	–	–	91,462	29,164	24,800
530	7	12	0	48	4.0	1	–	60	–	–	72,038	26,788	19,400
322	8	4	1	44	4.5	1	–	30	–	–	64,632	28,565	21,000
71	11	10	1	44	4.5	1	–	45	–	–	66,003	27,466	16,300
891	6	5	0	60	4.0	0	0.90	–	–	–	85,783	33,527	25,300
372	10	5	0	44	4.5	2	–	–	0.50	1.15	66,124	28,206	22,900
768	5	6	0	52	4.0	2	–	–	0.75	0.85	78,226	29,812	16,200
457	10	8	1	44	4.0	1	–	75	–	–	65,958	24,519	16,300
248	12	9	1	40	3.5	1	–	30	–	–	59,937	19,430	13,000
612	6	8	0	52	4.5	0	1	–	–	–	79,593	33,004	23,100
706	4	8	1	44	4.5	1	–	90	–	–	64,632	28,565	22,800
979	4	7	0	48	4.5	1	–	45	–	–	73,916	31,400	22,900
299	6	12	1	48	4.0	0	0.50	–	–	–	52,810	22,877	11,500
102	11	12	0	44	3.5	2	–	–	0.75	0.85	66,008	21,326	9920

Table 3
Carbon coefficient factors for steel and concrete types.

ID	Material	Source	WF	A1–A3	A4	A5	C2–C4	Total carbon coefficient
1	Structural steel "typical", global average (Gibbons et al., 2022)	ICE	0.01	1.55	0.005	0.0175	0.018	1.5905
2	Recycled structural steel sections, EA furnace (Beltrame SPA) (Beltrame Group, 2020)	EPD	0.01	0.3693	0.005	0.0057	0.018	0.398
3	Reused structural steel sections (EMR Group) (EMR, 2022)	EPD	0.01	0.0466	0.005	0.0025	0.018	0.0721
4	RC30 reinforced concrete CEM I, 0% GGBS (Gibbons et al., 2022)	ICE	0.053	0.129	0.005	0.0081	0.018	0.1601
5	C28/35 CIIIB ready mix concrete, 12% GGBS (Hanson UK) (Hanson UK, 2017)	EPD	0.053	0.0696	0.005	0.0049	0.018	0.0975

3. Dataset: inputs and outputs

The dataset was constructed to support surrogate modelling and was obtained from 1000 parametric diagrid tall-building configurations analysed under seven bidirectional ground-motion records, resulting in 7000 analysed cases. To avoid ambiguity, three levels of data representation are distinguished in this study: (i) the raw input variables defining each analysed case; (ii) the preprocessed input columns supplied to the neural network; and (iii) the output responses predicted by the surrogate model.

The raw input vector contains 22 variables. These are grouped into two categories: (1) ten architectural and geometric design parameters X_1–X_{10}, defining the building configuration; and (2) twelve record-specific ground-motion descriptors obtained from the PEER database after excluding redundant metadata. The ground-motion descriptors were included as explicit input variables of the surrogate model. Therefore, the seven records were not treated merely as repeated loading cases or as simple record labels, but as seismic inputs characterised by their own intensity- and record-related descriptors. This allowed the MIMO-FNN to learn response variations associated with the selected ground-motion characteristics, although the model should not be extrapolated beyond the range covered by the adopted record suite without further validation.

Before training, the 22 raw input variables were transformed into 45 preprocessed input columns through numerical scaling, categorical encoding and conditional masking of transformation-dependent geometric variables. This preprocessing step is described in Section 4. The distinction between raw variables and preprocessed columns is important because the neural network does not receive only the original X_1–X_{10} and ground-motion descriptors, but their encoded and scaled representation.

The output vector contains 91 response quantities. These are grouped into three categories: (1) 14 environmental and cost-related responses, quantifying sustainability and economic indicators; (2) 69 structural responses, referring to seismic performance quantities extracted from the OpenSeesPy time-history analyses, see Section 3.2; and (3) 8 geometry-dependent responses, related to form-based characteristics independent of the dynamic analysis, see Section 3.3.

A Python-based automation pipeline was developed to extract and organise the response quantities for each analysed model (md_1 to md_{1000}) and each ground motion (gm_1 to gm_7). The final surrogate-modelling dataset comprises 7000 analysed cases, 22 raw input variables, 45 preprocessed input columns and 91 output responses.

Tables 4 and 5 summarise the input and output groups used in the surrogate-modelling dataset, while the complete machine-readable lists of variable names and descriptions are provided in the supplementary files Features.xlsx and Responses.xlsx.

To further support reproducibility and transparency, the repository associated with the surrogate-modelling workflow is available at: https://github.com/pkazemis-a11y/tb_surrogate. The repository includes the accessible computational material related to data preprocessing, surrogate-model implementation and optimisation workflow, and is intended to help readers reproduce the main data-driven steps of the framework.

3.1. Environmental and cost responses

Environmental impact was assessed through embodied-carbon indicators and cost-related quantities, accounting for material production, transportation, construction processes and end-of-life stages. The embodied carbon was expressed in kg of CO_2 equivalent ($kgCO_2e$) and was computed by multiplying the relevant material quantity by the corresponding embodied-carbon coefficient, following the data reported by Kanyilmaz et al. (2023).

Life-cycle stages A1–A3 (production), A4–A5 (construction) and C2–C4 (end-of-life) were included to define the carbon coefficients summarised in Table 3.

For structural components, the embodied carbon $E_{\text{structure}}$ was computed from the structural mass and the corresponding mass-based embodied-carbon coefficient. In particular, $W_{\text{structure}}$ denotes the

Table 4
Raw input-variable groups used to define each analysed case of the surrogate-modelling dataset.

Group	Variable	Description	Preprocessing treatment
Building design	X_1	Number of sides of the top plan	Encoded as a discrete geometric descriptor
Building design	X_2	Number of sides of the bottom plan	Encoded as a discrete geometric descriptor
Building design	X_3	Top–bottom plan orientation	Encoded as a binary/categorical design descriptor
Building design	X_4	Total number of storeys	Scaled numerical variable
Building design	X_5	Inter-storey height	Scaled numerical variable
Building design	X_6	Vertical transformation type: tapered, twisted or curvilinear	Categorical encoding
Building design	X_7	Tapering ratio, active only when $X_6 = 0$	Scaled and paired with a validity mask
Building design	X_8	Twisting angle, active only when $X_6 = 1$	Scaled and paired with a validity mask
Building design	X_9	Vertical location of curvilinearity, active only when $X_6 = 2$	Scaled and paired with a validity mask
Building design	X_{10}	Curvilinearity scale factor, active only when $X_6 = 2$	Scaled and paired with a validity mask
Ground motion	G_1–G_{12}	Twelve record-specific ground-motion descriptors retained from the PEER metadata/intensity-measure set after excluding redundant metadata. These descriptors characterise the seismic input numerically rather than through a simple record label.	Scaled numerical variables
Total	**22 raw variables**	**Transformed into 45 preprocessed input columns before MIMO-FNN training**	

Table 5
Output-response groups predicted by the MIMO-FNN surrogate.

Output group	Number	Physical meaning	Examples
Structural responses	69	Global maximum response quantities extracted from OpenSeesPy time-history analyses, including translational, drift, stress, torsional, reaction-force and internal-force indicators.	Maximum acceleration, maximum displacement, inter-storey drift, von Mises stress, torsional response, base reactions and member-force indicators
Environmental and cost responses	14	Quantities related to embodied carbon, material use, construction cost and normalised sustainability/cost indicators.	Embodied carbon indicators, total cost and cost per unit area
Geometry-dependent responses	8	Quantities derived from the generated parametric geometry and structural sizing information, independent of the dynamic time-history analysis.	Total area, façade area, aspect ratio, building height, diagrid angle indicators and total mass
Total	**91**	**Complete multi-output response vector used for surrogate training and prediction**	

Table 6
Data related to cost analysis.

Item	Cost	Unit
Steel material UMC (Janisch, 2023)	1.5	€/ kg
Steel labour ULC (Janisch, 2023)	0.9	€/ kg
Precast RC slab material (Rete Ferroviaria Italiana, 2023)	50	€/ m^2
Precast RC slab labour (Rete Ferroviaria Italiana, 2023)	25	€/ m^2
Fixed joints coefficient (Simões, 1996)	1.6	–
Pinned joints coefficient (Simões, 1996)	1.2	–

structural weight obtained from the structural model, and $W_{\text{structure}}/g$ is the corresponding structural mass in kg. The coefficient $e_{\text{structure}}$ is the embodied-carbon factor in kgCO_2e/kg, as reported in Table 3. Therefore, the resulting embodied carbon has units of kgCO_2e:

$$E_{\text{structure}} = \frac{W_{\text{structure}}}{g}\, e_{\text{structure}}. \tag{1}$$

Area-normalised embodied-carbon indicators, expressed in kgCO_2e/m^2, were computed separately by dividing the relevant embodied-carbon quantities by the corresponding gross internal area. These area-normalised quantities were used only as sustainability indicators and not as the coefficient in the mass-based calculation above.

Similarly, the embodied carbon for flooring systems E_{floor} was calculated based on the floor weight, accounted for in terms of the floor area. The analysis considered multiple material options, including typical steel, recycled steel, reused steel, typical concrete and sustainable concrete variants. The calculations spanned across six data sets to cover all combinations of materials and six other data sets with a special focus on the embodied carbon per gross internal area in kgCO_2e/m^2.

The financial assessment covered structural components, joint assemblies and flooring systems, based on data in Table 6. Costs for concrete floors and the steel structural framework, including the core elements, diagrids and radial beams, were accounted for. The structural cost SC was computed from the material weight W_{mat} multiplied by the sum of unitary material and labour costs $UMC_{\text{mat}} + ULC_{\text{mat}}$:

$$SC = W_{\text{mat}} \times (UMC_{\text{mat}} + ULC_{\text{mat}}) \tag{2}$$

The joint cost JC was calculated as a fraction of the structural cost, using coefficients of 0.6 and 0.2 for fixed and pinned joints, respectively:

$$JC = SC \times \left(\frac{\text{Fixed joints}}{\text{Total joints}} \times 0.6 + \frac{\text{Pinned joints}}{\text{Total joints}} \times 0.2 \right). \tag{3}$$

The coefficients 0.6 and 0.2 represent incremental joint-cost fractions relative to the structural cost SC. Consequently, all-fixed and all-pinned cases correspond to total structural-plus-joint cost multipliers of 1.6 and 1.2, respectively.

The floor cost FC was computed based on the floor area, multiplied by the sum of unitary material and labour costs:

$$FC = \text{Area} \times (UMC_{\text{floor}} + ULC_{\text{floor}}) \tag{4}$$

The final dataset included data relevant to the total cost and total cost per unit area.

3.2. Structural responses

The structural responses, derived during post-processing, formed the core outcomes of this study. A total of 69 structural response metrics were extracted, focusing on global maxima across simulations. These included acceleration, displacement, inter-storey drift, stress levels in key structural components (e.g., diagrids), torsional effects and reaction forces.

3.3. Geometry-dependent responses

The eight geometry-dependent quantities, unlike the structural responses, are derived directly from the generated parametric geometry

and structural sizing information (e.g., total area, façade area and total mass). Their non-simulation origin justifies their separate categorisation. These quantities were computed for each building configuration and retained alongside the design parameters X_1–X_{10}. Although the dataset included full time-series, interval-based and per-floor maxima, only the global maxima were retained for surrogate modelling, due to their relevance for performance assessment. This choice simplified the model while preserving the potential for future granular analyses.

4. Surrogate modelling

4.1. Data preprocessing

Preprocessing is a critical phase in model development, given the sensitivity of ML algorithms to data scale and distribution. Its primary goal is to improve model performance by standardising input formats and magnitudes.

Particular emphasis was given to distinguishing between numerical and categorical features, based on data type and also on their relevance to the predictive task and model behaviour. Numerical features are measurable quantities on continuous or discrete scales, often with linear or monotonic effects on outputs. In contrast, categorical features represent distinct classes without a continuum. Some design variables, although represented by numerical values, describe discrete geometric choices rather than continuous physical quantities. For this reason, variables such as the number of sides of the top and bottom plans were treated as discrete design descriptors during preprocessing. This treatment was adopted to preserve the parametric meaning of the design space, not because of an assumed limitation of the feed-forward neural network in learning nonlinear dependencies. Similarly, plan orientation and vertical transformation method were treated as categorical design descriptors because they define distinct geometric classes rather than continuously varying physical quantities. The ground motions, however, were represented through record-specific descriptors rather than by a simple record label. Thus, the surrogate does not use a ground-motion ID as a substitute for the physical characteristics of the excitation. Nevertheless, the model should not be extrapolated to ground motions whose descriptors fall outside the range covered by the adopted record suite without additional validation. Floor-to-floor height, despite having few discrete options, was instead considered numerical due to its measurable influence on the structural responses.

Ten preprocessing strategies were applied and are collected in the Supplementary file PreprocessingStrategies.pdf, where a brief description of each algorithm is provided. Among them, the MinMaxScaler achieved the lowest root mean square error (RMSE, a metric that quantifies the discrepancy between a model's predicted values and the corresponding observed values (Sharma et al., 2022)), making it the optimal choice for both inputs and outputs. Its effectiveness stems from rescaling features to the [0, 1] range, aiding model convergence by maintaining consistent variance within the dataset. The MinMaxScaler's compatibility with NNs, due to its ability to place features on a common bounded scale and stabilise gradient-based training, also led to more efficient training and improved convergence.

To avoid data leakage, all preprocessing transformations were fitted only on the training subset within each split. In particular, the MinMaxScaler was not fitted on the full dataset before the train–validation–test division. For each cross-validation fold, the minimum and maximum values used by the scaler were computed exclusively from the training subset, and the fitted scaler was then applied unchanged to the corresponding validation and test subsets. Consequently, validation and test samples did not contribute to the scaling bounds used during model training or model selection.

To make the preprocessing workflow self-contained, Table 7 summarises the main preprocessing operations used before MIMO-FNN training. The supplementary file *PreprocessingStrategies.pdf* provides additional implementation details, while the table reports the essential operations required to understand the data transformation pipeline.

Table 7
Summary of the main preprocessing operations used in the surrogate-modelling workflow.

Operation	Role in the present study
Numerical scaling	Applied to numerical building variables, ground-motion descriptors and continuous outputs to improve neural-network convergence. Scaling parameters were fitted only on the training subset within each split.
Categorical encoding	Applied to discrete design descriptors such as plan orientation and vertical transformation type, so that distinct geometric classes could be represented without imposing artificial continuous distances.
Conditional masking	Applied to X_7–X_{10}, which are meaningful only for specific transformation types. Each conditional variable was paired with a binary validity indicator.
Output scaling	Applied to the 91 output responses to balance quantities with different physical units and numerical ranges during multi-output regression.
Fold-local preprocessing	For each cross-validation fold, preprocessing parameters were estimated only from the training subset and then applied to the corresponding validation and test subsets to avoid data leakage.

It is worth mentioning that, in this preprocessing phase, categorical features were transformed via one-hot encoding into binary data sets, enabling ML compatibility. The preprocessing pipeline also addressed the conditional nature of the geometric variables X7–X10. These variables are not missing values in the usual statistical sense; rather, they are structurally non-applicable for some transformation classes. Specifically, X7 is meaningful only for tapered geometries, X8 only for twisted geometries, and X9–X10 only for curvilinear geometries. Therefore, assigning a fixed dummy value to non-applicable entries would risk introducing artificial numerical states into the training data, while treating all variables as ordinary numerical inputs would obscure the physical logic of the parametric design space.

To avoid this issue, a conditional masking strategy was adopted. Each conditional variable was paired with a binary validity indicator, where 1 denotes that the corresponding geometric variable is active and physically meaningful for the selected transformation type, while 0 denotes that the variable is inactive because of the design logic. In this way, the MIMO-FNN receives both the numerical value of an active parameter and an explicit indication of whether that parameter is applicable. This representation prevents the network from interpreting non-applicable placeholders as real geometric values. The adopted strategy can be interpreted as a physically informed encoding of the conditional design space. A simple dummy-value strategy would replace inactive entries with arbitrary numerical constants, but the network would not be able to distinguish whether such constants represent actual geometric values or placeholders. One-hot encoding of the transformation class X6 would identify whether the model is tapered, twisted or curvilinear, but it would still not explicitly flag the validity of X7, X8, X9 and X10 at the variable level. The proposed masking strategy combines these two requirements: it retains the transformation-class information and separately identifies the active/inactive status of each conditional geometric variable. This is especially important in the present multi-output problem, where the same input representation is used to predict 91 structural, environmental and geometry-dependent responses. The comparison among these preprocessing alternatives is summarised in Table 8.

This comparison clarifies that the role of the masking strategy is not merely to increase the number of input columns, but to preserve the conditional semantics of the design variables. For instance, a value assigned to X8 has a physical interpretation only when the selected transformation is twisting; the same numerical entry is meaningless for tapered or curvilinear models. Similarly, X9 and X10 describe the location and intensity of curvilinearity and are not physically defined

Table 8
Comparison of alternative preprocessing strategies for the conditional geometric variables X7–X10.

Preprocessing strategy	Transformation class encoded	Conditional variable value retained	Applicability explicitly encoded	Risk of spurious numerical interpretation
Dummy-value encoding	Partially	Yes, after imputation	No	High
One-hot encoding of X6 only	Yes	Yes, after imputation	No, not at variable level	Moderate
Proposed conditional masking	Yes	Yes	Yes	Low

for tapered or twisted configurations. The binary masks therefore separate non-applicability imposed by design logic from genuine numerical values. This distinction improves the consistency of the input representation used by the surrogate model and reduces the possibility of learning artificial relationships caused by arbitrary placeholders.

The preprocessing steps discussed above transformed the original 22 input variables into 45 input columns after categorical encoding, numerical scaling and the introduction of conditional validity masks. The increased dimensionality is therefore a consequence of preserving both the value and the applicability status of the conditional geometric variables.

4.2. Mathematical formulation of the MIMO-FNN surrogate

After preprocessing, each analysed case is represented by an input vector

$$\boldsymbol{x}_i = \left[\boldsymbol{x}_i^b,\ \boldsymbol{x}_i^g\right] \in \mathbb{R}^{45}, \tag{5}$$

where $\boldsymbol{x}_i^b$ contains the preprocessed building-design variables and $\boldsymbol{x}_i^g$ contains the ground-motion descriptors. The corresponding output vector is

$$\boldsymbol{y}_i = \left[\boldsymbol{y}_i^s,\ \boldsymbol{y}_i^e,\ \boldsymbol{y}_i^q\right] \in \mathbb{R}^{91}, \tag{6}$$

where $\boldsymbol{y}_i^s \in \mathbb{R}^{69}$ denotes the structural responses, $\boldsymbol{y}_i^e \in \mathbb{R}^{14}$ denotes the environmental and cost-related responses, and $\boldsymbol{y}_i^q \in \mathbb{R}^8$ denotes the geometry-dependent responses. The surrogate model is therefore formulated as the nonlinear multi-input, multi-output mapping

$$\hat{\boldsymbol{y}}_i = f_{\boldsymbol{\theta}}(\boldsymbol{x}_i), \qquad f_{\boldsymbol{\theta}} : \mathbb{R}^{45} \to \mathbb{R}^{91}, \tag{7}$$

where $f_{\boldsymbol{\theta}}$ is the MIMO-FNN and $\boldsymbol{\theta}$ denotes its trainable weights and biases. The two input groups are first processed through separate dense branches,

$$\boldsymbol{h}_i^b = \phi_b(\boldsymbol{x}_i^b; \boldsymbol{\theta}_b), \qquad \boldsymbol{h}_i^g = \phi_g(\boldsymbol{x}_i^g; \boldsymbol{\theta}_g), \tag{8}$$

and the resulting latent representations are concatenated as

$$\boldsymbol{h}_i = \left[\boldsymbol{h}_i^b,\ \boldsymbol{h}_i^g\right]. \tag{9}$$

The shared dense layers then map $\boldsymbol{h}_i$ to the final multi-response prediction,

$$\hat{\boldsymbol{y}}_i = \psi(\boldsymbol{h}_i; \boldsymbol{\theta}_\psi). \tag{10}$$

The model parameters were obtained by minimising the root mean square error over the training set,

$$\mathcal{L}_{\mathrm{RMSE}}(\boldsymbol{\theta}) = \sqrt{\frac{1}{n\,m}\sum_{i=1}^{n}\sum_{j=1}^{m}\left(\hat{y}_{ij} - y_{ij}\right)^2}, \qquad m = 91, \tag{11}$$

where n is the number of training samples and m is the number of output responses. The training problem can therefore be written as

$$\boldsymbol{\theta}^* = \arg\min_{\boldsymbol{\theta}} \mathcal{L}_{\mathrm{RMSE}}(\boldsymbol{\theta}) + \lambda\|\boldsymbol{\theta}\|_2^2, \tag{12}$$

where the second term denotes L2 regularisation through weight decay. This formulation makes explicit that the proposed surrogate is a high-dimensional multi-response regression operator trained to replace repeated FE-based response evaluations during early-stage design exploration.

The computational contribution of this formulation is not the isolated use of a feed-forward neural network, but the construction of a reproducible surrogate operator that maps a conditionally defined diagrid design space and record-specific ground-motion descriptors to a large set of structural, environmental and geometry-dependent responses. This enables multi-objective optimisation to be performed through inexpensive neural-network inference rather than repeated FE time-history analyses.

4.3. Multi input multi output feed-forward neural network training

A MIMO-FNN was specifically designed for the simultaneous prediction of multiple response quantities from diverse input parameters. A multi-input layer architecture was adopted to separately process independent input sets, such as building characteristics and ground motion descriptors, allowing for a more accurate representation of each effect.

Similarly, the output layer was structured into separate branches for structural, environmental and geometry-dependent responses. Within the structural outputs, further distinctions were applied (e.g., inter-storey drift vs. acceleration), enabling tailored prediction paths for distinct response types, as detailed in the supplementary file Responses.xlsx.

MIMO-FNNs are particularly effective at capturing complex relationships among a broad range of response categories, with emphasis on those produced by seismic structural FE simulations. By adjusting weights through backpropagation via optimisers like Adam, see (Ilboudo et al., 2020; Lydia & Francis, 2019), these networks learn nuanced patterns within the dataset. Their architecture enables simultaneous prediction of multiple outputs, effectively solving multiple regression tasks and linking input features to each response group.

While MIMO-FNNs effectively capture statistical patterns, they lack an inherent understanding of the physical laws governing structural dynamics. Their predictions are based on data fidelity and are limited by the representativeness of the training set. Networks trained on a broad spectrum of seismic scenarios and structural responses can generalize effectively. However, they can fail to maintain the same accuracy in cases where scenarios differ markedly from those in the training dataset because, as mentioned above, the physics is not encoded.

For this reason, the physical credibility of the present surrogate should not be interpreted in the same way as that of a physics-informed neural network. In the current framework, mechanical consistency is inherited indirectly from the FE-generated training data and from the explicit treatment of admissible geometric conditions through encoding and masking, rather than from constraint-based loss terms or embedded governing equations. The model is therefore intended as a data-driven predictor for rapid design-space exploration. Its outputs are useful for screening and prioritising alternatives, but candidate solutions identified by the optimiser should still be re-analysed with high-fidelity FE models before being adopted in practical engineering workflows. This clarification is important to avoid overstating the reliability of the surrogate beyond the scope supported by the present study.

Two levels of consistency control are nevertheless embedded in the workflow. First, the input design variables are constrained to the admissible discrete ranges of the parametric diagrid design space, and transformation-dependent variables are activated only when physically

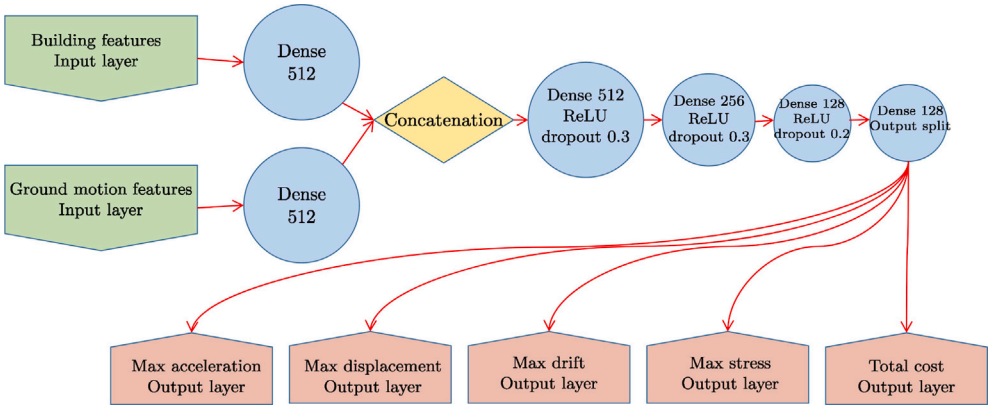


Fig. 2. Scheme of the MIMO-FNN.

meaningful. For example, tapering intensity is used only for tapered geometries, twisting angle only for twisted geometries, and curvilinearity parameters only for curvilinear configurations. Second, the surrogate is trained exclusively on responses generated from mechanics-based FE time-history analyses, so that the learned mapping remains tied to a structural simulation database rather than to arbitrary response labels. These checks do not replace explicit physics-informed learning, but they reduce the risk of physically inconsistent input representations and support the use of the model as a preliminary screening tool.

4.4. Architecture specifications and hyperparameter optimisation

To improve the reproducibility of the proposed surrogate model, the architecture of the adopted MIMO-FNN is described here explicitly. After preprocessing, the original 22 input variables were transformed into 45 input columns through numerical scaling, one-hot encoding of categorical variables and conditional masking of context-dependent geometric variables. These preprocessed inputs were organised into two input streams: one stream corresponding to the building-design variables and one stream corresponding to the ground-motion descriptors. This separation was adopted to allow the network to learn design-related and excitation-related feature representations before their interaction was modelled in the shared latent space.

The two input streams were first passed through separate fully connected input branches and were then concatenated. The concatenated feature vector was subsequently processed by a pyramidal feed-forward architecture composed of four hidden dense layers with 512, 256, 128 and 128 neurons, respectively. Rectified linear unit (ReLU) activation functions were used in all hidden layers because of their computational efficiency and their ability to mitigate the vanishing-gradient problem in deep regression networks. Dropout regularisation was applied after the first three hidden layers, with dropout rates equal to 0.3, 0.3 and 0.2, respectively. The final layer was divided into output branches corresponding to the considered response groups, namely structural, environmental and geometry-dependent responses. Since the problem is formulated as a continuous multi-output regression task, linear activation functions were adopted in the output layers.

The resulting MIMO-FNN therefore maps the preprocessed input vector to 91 output responses, including 69 structural responses, 14 environmental and cost-related responses and 8 geometry-dependent quantities. The model was implemented in PyTorch and trained using the Adam optimiser. The RMSE was used as the loss function, consistent with the objective of minimising prediction errors in the scaled response space. The adopted learning rate was 2×10^{-5}, the batch size was 32, the weight-decay coefficient was 1×10^{-4}, and the training was carried out for 100 epochs. The schematic representation of the adopted MIMO-FNN is shown in Fig. 2, while the main architecture and training hyperparameters are summarised in Table 9.

Table 9
Architecture and training hyperparameters of the adopted MIMO-FNN surrogate model.

Item	Specification
Original input variables	22
Preprocessed input columns	45
Input grouping	Building-design variables and ground-motion descriptors
Output variables	91 responses
Output grouping	69 structural responses, 14 environmental and cost-related responses and 8 geometry-dependent responses
Network type	Multi-input, multi-output feed-forward neural network
Input streams	Two parallel input streams processed through separate dense layers before concatenation
Shared hidden architecture	Pyramidal fully connected architecture
Number of shared hidden layers	4
Hidden-layer widths	512, 256, 128 and 128 neurons
Hidden-layer activation function	ReLU
Output-layer activation function	Linear activation for continuous regression outputs
Dropout rates	0.3 after the first hidden layer, 0.3 after the second hidden layer and 0.2 after the third hidden layer
Loss function	RMSE
Optimiser	Adam
Learning rate	2×10^{-5}
Batch size	32
Weight decay / L2 regularisation	1×10^{-4}
Number of epochs	100
Cross-validation	2-fold to 5-fold cross-validation examined
Implementation framework	PyTorch
Input/output scaling	MinMaxScaler for both inputs and outputs

The 7000 analysed cases were divided using a grouped strategy at the building-configuration level. All seven ground-motion cases associated with the same building configuration were assigned to the same subset. Therefore, if a building configuration was included in the test subset, none of its seven ground-motion analyses appeared in the training or validation subsets. This grouped splitting strategy avoids information leakage associated with repeated analyses of the same geometry and provides a more meaningful assessment of the surrogate's ability to predict responses for unseen building configurations.

MSE, RMSE and custom weighted loss functions were tested as loss functions: they all failed to resolve the issue related to inflated loss values until an appropriate scaling of input and output was adopted, along with further optimisation focused on (Brunton & Kutz, 2024): i) tuning of

learning rate, batch size and optimiser choice (among stochastic gradient descent, RMSprop and Adam); ii) layer size and depth, dropout rates and activation functions (among LeakyReLU, ELU and ReLU); iii) regularisation techniques, such as L1/L2 ones, to curb overfitting (Tibshirani et al., 2010), supported by feature engineering, k-fold cross-validation and early stopping refinements.

Model architecture had the highest impact, as changes in layer size, depth and activation functions strongly affected the model capacity to capture complex relationships. Learning rate and optimiser had a medium-to-high impact, influencing convergence speed and training stability. Regularisation and dropout techniques had a moderate influence by improving generalisability and reducing overfitting. Batch size and number of epochs had a medium-to-low impact, affecting learning dynamics and exposure to data. Lastly, k-fold cross-validation had low impact, primarily contributing to evaluation robustness rather than predictive performance.

By addressing data scaling first and key hyperparameter tuning next, major training challenges were resolved, and model performance was significantly boosted. Consequently, the use of automated hyperparameter optimisation methods, like Bayesian optimisation, random search, or grid search, was then deemed unnecessary, as the NN configuration already achieved low RMSE with a remarkable balance of computational efficiency and accuracy. The learning rate value of 2×10^{-5} was selected through extensive testing to balance convergence speed and stability, preventing overshooting during optimisation. A weight decay of 1×10^{-4} incorporated L2 regularisation, which mitigated overfitting by penalising large weights and promoting model simplicity. A batch size of 32 was selected to balance computational efficiency and learning stability; it proved to be small enough for frequent updates and convergence, yet large enough to leverage vectorised operations without excessive noise. The training was conducted over 100 epochs, ensuring ample opportunity for the model to adjust its weights and biases towards minimising the loss function.

The MIMO-FNN was developed within the PyTorch framework, here chosen for its dynamic computation graphs and efficient tensor operations. GPU acceleration further enabled parallel processing, significantly reducing computation time and supporting faster model iteration over large datasets. The related implementation material is also made available through the project repository: https://github.com/pkazemis-a11y/tb_surrogate. This additional material complements the architectural description and hyperparameter summary in Table 9, thereby improving the reproducibility of the MIMO-FNN surrogate.

The choice of the optimiser played a crucial role in optimising NN training. Adam was employed for its adaptive learning rates and momentum-based updates, which enhanced convergence and managed sparse gradients effectively. It also consistently outperformed the other tested optimisers in terms of training stability and efficiency.

Fig. 3 shows the training and validation loss curves of the MIMO-FNN. Ideally, both curves decline and converge to low values, indicating effective learning and generalisation. Increasing the cross-validation folds from 2 to 5 aimed to examine the stability of the training procedure; the absence of significant performance gains indicated that the selected MIMO-FNN configuration was stable within the considered dataset. This statement refers to the internal predictive stability of the surrogate over the available 7000 simulations and should not be interpreted as a claim of global seismic robustness with respect to all possible ground-motion scenarios.

The RMSE trends show consistent declines in both training and validation errors, indicating effective learning within the sampled design and ground-motion space. The low final test RMSE and the absence of overfitting or underfitting support the predictive accuracy of the surrogate for the available simulation database. Stable performance across cross-validation folds further demonstrates the internal consistency of the MIMO-FNN, while broader seismic generalisation would require additional validation using a larger set of ground-motion records.

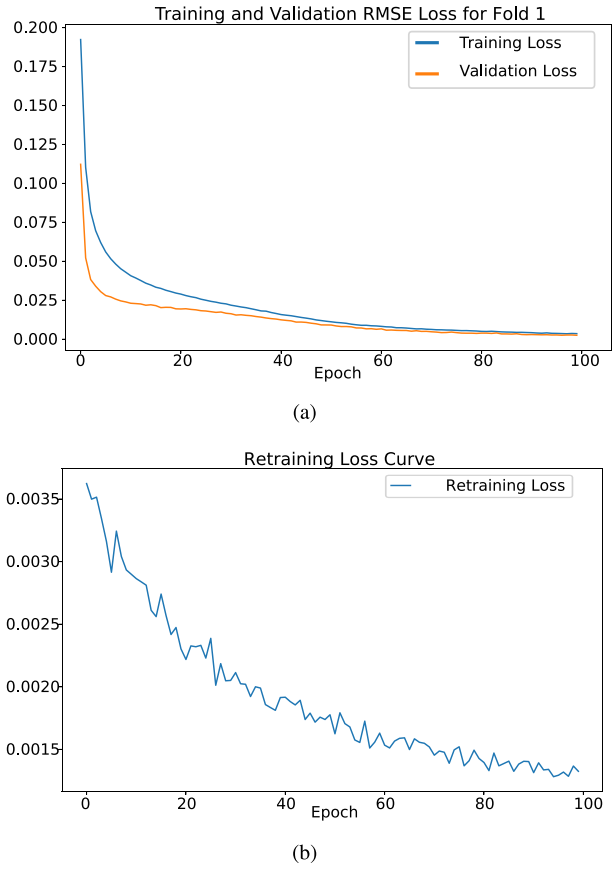


Fig. 3. MIMO-FNN: Training-validation loss for (a) an exemplary fold and (b) the full dataset.

5. Exploring design space through optimisation for generating new solutions

The optimisation stage was designed as a demonstrative early-stage design-space screening exercise. It should not be interpreted as a complete seismic-resilience assessment. The optimisation was performed for one selected ground-motion descriptor set in order to demonstrate how the trained MIMO-FNN can be coupled with NSGA-II to generate candidate geometries at low marginal computational cost. A full resilience-oriented design procedure would require optimisation and verification under several unseen ground-motion records and, preferably, nonlinear response-history analyses.

This section applies optimisation techniques to the surrogate model described earlier to generate new candidate design solutions. Similar to generative AI, the approach explores structurally efficient regions of the design space by evaluating surrogate-predicted responses. Since these responses are predicted by the MIMO-FNN rather than recomputed through new FE time-history analyses, the optimisation results should be interpreted as surrogate-screened preliminary candidates for early-stage design-space reduction. This procedure significantly reduces the computational cost of design exploration, while the final engineering acceptance of any selected candidate still requires subsequent high-fidelity FE verification and code-based assessment.

To demonstrate the feasibility of integrating multi-objective optimisation within the proposed framework, the surrogate model was coupled with the optimisation tool to effectively generate candidate design solutions. The term “candidate” is used deliberately: the NSGA-II procedure identifies promising regions of the design space according to the MIMO-FNN predictions, but it does not by itself constitute a final FE verification of the optimised geometries. Given the exploratory nature of the optimisation landscape, the emphasis is on methodological

applicability and rapid screening rather than on exhaustive design certification.

This study employs genetic algorithm-based methods, namely NSGA-II. The procedure is initialised with a diverse population whose fitness scores are evaluated. Individuals are ranked into Pareto fronts based on non-domination and crowding distance, to maintain solution diversity. Parent selection is performed via tournament selection, followed by crossover and random mutation to generate offspring. The combined parent–offspring population is then sorted and the best individuals are selected for the next generation, preserving both optimality and diversity in the solution process.

The complete discrete design space contains 61,200 admissible configurations. Once the surrogate is trained, exhaustive neural-network evaluation of this finite set would be computationally feasible. However, direct FE analysis of the complete space under the seven ground-motion records would require $61{,}200 \times 7 = 428{,}400$ bidirectional time-history simulations. The computational bottleneck is therefore the generation of high-fidelity FE data, not the subsequent surrogate evaluation.

The use of NSGA-II in this study is intended to demonstrate a scalable surrogate-assisted multi-objective optimisation workflow. Exhaustive surrogate screening of the finite discrete design space would also be possible and complementary. However, the evolutionary framework is more readily extendable to larger, continuous or more highly constrained design spaces, additional objectives, additional design variables and multi-scenario optimisation.

In a conventional FE-based structural optimisation loop, each candidate geometry typically requires model generation, dynamic analysis and response post-processing before its objective values can be evaluated. In the present workflow, the trained MIMO-FNN provides low-cost predictions of the relevant response quantities during the optimisation stage, while NSGA-II operates on the selected surrogate-predicted objectives to retain diverse trade-off solutions within the discrete design space defined by X1–X10. The two components therefore have distinct roles: the surrogate supports efficient candidate evaluation, whereas the evolutionary search explores competing objectives. This arrangement enables early-stage screening of multiple candidate geometries against several performance and material-demand criteria. The complete optimisation procedure is summarised in Fig. 4.

With the aim of showing a consistent example of the optimisation procedure, each individual in the population was evaluated across five distinct objectives, enabling trade-off solutions among competing criteria. Hence, from the initial set of 91 structural and performance responses, 10 were shortlisted: ‘Overall Max Acc’, ‘Max Displacement’, ‘Overall Max Drift’, ‘Total Max Von Mises tot’, ‘Overall Max Torsion’, ‘Total Max Magnitude R’, ‘Total Max Magnitude M’, ‘Total Mass’, ‘Total costs/Total gross area (TGA)’ and ‘EC GIA (Steel3, floor 2)’. This selection was informed by prior studies and further refined through correlation analysis. Using the Spearman correlation metric, the five representative responses finally retained for the optimisation example were identified as: Total Mass, Total Max Von Mises tot, Max Displacement, Overall Max Drift and Overall Max Acc. Overall Max Torsion was considered during the preliminary response-screening stage, but it was not retained among the final five objectives used in the optimisation results reported in Tables 11 and 12.

Ground motion features were selected as a distinct subset, while the remaining features were assigned to building design. Unique values among building features were extracted to ensure diversity in the initial population. The initial solution was selected from the pool of models used to train the MIMO-FNN. Because the objective of this section is to demonstrate the surrogate-assisted optimisation workflow, the optimisation was performed for one selected ground-motion descriptor set. The corresponding building configurations were used to initialise and evaluate the NSGA-II search through the trained MIMO-FNN. The resulting candidates are therefore conditional on the selected seismic-input scenario and are not claimed to represent generally resilient designs.

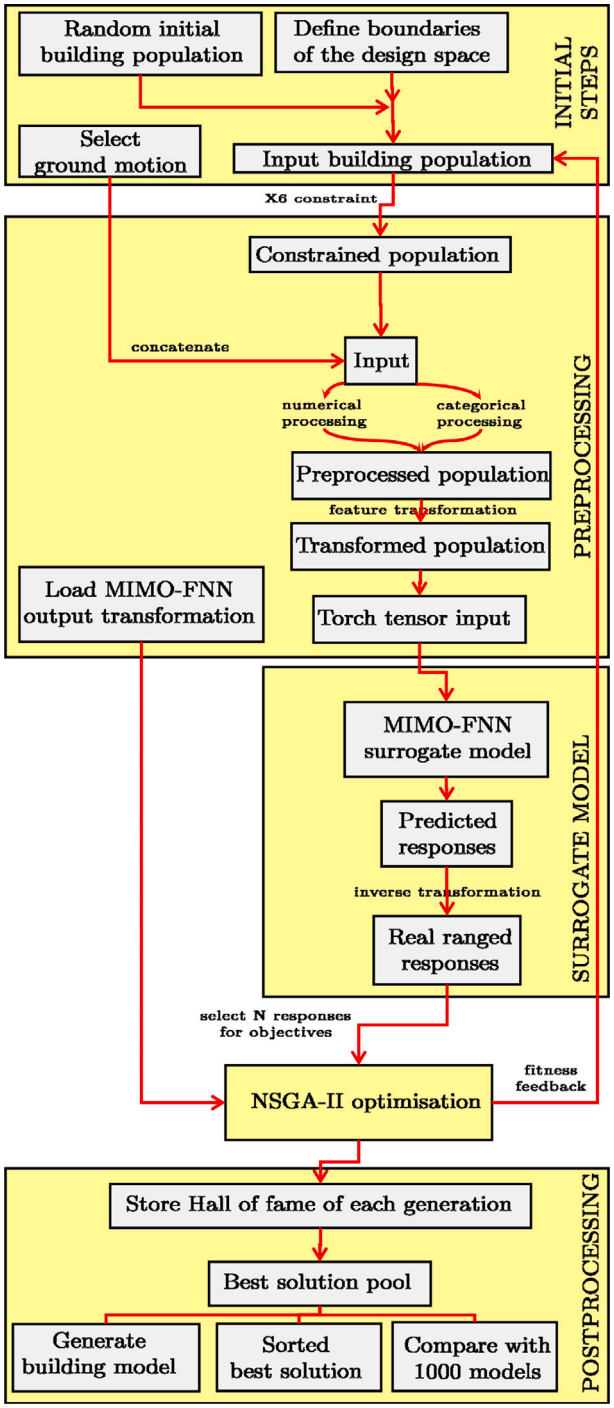


Fig. 4. Flowchart of the entire optimisation procedure.

A more complete seismic-design workflow would repeat the optimisation and verification under multiple unseen ground-motion records. The resulting optimised solutions should therefore be interpreted as preliminary candidate geometries identified under the selected seismic-input scenario. Before practical adoption, the shortlisted candidates should be re-analysed using high-fidelity FE models under additional ground-motion records not used during surrogate training.

The optimisation process was guided by the 10 building variables X1 to X10, with each generated population constrained within predefined boundaries. Design variable boundaries were subsequently defined to

specify allowable ranges or discrete options. As the default mutation function of NSGA-II may occasionally yield out-of-bound solutions, a custom mutation function was implemented to enforce compliance with the discrete parameter ranges. The crossover was then applied to exchange segments of design variables between individuals, enabling exploration of the solution space and preserving population diversity. Both mutation and crossover were employed to enhance optimisation effectiveness through broader search capability.

A specific ground motion descriptor set corresponding to one of the seven available records was selected and the corresponding pre-trained preprocessors and normalisers for the MIMO-FNN were loaded. This ensured consistent data processing and appropriate scaling and encoding of inputs, which are essential for accurate NN predictions. Each individual fitness was evaluated and the processed features were then fed into the MIMO-FNN model to predict the 91 selected quantities, from which the five selected responses were retained as the optimisation objectives and used to evaluate fitness.

The built-in Hall of Fame function was employed to track top-performing individuals across generations, with duplicates removed to ensure solution uniqueness. Approximately 2000 candidate evaluations were generated over 10 generations, resulting in 408 unique candidate designs after duplicate removal. The NSGA-II search was carried out as a multi-objective procedure using the selected response quantities as separate objectives. After the Pareto-based search, each selected response was normalised and an average normalised score was computed only as a post-processing indicator for ranking, comparing and visualising the generated candidate designs. This scalar indicator was not used to replace the multi-objective nature of NSGA-II. It provides a transparent summary index for inspecting the candidate geometries, while the final selection of a design may still depend on alternative stakeholder-defined weights or on specific architectural, structural and sustainability priorities. The solutions were then sorted accordingly, to generate 408 tall building models and visualise the optimised designs via Design Explorer (Tomasetti, 2025), a web-based tool for exploring multidimensional design spaces.

5.1. Computational cost and time-saving benefits of the surrogate workflow

A primary motivation for introducing the MIMO-FNN surrogate is the reduction of computational overhead during design-space exploration. It is important to distinguish between the one-time cost of generating the training database and the marginal cost of evaluating new candidate designs during optimisation. The proposed workflow does not eliminate the initial FE simulation cost required to construct a reliable training database. In the present study, this database was generated from 1000 parametric building configurations, each analysed under seven bidirectional ground-motion records, resulting in 7000 OpenSeesPy time-history simulations. This initial database-generation stage represents the main offline computational investment of the framework.

The computational advantage appears in the subsequent optimisation stage. In a conventional FE-based optimisation loop, each candidate design generated by NSGA-II would require model reconstruction, FE time-history analysis and post-processing of the response quantities. If N_{cand} denotes the number of candidate designs evaluated during optimisation and N_{gm} denotes the number of ground-motion records, the approximate FE-based computational demand can be expressed as

$$C_{\mathrm{FE,opt}} \approx N_{\mathrm{cand}} N_{\mathrm{gm}} T_{\mathrm{FE}}, \tag{13}$$

where T_{FE} is the average computational time required for one FE time-history simulation and its associated post-processing. In contrast, the surrogate-assisted optimisation replaces each FE solve by a forward pass of the trained MIMO-FNN, so that the corresponding marginal computational demand is

$$C_{\mathrm{sur,opt}} \approx N_{\mathrm{cand}} T_{\mathrm{NN}}, \tag{14}$$

where T_{NN} is the inference time of the neural network for one candidate. Since T_{NN} is several orders of magnitude smaller than the time needed to generate, run and post-process an FE time-history model, the trained surrogate makes repeated design exploration computationally practical.

In the optimisation example considered in this study, approximately 2000 Pareto-optimal candidate evaluations were produced over 10 NSGA-II generations, from which 408 unique candidate designs were retained after duplicate removal. If these approximately 2000 candidate evaluations had been performed directly using the selected FE ground-motion case, approximately 2000 additional FE time-history analyses would have been required. If the same optimisation stage were repeated for all seven ground-motion records, the corresponding FE burden would increase to approximately 14,000 additional time-history analyses. In the proposed workflow, these repeated FE calls are replaced by MIMO-FNN predictions of the 91 response quantities. The comparison is summarised in Table 10.

This comparison shows that the surrogate should be understood as an amortised computational tool. The initial FE database remains necessary to provide physically meaningful training data, but once the MIMO-FNN is trained, the marginal cost of evaluating additional design candidates is drastically reduced. Therefore, the proposed framework is particularly suitable for early-stage design, where many alternative geometries must be screened rapidly before a smaller number of promising candidates are selected for subsequent high-fidelity FE verification.

5.2. Surrogate-error propagation in the optimisation score

Because the NSGA-II search is driven by MIMO-FNN predictions, approximation errors in the surrogate can affect the ranking of the generated candidates. To clarify this point, the optimisation score used for ranking can be interpreted as a function of the predicted responses. Let

Table 10
Quantitative comparison of the FE-based and surrogate-assisted computational workflows.

Stage	Traditional FE-based workflow	Surrogate-assisted workflow	Computational implication
Training/database generation	1000 building models × 7 ground-motion records = 7000 FE time-history simulations	Same 7000 FE simulations used once to construct the supervised training database	No saving at the initial offline data-generation stage
Model training	Not applicable	MIMO-FNN training for 100 epochs using the 7000-sample database	One-time offline cost after FE database generation
Candidate evaluation during NSGA-II	One new FE model generation, dynamic analysis and post-processing step for each candidate design	One MIMO-FNN forward pass for each candidate design	FE analyses are replaced by neural-network inference
Optimisation example in this study	Approximately 2000 additional FE analyses for the selected ground-motion case; up to approximately 14,000 FE analyses if repeated for all seven records	Approximately 2000 MIMO-FNN forward evaluations, followed by duplicate removal to obtain 408 unique candidates	Large reduction in marginal computational cost during design exploration
Predicted response quantities	Responses obtained after FE analysis and post-processing	91 structural, environmental and geometry-dependent responses predicted simultaneously	Rapid multi-response screening of candidate geometries

$\hat{y}_{ik}$ denote the MIMO-FNN prediction of the kth selected response for the ith candidate design, and let y_{ik} denote the corresponding value that would be obtained from a direct FE analysis. The prediction error is therefore

$$e_{ik} = \hat{y}_{ik} - y_{ik}. \tag{15}$$

For the five selected optimisation responses, the normalised predicted response can be expressed as

$$\hat{z}_{ik} = \frac{\hat{y}_{ik} - y_k^{\min}}{y_k^{\max} - y_k^{\min}}, \tag{16}$$

where $y_k^{\min}$ and $y_k^{\max}$ are the minimum and maximum reference values used for normalisation of the kth response. The average normalised response used for ranking the candidates is then

$$\hat{J}_i = \frac{1}{q}\sum_{k=1}^{q} \hat{z}_{ik}, \qquad q = 5. \tag{17}$$

If the same candidate were re-analysed by a high-fidelity FE model, the corresponding FE-based score would be

$$J_i = \frac{1}{q}\sum_{k=1}^{q} \frac{y_{ik} - y_k^{\min}}{y_k^{\max} - y_k^{\min}}. \tag{18}$$

The difference between the surrogate-based and FE-based ranking scores is therefore

$$\Delta J_i = \hat{J}_i - J_i = \frac{1}{q}\sum_{k=1}^{q} \frac{e_{ik}}{y_k^{\max} - y_k^{\min}}. \tag{19}$$

No additional FE time-history analyses were performed for the final NSGA-II-generated candidates in the present study. The optimisation results should therefore be interpreted as surrogate-screened candidate designs, not as independently FE-verified optimum solutions. The predictive capability of the MIMO-FNN was assessed using cross-validation and held-out test data defined at the building-configuration level, so that cases associated with the same building configuration were not simultaneously used for training and testing. This procedure supports the use of the surrogate for early-stage screening within the sampled design and ground-motion domain. However, final verification of the optimised candidates would require reconstructing the selected geometries and performing new high-fidelity FE time-history analyses, preferably under additional ground-motion records and nonlinear modelling assumptions.

This relation shows that the influence of the surrogate error on the final ranking depends not only on the absolute prediction error of each response, but also on the response range used for normalisation. Responses with small normalisation ranges are more sensitive to prediction error, whereas responses with large ranges are less affected by the same absolute error. Accordingly, the optimised solutions reported below should be interpreted as candidates whose ranking is conditional on the predictive accuracy of the surrogate within the sampled design space.

In the present study, the internal predictive consistency of the MIMO-FNN was assessed through training-validation curves and cross-validation. These results support the use of the surrogate for rapid screening of alternatives within the available database. Nevertheless, the error-propagation relation above makes it clear that the NSGA-II outputs are not substitutes for direct FE verification. A rigorous verification test of the final candidates would require reconstructing the selected optimised geometries, performing new high-fidelity FE time-history analyses and comparing the FE-computed responses with the MIMO-FNN predictions. Such a verification step is necessary before practical design adoption and is identified here as a required continuation of the workflow. The surrogate-predicted responses of the optimised solutions were subsequently compared to the responses of the previously generated samples. For consistency, the five selected responses were extracted from both sets and the MIMO-FNN outputs were interpreted in the same physical response space after inverse transformation. Statistical indicators, namely mean, median, minimum and maximum, were computed for both (O) optimised and (S) initially generated sets, as shown in Table 11. The comparison between the O and S solutions therefore offers insight into the effectiveness of the surrogate-assisted screening process, while not representing an independent FE verification of the final optimised candidates.

It is worth noting that the optimisation does not reduce every response indicator for every individual candidate. Since NSGA-II was used in a multi-objective setting, the retained candidates represent trade-off solutions among acceleration, displacement, inter-storey drift, von Mises stress and total mass. For this reason, some individual solutions may exhibit a larger value of one response while still being retained because of improvements in other objectives. In Table 11, the mean and median values of Overall_Max_Acc decrease from 0.865 to 0.742 m/s^2 and from 0.891 to 0.710 m/s^2, respectively, indicating an overall reduction in acceleration demand within the optimised population. However, the maximum value increases from 1.324 to 1.449 m/s^2, because no explicit hard upper-bound constraint was imposed on acceleration during the surrogate-assisted search. Therefore, the optimised set should be interpreted as a population of trade-off candidates rather than as a set in which every response is monotonically improved for every design.

To assess the practical significance of the optimisation stage, the initially generated 1000 design samples were used as a non-optimised baseline population. This baseline represents the response distribution obtained from admissible parametric design-space sampling without surrogate-assisted evolutionary search. The optimised population was then compared with this baseline using the five selected objective responses. This comparison does not prove superiority over every possible rule-based engineering design strategy, but it provides a transparent reference for evaluating whether the NSGA-II search shifts the candidate population toward more favourable regions of the explored design space.

Based on the mean values reported in Table 11, the optimised population reduced the mean overall maximum acceleration from 0.865 to 0.742 m/s^2; the mean maximum displacement from 78.1 to 72.5 mm; the mean overall maximum drift from 2.2 to 2.0 mm; the mean maximum von Mises stress from 128 to 115 MPa; and the mean total mass from 22,100 to 19,010 tons. These changes correspond approximately to reductions of 14.2%, 7.2%, 9.1%, 10.2% and 14.0%, respectively. The simultaneous reduction of both response-demand indicators and total mass suggests that the surrogate-assisted optimisation can identify favourable trade-off regions rather than merely improving one response at the expense of all others.

Table 11
Comparison between design solutions generated (O) by the optimisation process and (S) those from the initial 1000 design samples.

Selected response for objective function	O-mean	S-mean	O-median	S-median	O-min	S-min	O-max	S-max
Overall_Max_Acc (m/s^2)	0.742	0.865	0.710	0.891	0.294	0.551	1.449	1.324
Max_Displacement (mm)	72.5	78.1	70.1	76.5	41.5	26.1	116	192
Overall_Max_Drift (mm)	2.0	2.2	2.0	2.2	1.4	0.9	2.5	3.3
Max_Von_Mises_tot (MPa)	115	128	115	127	98	77	148	517
TotalMass (tons)	19,010	22,100	18,640	20,500	7613	6290	37,390	58,500

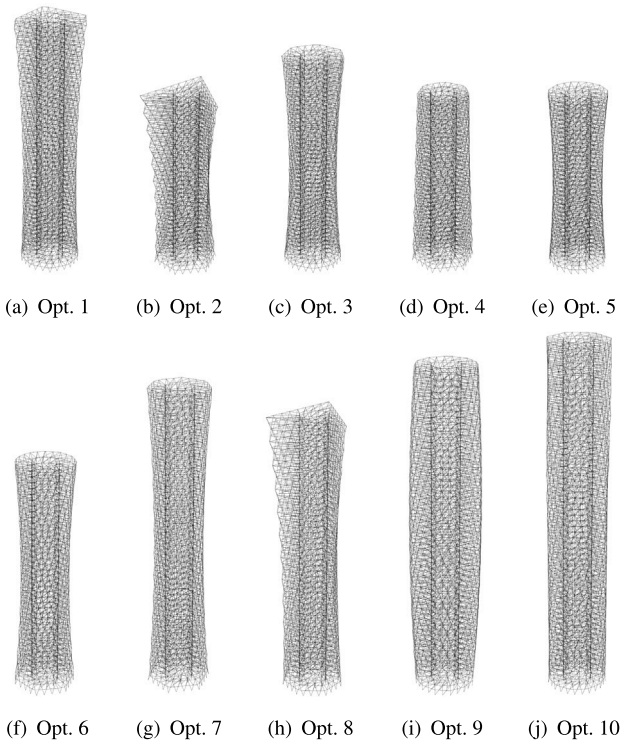


Fig. 5. Visualisation of 10 optimised models (labelled Opt. 1–10), whose specifications are in Table 12.

Nevertheless, these improvements should be interpreted cautiously. The comparison is made against the sampled baseline available in the present database, not against all possible simplified design rules or engineering heuristics. Defining a universal rule-based baseline for the considered class of diagrid tall buildings is non-trivial because the design space includes different polygonal plans, storey numbers, inter-storey heights and mutually exclusive vertical transformations such as tapering, twisting and curvilinearity. For this reason, the present comparison is used to demonstrate the practical usefulness of the optimisation as a design-space reduction mechanism. A more exhaustive practical validation would require comparison with additional rule-based design strategies and direct FE verification of selected optimised candidates.

The optimised solutions and their corresponding responses were imported into Grasshopper™, where a script generated the associated tall building models. Data visualisation supported the comparison with real samples, facilitating a clear presentation of the optimisation outcomes. This process ensured effective documentation and communication of results. Fig. 5 illustrates 10 optimised solutions, with their characteristics summarised in Table 12. From a practical engineering standpoint, the optimised solutions in Table 12 and Fig. 5 should be interpreted as a reduced set of promising design alternatives rather than as final designs ready for direct implementation. Their main role is to support early-stage design-space reduction by identifying candidate geometries that offer favourable trade-offs among acceleration, displacement, inter-storey drift, von Mises stress and total mass. In a practical workflow, these candidates can be used to guide discussions between architects and structural engineers, prioritise a limited number of configurations for detailed study, and reduce the number of expensive FE analyses required in subsequent design stages.

Before practical adoption, however, the shortlisted candidates should be re-analysed using high-fidelity FE models and checked against the relevant design-code requirements, serviceability limits, strength criteria, connection behaviour and constructability constraints. The present optimisation stage therefore acts as a computational decision-support layer placed before detailed engineering verification. This interpretation is consistent with the early-stage scope of the study and avoids treating surrogate-predicted optimum candidates as final certified structural designs.

5.3. Design optimised solutions

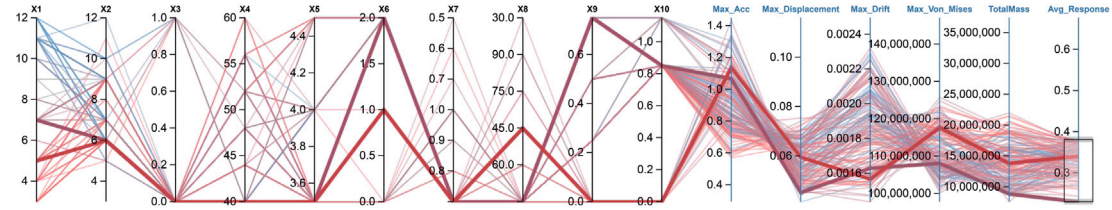
Fig. 6(a) (top) presents a parallel-coordinates plot of the design features (inputs X1 to X10) and the selected responses (outputs). The final response shown is the average normalised response, which serves as a structural performance indicator. This visualisation enables a simple examination of the value ranges across both input and output parameters.

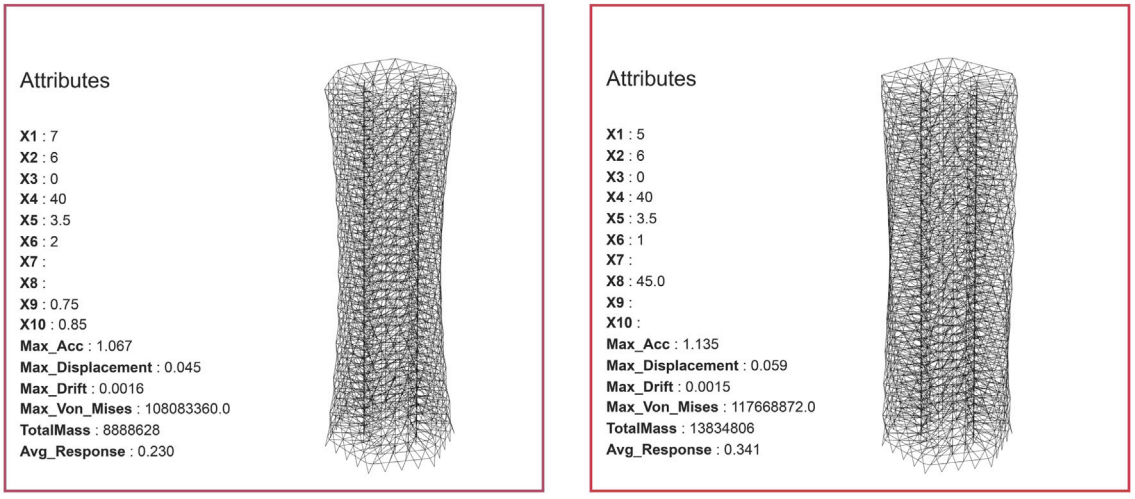
The optimised solutions are categorised into three subsets based on percentiles of the average normalised response: 0 to 33.33, 33.33 to 66.66 and 66.66 to 100 percentiles. Each subset contains two representative models, altogether corresponding to percentiles 0, 20, 40, 60, 80, and 100 across the entire dataset. Feature and response ranges and their representative models are illustrated in Figs. 6(a), 7(a) and 8(a), while feature and response distributions are depicted in Figs. 6(b), 7(b) and 8(b).

In the distribution plots of Figs. 6(b)–8(b), the bars indicate the number of optimised candidates associated with each admissible value or interval of a given design variable or response quantity, whereas the superimposed lines provide only a smoothed visual indication of the distribution tendency. For discrete variables with a small number of admissible values, the empty spaces between bars do not correspond to unselected or non-optimal intermediate values; rather, they indicate values that are not part of the predefined design space and were therefore not tested. These plots should consequently be interpreted as frequency distributions over the admissible tested values, not as continuous design-variable distributions. The first subset (0-33.33%) is depicted in Fig. 6(a). The bottom-left picture displays the

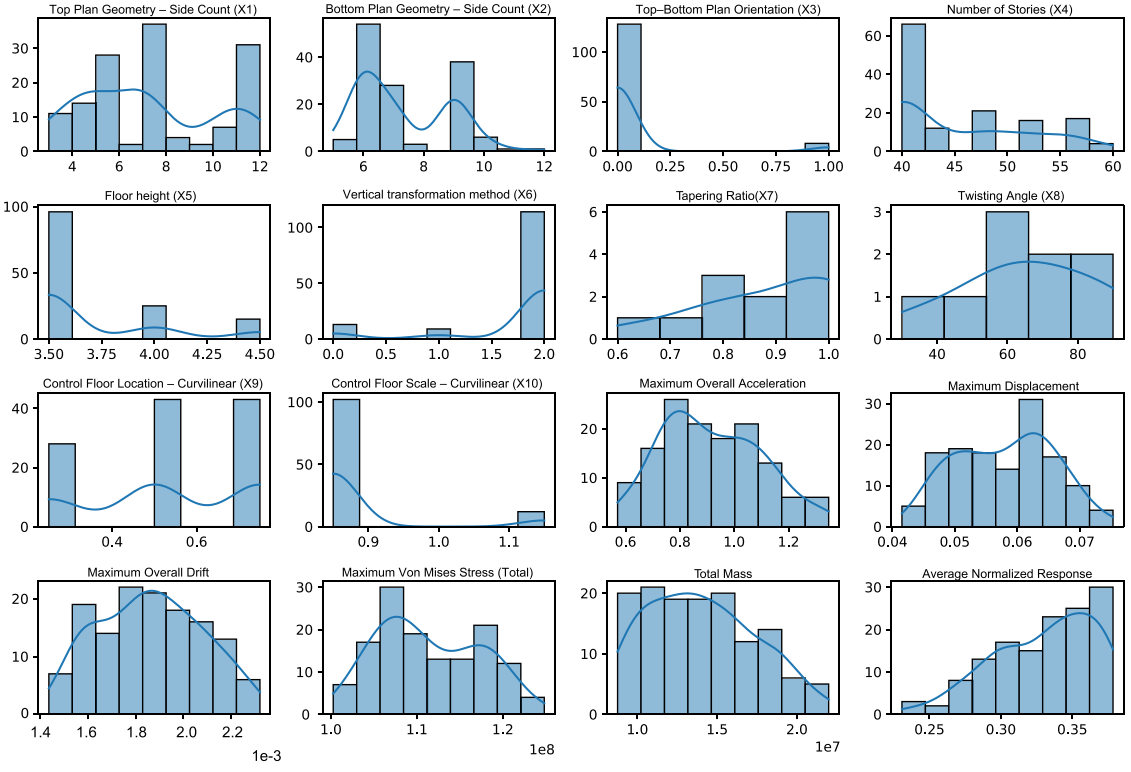
Table 12
Building features and five selected responses for the 10 optimised models. The last column shows the average normalised score of all five responses. The entry "–" is used for a non-context applicable feature.

Model	X1	X2	X3	X4	X5	X6	X7	X8	X9	X10	Max acc (m/s^2)	Max displ. (mm)	Max i.s. drift (mm)	Max von Mises (MPa)	Total mass (tons)	Average normalised response
Opt_1	7	8	0	52	4.5	2	–	–	0.75	0.85	0.600	74.1	2.07	111	17,000	0.377
Opt_2	5	6	0	60	3.5	0	0.8	–	–	–	0.763	70.7	1.79	122	23,000	0.425
Opt_3	5	9	0	52	4.5	2	–	–	0.75	0.85	0.793	62.6	2.26	104	15,300	0.373
Opt_4	4	7	0	52	3.5	2	–	–	0.50	0.85	0.716	60.1	1.91	108	15,300	0.303
Opt_5	3	6	0	44	4.0	2	–	–	0.75	0.85	0.718	64.0	1.83	114	14,800	0.322
Opt_6	3	10	0	60	4.5	0	0.9	–	–	–	0.423	93.3	2.07	123	30,000	0.535
Opt_7	5	7	0	40	3.5	2	–	–	0.75	0.85	0.954	52.3	1.90	106	11,300	0.290
Opt_8	11	9	1	52	4.0	2	–	–	0.50	0.85	0.949	62.5	2.18	105	14,200	0.384
Opt_9	8	7	0	48	4.0	2	–	–	0.75	0.85	0.743	66.9	2.17	102	14,300	0.344
Opt_10	7	6	0	60	4.5	1	–	90.0	–	–	0.356	106.0	2.16	117	32,900	0.567



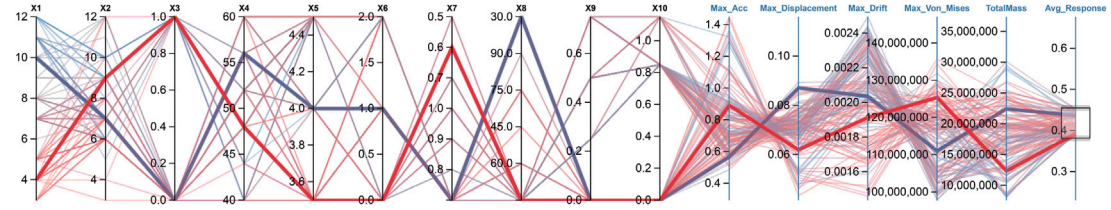


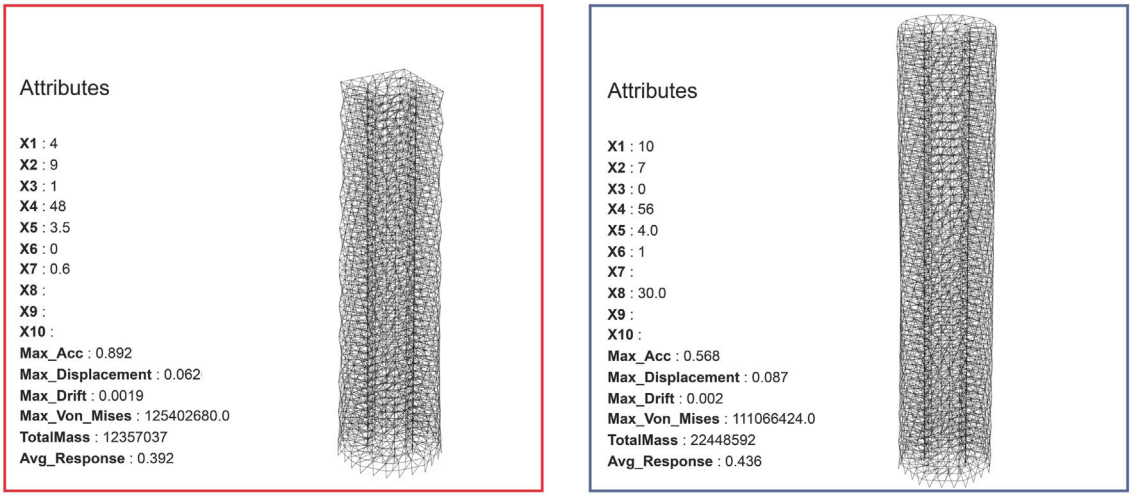
(a)



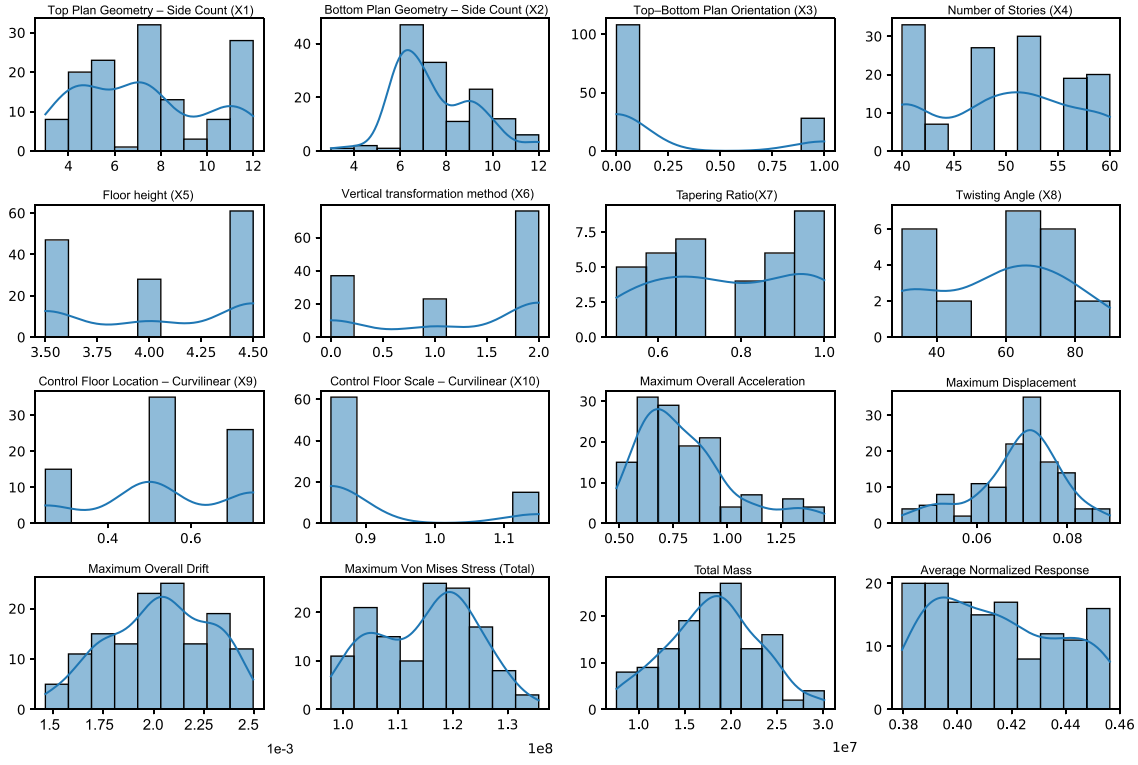
(b)

Fig. 6. First subset of the optimised models (0–33.33%): (a) visualisation and (b) distributions of features and responses. In (b), bars show counts and lines show smoothed distribution trends; empty spaces for discrete variables indicate non-admissible values that were not tested.



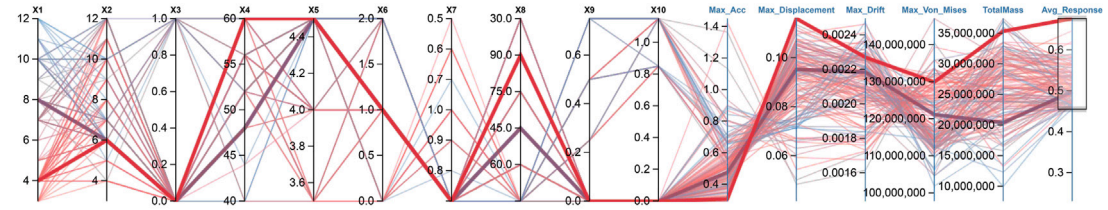


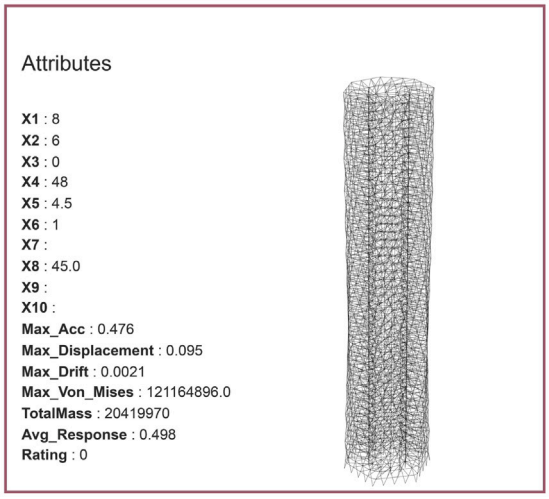
(a)

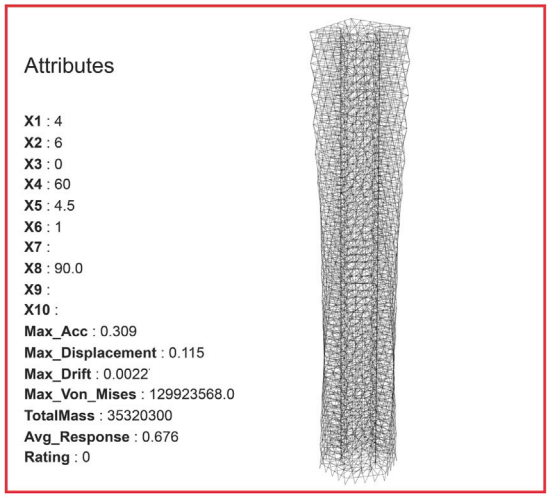


(b)

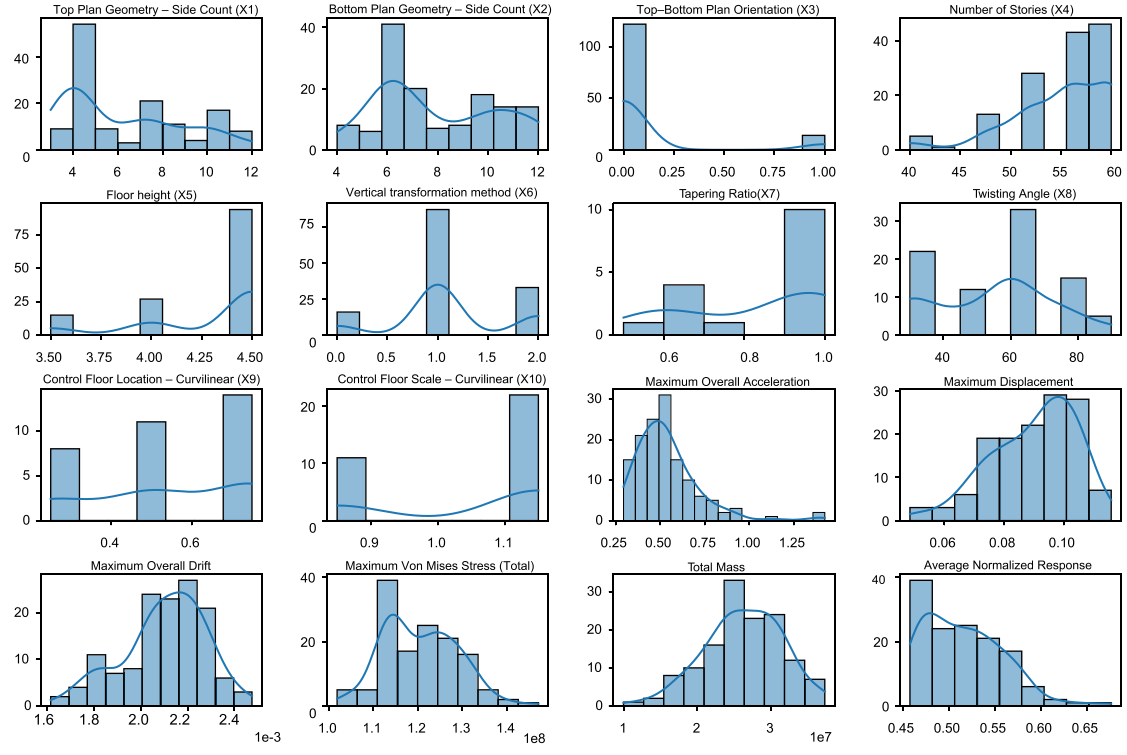
Fig. 7. Second subset of the optimised models (33.33–66.66%): (a) visualisation and (b) distributions of features and responses. In (b), bars show counts and lines show smoothed distribution trends; empty spaces for discrete variables indicate non-admissible values that were not tested.







(a)



(b)

Fig. 8. Third subset of the optimised models (66.66–100%): (a) visualisation and (b) distributions of features and responses. In (b), bars show counts and lines show smoothed distribution trends; empty spaces for discrete variables indicate non-admissible values that were not tested.

best-performing structural model, with the lowest average normalised response of 0.23. Its features and responses are highlighted in purple above. This model, representing the 0th percentile, has 40 floors with a 3.5-meter floor-to-floor height. It adopts a curvilinear transformation, with the control floor located at 75% of the total height and a floor scale factor of 0.85. In contrast, the bottom-right panel shows the model in the 20th percentile, with an average normalised response of 0.341. Highlighted in red, this model also features 40 floors and a 3.5-meter inter-storey height but uses a twisted transformation with a 45-degree rotation angle. According to the distributions reported in Fig. 6(b), in this subset (0-33.33%) X1 and X2 (number of sides in the top/bottom plan) vary, while most models have X3 set to 0, indicating a corner-to-corner orientation. The majority feature 40 floors with a 3.5-meter floor-to-floor height. Nearly all models are curvilinear, with few tapered or twisted alternatives. Most curvilinear designs place the control floor in the lower half, typically with a concave shape (scale of 0.85).

The second subset (33.33–66.66%) is depicted in Fig. 7(a). The bottom-left picture displays the model at the 40th percentile of the average normalised response (0.392), with red-highlighted features above. It has 48 floors, a 3.5-meter floor-to-floor height and a tapered form—its top floor covers 60% of the bottom plan area. The bottom-right picture shows the 60th percentile model (0.436), highlighted in blue. It features 56 floors, a 4-meter floor-to-floor height and a 30-degree twisting angle. According to the distributions shown in Fig. 7(b), in this subset X1 and X2 vary, with six-sided bottom plans being most common. Most models have X3 set to 0, indicating a corner-to-corner orientation. Buildings with 40, 48 and 52 floors dominate, typically featuring a 4.5-meter floor height. Curvilinear forms are most prevalent, followed by tapered and twisted shapes. Among tapered models, non-prismatic forms are frequent, while twisted ones typically have a 60-degree angle. Most curvilinear models place the control floor in the lower half and are concave, with the control floor scale of 0.85.

The third subset (66.66-100%) is shown in Fig. 8(a). In the bottom-left picture, the model at the 80th percentile of the average normalised response (0.498) is shown, with features and responses highlighted in purple. It has 48 floors, a 4.5-meter floor height and a 45-degree twist. The bottom-right picture displays the 100th percentile model—the least optimal among the optimised set—with an average normalised response of 0.676. Highlighted in red, this model features the maximum number of floors (60), the tallest floor height (4.5 meters) and a 90-degree twist. According to Fig. 8(b), X1 commonly corresponds to 4-sided polygons, while X2 often features 5-sided ones. X3 is mostly 0, indicating a corner-to-corner orientation. This subset includes taller buildings—typically 56 or 60 stories—with a floor height of 4.5 meters. Twisted models dominate, followed by curvilinear and tapered forms. Non-prismatic shapes are frequent in tapered models and a 60-degree twist is most common. In curvilinear designs, the control floor is generally in the bottom half and often convex, with a scale of 1.15.

6. Conclusion

This study proposes a surrogate-assisted computational workflow for early-stage multi-objective seismic design optimisation of diagrid tall buildings. The optimisation addresses design-space exploration rather than vibration control: the design variables are the architectural and geometric descriptors of the diagrid configurations, while the selected objectives are related to acceleration, displacement, inter-storey drift, von Mises stress and total mass. The main contribution is therefore not the isolated use of finite-element analysis, feed-forward neural networks or NSGA-II, but their integration into a reproducible workflow for rapidly screening high-dimensional diagrid design alternatives under seismic actions.

The computational framework combined automated parametric modelling, OpenSeesPy time-history simulation, conditional data preprocessing, MIMO-FNN surrogate modelling and NSGA-II-based candidate generation. A database of 1000 diagrid tall-building configurations was generated and analysed under seven bidirectional ground-motion records, resulting in 7000 time-history simulations. The original 22 input variables were expanded to 45 preprocessed input columns through numerical scaling, categorical encoding and conditional masking of transformation-dependent geometric variables. This preprocessing was necessary to preserve the physical meaning of conditional parameters associated with tapered, twisted and curvilinear geometries.

The adopted MIMO-FNN mapped the 45-dimensional preprocessed input space to 91 output responses, including 69 structural responses, 14 environmental and cost-related indicators, and 8 geometry-dependent quantities. The surrogate used a pyramidal architecture with four hidden layers of 512, 256, 128 and 128 neurons, ReLU activation functions, dropout regularisation and linear output layers for continuous multi-output regression. The model was trained for 100 epochs with a learning rate of 2×10^{-5}, a batch size of 32, a weight-decay coefficient of 1×10^{-4} and up to 5-fold cross-validation. The training and validation trends indicated stable predictive behaviour within the considered design and ground-motion space, without evidence of overfitting.

After training, the surrogate was coupled with NSGA-II to perform rapid multi-objective design exploration. Approximately 2000 optimisation candidates were evaluated over 10 generations and then reduced to 408 unique candidate geometries after duplicate removal. In a conventional FE-based optimisation loop, these evaluations would have required approximately 2000 additional time-history analyses for the selected ground-motion case, or up to approximately 14,000 additional analyses if repeated for all seven records. In the proposed workflow, these repeated FE calls were replaced by MIMO-FNN forward evaluations, allowing rapid prediction of all 91 response quantities. This demonstrates the practical value of the surrogate as an amortised computational tool: the initial FE database remains necessary, but the marginal cost of subsequent design exploration is substantially reduced.

The optimised candidates showed improved average behaviour relative to the initially generated samples for several selected responses. The mean overall maximum acceleration decreased from 0.865 to 0.742 m/s^2, the mean maximum displacement from 78.1 to 72.5 mm, the mean inter-storey drift from 2.2 to 2.0 mm, the mean von Mises stress from 128 to 115 MPa and the mean total mass from 22,100 to 19,010 tons. These changes correspond approximately to reductions of 14.2%, 7.2%, 9.1%, 10.2% and 14.0%, respectively. These results indicate that the surrogate-assisted optimisation can identify favourable regions of the design space by simultaneously considering structural response and material efficiency.

These findings should nevertheless be interpreted within the scope of the study. The proposed framework is intended as a preliminary screening and decision-support tool for early-stage design, especially when many candidate geometries must be explored quickly under moderate seismic demand. Its predictions and optimisation outputs are not intended to replace nonlinear FE analyses, code-based checks or final engineering verification. The optimised candidates reported in this paper were ranked using surrogate-predicted responses; therefore, their final acceptance requires direct FE re-analysis to quantify the discrepancy between MIMO-FNN predictions and FE-computed responses.

A further limitation concerns the number of ground-motion records. Although the seven-record suite enabled the construction of a computationally demanding database of 7000 time-history simulations, it cannot fully represent the high stochastic variability of seismic excitation. The results should therefore be interpreted as demonstrating stable surrogate behaviour within the considered record set, not as evidence of universal seismic robustness. Future developments should extend the database using additional PEER NGA-West2 records, broader intensity levels, nonlinear structural models, physics-aware constraints, direct FE verification of optimised candidates and comparison with simpler rule-based or baseline design strategies.

Declaration of competing interest

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

Appendix A. Supplementary data

Supplementary data for this article can be found online at doi:10.1016/j.cacaie.2026.100181.

References

Adeli, H., & Kumar, S. (1995). Concurrent structural optimization on massively parallel supercomputer. *Journal of Structural Engineering*, *121*(11), 1588–1597.

Adeli, H., & Park, H. S. (1996). Fully automated design of super-high-rise building structures by a hybrid AI model on a massively parallel machine. *AI Magazine*, *17*(3), 87–93.

Asadi, E., & Adeli, H. (2017). Diagrid: An innovative, sustainable and efficient structural system. *The Structural Design of Tall and Special Buildings*, *26*(8), e1358.

Asadi, E., & Adeli, H. (2018). Nonlinear behavior and design of mid-to-highrise diagrid structures in seismic regions. *Engineering Journal, AISC*, *55*, 161–180.

Bailey, E. T., & Caldas, L. (2023). Operative generative design using non-dominated sorting genetic algorithm II (NSGA-II). *Automation in Construction*, *155*, 105026.

Balmer, V., Kuhn, S. V., Bischof, R., Salamanca, L., Kaufmann, W., Perez-Cruz, F., & Kraus, M. A. (2024). Design space exploration and explanation via conditional variational autoencoders in meta-model-based conceptual design of pedestrian bridges. *Automation in Construction*, *163*, 105411.

Beltrame Group. (2020). *Environmental product declaration*. Beltrame Group, Technical Report.

Bhatta, S., & Dang, J. (2024). Multiclass seismic damage detection of buildings using quantum convolutional neural network. *Computer-Aided Civil and Infrastructure Engineering*, *39*(3), 406–423.

Brunton, S. L., & Kutz, J. N. (2024). *Data-driven science and engineering. Machine learning, dinamical systems and control* (Second. ed.). Cambridge University Press.

Chen, Z.-H., Entezami, A., Sarmadi, H., Xu, Z.-D., Gao, W., Liu, L., & Behkamal, B. (2026). Transformer-based contrastive denoising autoencoder for unsupervised detection of wind-induced vibration changes in cable-supported bridges. *Computer-Aided Civil and Infrastructure Engineering*, 100114.

Chen, Z.-H., Xu, Z.-D., Entezami, A., Sarmadi, H., Gao, W., & Behkamal, B. (2026). Dynamic monitoring of bridges subjected to freezing-induced variability by kernel-enhanced deep anomaly detection. *Engineering Structures*, *362*, 122971.

Danhaive, R., & Mueller, C. T. (2021). Design subspace learning: Structural design space exploration using performance-conditioned generative modeling. *Automation in Construction*, *127*, 103664.

EMR. (2022). *Reusable steel environmental product declaration programme*. EMR, Technical Report.

Entezami, A., Kiarad, H., Sarmadi, H., & Behkamal, B. (2025). Dynamic response prediction of heritage masonry buildings by dual-phase hybrid regression modeling under partial monitoring parameters. *Developments in the Built Environment*, *24*, 100802.

Entezami, A., & Sarmadi, H. (2025). Machine learning-aided prediction of windstorm-induced vibration responses of long-span suspension bridges. *Computer-Aided Civil and Infrastructure Engineering*, *40*(8), 1043–1060.

Entezami, A., Sarmadi, H., Kiarad, H., & Behkamal, B. (2026). Residual-informed deep meta regressor for predicting modal frequencies of historical masonry buildings under environmental and operational data scarcity. *Mechanical Systems and Signal Processing*, *242*, 113675.

Gauch, H. L., Hawkins, W., Ibell, T., Allwood, J. M., & Dunant, C. F. (2022). Carbon vs. cost option mapping: A tool for improving early-stage design decisions. *Automation in Construction*, *136*, 104178.

Gibbons, O. P., Orr, J. J., Archer-Jones, C., Arnold, W., & Green, D. (2022). *How to calculate embodied carbon*. Institution of Structural Engineers (ISTRUCTE).

Granadeiro, V., Duarte, J. P., Correia, J. R., & Leal, V. M. S. (2013). Building envelope shape design in early stages of the design process: Integrating architectural design systems and energy simulation. *Automation in Construction*, *32*, 196–209.

Guner, T., Bursi, O. S., & Erlicher, S. (2023). Optimization and performance of metafoundations for seismic isolation of small modular reactors. *Computer-Aided Civil and Infrastructure Engineering*, *38*(12), 1558–1582.

Hanson UK. (2017). Environmental product declaration. *Environmental Product Declaration Series*, *44*(456).

Ilboudo, W. E. L., Kobayashi, T., & Sugimoto, K. (2020). *Tadam: A robust stochastic gradient optimizer*. arXiv preprint arXiv:2003.00179.

Janisch, A. (2023). *Kostenfaktoren im stahlbau - was kostet ein kg stahlbau 2023?* https://jactio.com/en/kostenfaktoren-im-stahlbau/.

Kanyilmaz, A., Dang, V. H., & Kondratenko, A. (2023). How does conceptual design impact the cost and carbon footprint of structures? *Structures*, *58*, 105102.

Kazemi, P., Entezami, A., & Ghisi, A. (2024b). Machine learning techniques for diagrid building design: Architectural–structural correlations with feature selection and data augmentation. *Journal of Building Engineering*, *86*, 108766.

Kazemi, P., Ghisi, A., & Entezami, A. (2024a). Artificial intelligence-powered computational strategies in selecting and augmenting data for early design of tall buildings with outer diagrids. *Buildings*, *14*(4), 1118.

Kazemi, P., Turrin, M., Andriotis, C., Entezami, A., Mariani, S., & Ghisi, A. (2025). Integrating dynamic seismic simulation in architectural design: A parametric approach for tall buildings with outer diagrids. In *Computation and earthquake disaster management: A perspective in built environment education*. TU Delft Open Publisher.

Lydia, A., & Francis, S. (2019). Adagrad - an optimizer for stochastic gradient descent. *International Journal of Information and Computing Science*, *6*(3), 566–568.

McKenna, F., Scott, M. H., & Fenves, G. L. (2010). Nonlinear finite-element analysis software architecture using object composition. *Journal of Computing in Civil Engineering*, *24*, 95–107.

Nabian, M. A., & Meidani, H. (2018). Deep learning for accelerated seismic reliability analysis of transportation networks. *Computer-Aided Civil and Infrastructure Engineering*, *33*(6), 443–458.

Negendahl, K. (2015). Building performance simulation in the early design stage: An introduction to integrated dynamic models. *Automation in Construction*, *54*, 39–53.

Ning, C., & Xie, Y. (2024). Convolutional variational autoencoder for ground motion classification and generation toward efficient seismic fragility assessment. *Computer-Aided Civil and Infrastructure Engineering*, *39*(2), 165–185.

Noureldin, M., Abuhmed, T., Saygi, M., & Kim, J. (2023). Explainable probabilistic deep learning framework for seismic assessment of structures using distribution-free prediction intervals. *Computer-Aided Civil and Infrastructure Engineering*, *38*(12), 1677–1698.

Pacific Earthquake Engineering Research Center. (2024). *PEER ground motion database*. https://ngawest2.berkeley.edu/.

Pan, W., Sun, Y., Turrin, M., Louter, C., & Sariyildiz, S. (2020). Design exploration of quantitative performance and geometry typology for indoor arena based on self-organizing map and multi-layered perceptron neural network. *Automation in Construction*, *114*, 103163.

Rete Ferroviaria Italiana. (2023). *Tariffa dei prezzi FA. opere civili e industriali. Edizione 2024. (In Italian)*. Murgia.

Shabbir, K., Noureldin, M., & Sim, S.-H. (2025). Data-driven model for seismic assessment, design, and retrofit of structures using explainable artificial intelligence. *Computer-Aided Civil and Infrastructure Engineering*, *40*(3), 281–300.

Sharma, D. K., Chatterjee, M., Kaur, G., & Vavilala, S. (2022). 3 - deep learning applications for disease diagnosis. In D. Gupta, U. Kose, A. Khanna, & V. E. Balas (Eds.), *Deep learning for medical applications with unique data* (pp. 31–51). Academic Press. https://www.sciencedirect.com/science/article/pii/B9780128241455000058.

Simões, L. M. C. (1996). Optimization of frames with semi–rigid connections. *Computers & Structures*, *60*(4), 531–539.

Singh, M. M., Deb, C., & Geyer, P. (2022). Early-stage design support combining machine learning and building information modelling. *Automation in Construction*, *136*, 104147.

Tibshirani, R., Hastie, T., & Friedman, J. (2010). Regularized paths for generalized linear models via coordinate descent. *Journal of Statistical Software*, *33*.

Tomasetti, T. (2025). *Design explorer 2*. https://tt-acm.github.io/DesignExplorer/.